

SGD-Net: Efficient Model-Based Deep Learning With Theoretical Guarantees

Jiaming Liu¹, *Graduate Student Member, IEEE*, Yu Sun¹, *Graduate Student Member, IEEE*, Weijie Gan, *Graduate Student Member, IEEE*, Xiaojian Xu¹, *Graduate Student Member, IEEE*, Brendt Wohlberg¹, *Senior Member, IEEE*, and Ulugbek S. Kamilov¹, *Senior Member, IEEE*

Abstract—Deep unfolding networks have recently gained popularity for solving imaging inverse problems. However, the computational and memory complexity of data-consistency layers within traditional deep unfolding networks scales with the number of measurements, limiting their applicability to large-scale imaging inverse problems. We propose SGD-Net as a new methodology for improving the efficiency of deep unfolding through stochastic approximations of the data-consistency layers. Our theoretical analysis shows that SGD-Net can be trained to approximate batch deep unfolding networks to an arbitrary precision. Our simulations on intensity diffraction tomography and sparse-view computed tomography show that SGD-Net can match the performance of the traditional batch network at a fraction of training and testing complexity.

Index Terms—Model-based deep learning, regularization by denoising, plug-and-play priors, stochastic optimization.

I. INTRODUCTION

THE recovery of an unknown image from a set of noisy measurements is a central problem in computational imaging. The recovery is traditionally formulated as an *inverse problem* that combines a physical-model characterizing the imaging system with a regularizer imposing a prior knowledge on the unknown image. Over the past years, many regularizers have been proposed as imaging priors, including those based on transform-domain sparsity, low-rank penalty, and dictionary learning [1]–[6].

There has been considerable recent interest in using *deep learning (DL)* in the context of imaging inverse problems (see recent reviews [7]–[9]). Instead of explicitly defining a regularizer,

Manuscript received January 22, 2021; revised May 18, 2021; accepted May 27, 2021. Date of publication June 1, 2021; date of current version June 21, 2021. This work was supported in part by NSF award CCF-1813910 and in part by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project 20200061DR. (Corresponding author: Ulugbek S. Kamilov.)

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).

Yu Sun, Weijie Gan, and Xiaojian Xu 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; weijie.gan@wustl.edu; xiaojianxu@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 and the Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis MO 63130 USA (e-mail: kamilov@ieee.org).

Digital Object Identifier 10.1109/TCI.2021.3085534

the traditional DL approach is based on training a *convolutional neural network (CNN)* architecture, such as U-Net [10], to invert the measurement operator by exploiting the natural redundancies in the imaging data [11]–[15]. *Plug-and-play priors (PnP)* [16] and *regularization by denoising (RED)* [17] are two well-known alternative approaches to the traditional DL that enable the integration of pre-trained CNN denoisers, such as DnCNN [18], as image priors within iterative algorithms. When equipped with advanced CNN denoisers, PnP/RED provides excellent performance by exploiting both the implicit prior, characterized by the denoiser, and the measurement model [19]–[27]. *Deep unfolding* is a related approach that interprets the iterations of an image recovery algorithm as layers of a CNN and trains it end-to-end in a supervised fashion (see “*Unrolling*” in [7] or “*Neural networks and analytical methods*” in [8]). Unlike in PnP/RED, the CNN in deep unfolding is trained jointly with the measurement model, leading to an image prior optimized for a given inverse problem [28]–[33]. Despite the recent popularity of deep unfolding, the training of such networks can be a significant practical challenge in applications that require processing of a large number of measurements. Specifically, the data-consistency layers of these recursive neural networks are based on *batch* processing, which means that they process the *entire set of measurements* at each layer. While this type of batch data processing is known to be suboptimal in traditional large-scale optimization [34]–[37], the issue has never been addressed in the context of designing deep unfolding networks.

We address this issue by proposing SGD-Net as the first deep unfolding methodology to adopt *stochastic processing* of measurements within data-consistency layers. This improves the efficiency of SGD-Net compared to its batch counterparts on large datasets during both training and testing. We implement SGD-Net by unfolding the gradient-based RED algorithm and introducing stochastic approximations to its data-consistency layers. The CNN within SGD-Net is trained in an end-to-end fashion to remove artifacts due to the imaging system and stochastic processing of the measurements. We present a theoretical analysis of SGD-Net that establishes that the network can be trained to approximate the corresponding batch deep unfolding network to desired accuracy. We also demonstrate the practical relevance of SGD-Net by reconstructing images in *intensity diffraction tomography (IDT)* [26], [38] and *sparse-view computed tomography (CT)* [39]. Our simulations demonstrate the effectiveness of SGD-Net in achieving comparable imaging

quality to batch deep unfolding networks at a fraction of computational complexity. SGD-Net thus addresses an important gap in the current literature on deep unfolding by providing an efficient framework applicable to a wide variety of imaging problems.

II. BACKGROUND

A. Inverse Problems in Computational Imaging

Computational imaging problems can usually be posed as the reconstruction of an unknown image $\mathbf{x} \in \mathbb{C}^n$ from a set of corrupted measurements $\mathbf{y} \in \mathbb{C}^m$. The reconstruction is often formulated as an inverse problem

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} f(\mathbf{x}) \quad \text{with} \quad f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x}), \quad (1)$$

where g is the data-fidelity term that quantifies consistency with \mathbf{y} and h is the regularizer. For example, two widely used functions in the context of imaging problems are the least squares and total variation (TV) [1]

$$g(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2 \quad \text{and} \quad r(\mathbf{x}) = \tau \|\mathbf{Dx}\|_1, \quad (2)$$

where \mathbf{A} is the measurement operator and \mathbf{D} is the discrete gradient operator. The data-fidelity term in (2) assumes a linear measurement model $\mathbf{y} = \mathbf{Ax} + \mathbf{e}$, where the measurement operator $\mathbf{A} \in \mathbb{C}^{m \times n}$ characterizes the response of the imaging system and $\mathbf{e} \in \mathbb{C}^m$ is the noise, which is often assumed to be independent and identically distributed (i.i.d.) Gaussian.

Many popular regularizers, such as the ones based on the ℓ_1 -norm, are nondifferentiable. Proximal algorithms [40], such as the *proximal gradient method (PGM)* [41]–[45] and the *alternating direction method of multipliers (ADMM)* [46]–[48], enable efficient minimization of nonsmooth functions without differentiating them by using the *proximal operator*

$$\text{prox}_{\mu h}(\mathbf{z}) := \arg \min_{\mathbf{x}} \left\{ \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 + \mu h(\mathbf{x}) \right\}, \quad (3)$$

where $\mu > 0$ is a parameter. Note that the proximal operator can be interpreted as the regularized image denoiser for AWGN with noise of variance μ .

B. Image Reconstruction Using Deep Learning

Recently, deep learning has gained popularity due to its effectiveness for solving imaging inverse problems. A widely used approach first brings the measurements to the image domain and then trains a deep network to map the corresponding low-quality images $\{\tilde{\mathbf{x}}_j\}$ to their clean target versions $\{\mathbf{x}_j\}$ by solving an optimization problem [11], [14], [15]

$$\arg \min_{\theta} \frac{1}{M} \sum_{j=1}^M \mathcal{L}(\mathcal{T}_{\theta}(\tilde{\mathbf{x}}_j) - \mathbf{x}_j), \quad (4)$$

where \mathcal{T}_{θ} represents the CNN parametrized by θ trained under the loss function \mathcal{L} . Popular loss functions include the ℓ_2 -norm and ℓ_1 -norm [49]. For example, in the context of sparse-view CT, prior methods have trained \mathcal{T}_{θ} for mapping a filtered back-projected image $\tilde{\mathbf{x}}$ to a reconstruction \mathbf{x} from a fully-sampled groundtruth data [7], [15].

The idea of end-to-end inversion can be refined by including the measurement operator into the CNN architecture. Inspired by LISTA [50], the corresponding *unfolding algorithms* interpret iterations of a regularized inversion as layers of a CNN and train it end-to-end in a supervised fashion [51]–[55]. In the context of compressive sensing, ADMM-CSNet [56] and ISTA-Net⁺ [28] have considered jointly training the image transforms and shrinkage functions within an unfolded algorithm. A related class of methods [31], [32], [57], have included a full CNN as a trainable regularizer within an unfolded algorithm. Such unfolding algorithms have been shown to be effective in a number of problems [32], [58] and are closely related to the PnP/RED methods (discussed in Section II-C) that also combine the measurement operator and the imaging prior. Their main differences is that the former optimize the parameters in an end-to-end manner, and generally produce higher-quality results with fewer iterations. However, model-based deep unfolding architectures require the storage of all the measurements, parameters of the measurement operator, and intermediary activation maps at every iteration of the unfolded algorithm. This limits their ability to solve inverse problems where one needs to process high-dimensional data (for example, see relevant discussion in the section “*Memory Requirements*” in [57]). It is worth mentioning that several recent publications have explored memory-efficient training in model-based deep learning [59]–[62]. For example, [59] proposed to reduce the memory usage in unfolding networks through reverse recalculations. On the other hand, [60] proposed a memory-efficient architecture that relies on per-iteration refinement of CNNs. Those strategies are fundamentally different from our usage of stochastic approximations within data-consistency layers. In fact, all those approaches are fully complimentary and can be combined with our work for efficient training of model-based deep architectures.

C. Using Denoisers as Image Priors

Since the proximal operator (3) is mathematically equivalent to regularized image denoising, there has been considerable interest in developing denoiser-based iterative algorithms such as PnP [16] and RED [17]. The key idea in PnP is to replace the proximal operator with an advanced image denoiser H_{σ} , where $\sigma > 0$ controls the strength of denoising. This simple replacement enables PnP to regularize the problem by using advanced denoisers, such as BM3D [63] or DnCNN [18], that do not correspond to any explicit h . Recent studies have confirmed the effectiveness of PnP in a range of imaging applications [20], [64]–[66].

The RED framework is an alternative scheme where the denoiser can sometimes lead to an explicit regularization function [17]. In the most general setting, RED algorithms seek a fixed point \mathbf{x}^* that satisfies

$$G(\mathbf{x}^*) = \nabla g(\mathbf{x}^*) + \tau(\mathbf{x}^* - H_{\sigma}(\mathbf{x}^*)) = 0, \quad (5)$$

where $\tau > 0$ is the regularization parameter. Equivalently, \mathbf{x}^* satisfies

$$\mathbf{x}^* \in \text{zer}(G) := \{\mathbf{x} \in \mathbb{C}^n : G(\mathbf{x}) = \mathbf{0}\}. \quad (6)$$

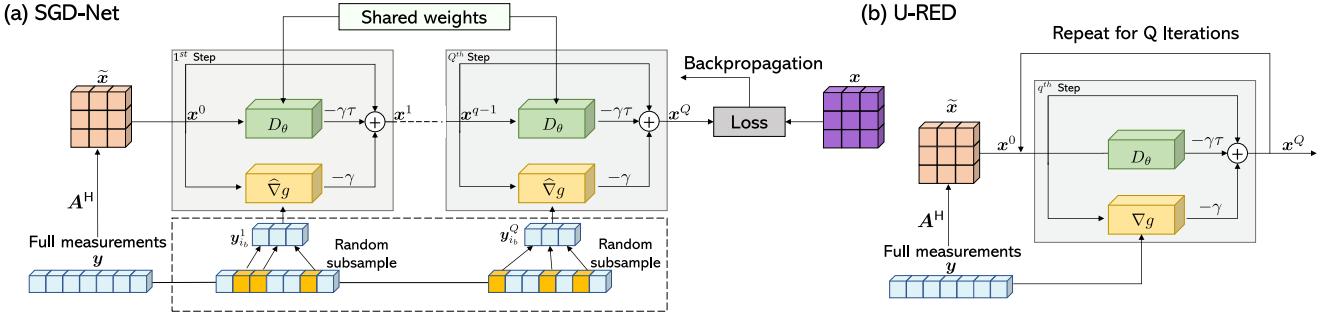


Fig. 1. Architecture of (a) *SGD-Net* with Q steps and (b) its batch variant that we label as *Unfolded RED* (*U-RED*). Data-consistency layers within SGD-Net rely on minibatch gradients with $B \ll I$ measurements, while those in U-RED rely on the full batch gradient using I measurements. The learned operator blocks within SGD-Net and U-RED share their weights across steps and are trained in an end-to-end fashion by accounting for the data-consistency layers.

When the denoiser is locally homogeneous and has a symmetric Jacobian [17], [23], the term $\tau(\mathbf{x} - H_\sigma(\mathbf{x}))$ corresponds to the gradient of the RED regularizer $h(\mathbf{x}) = (\tau/2)\mathbf{x}^\top(\mathbf{x} - H_\sigma(\mathbf{x}))$, which enables a simple interpretation of RED as an instance of (1). The excellent performance of RED together with learned CNN denoisers has been reported in super-resolution, phase retrieval, and compressed sensing [67], [68]. Additionally, prior work has developed a scalable *online* variant of RED, known as SIMBA, that is well suited for tomographic applications with a large number of projections [26]. However, unlike deep unfolding, the CNN in RED is not jointly trained with the measurement model, limiting its ability to capture the non-iid nature of the artifacts within iterations [31]. SGD-Net, introduced in the next section, is a natural extension of SIMBA [26] towards CNNs trained for artifact-removal, which leads to a scalable and end-to-end trainable deep network.

III. PROPOSED METHOD

We now introduce SGD-Net, which adopts stochastic processing of measurements within data-consistency layers. As demonstrated by our results in Section V, SGD-Net is ideal for data-intensive applications where the object features are difficult to characterize by a pre-trained CNN denoiser.

A. Stochastic Data-Consistency Layers

Consider a data-fidelity g that consists of a set of $I \geq 1$ component functions $\{g_i(\mathbf{x})\}$, where each g_i depends on some subset \mathbf{y}_i of the measurements in \mathbf{y} . For example, in tomographic imaging each \mathbf{y}_i corresponds to a single projection of an object along a specific angle [39]. When g corresponds to the least-squares penalty, the data-consistency layer within the deep unfolding network becomes

$$\nabla g(\mathbf{x}) = \frac{1}{I} \sum_{i=1}^I \nabla g_i(\mathbf{x}) = \frac{1}{I} \sum_{i=1}^I \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{x} - \mathbf{y}_i), \quad (7)$$

where $(\cdot)^\top$ denotes the conjugate transpose operation. Note that the complexity of ∇g scales *linearly* with the total number of components I . This means that when $I \rightarrow \infty$, the memory requirements or computation time of the traditional batch deep unfolding becomes impractical. The central idea of SGD-Net,

summarized in Fig 1, is to approximate the gradient at every step within the deep unfolding network with an average of $1 \leq B \ll I$ component gradients, which makes SGD-Net independent of the total number of components I . The corresponding *minibatch* gradient is computed as

$$\widehat{\nabla}g(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \nabla g_{i_b}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \mathbf{A}_{i_b}^\top (\mathbf{A}_{i_b} \mathbf{x} - \mathbf{y}_{i_b}), \quad (8)$$

where i_1, \dots, i_B are independent random indices that are selected uniformly from the set $\{1, \dots, I\}$. The minibatch size parameter $B \geq 1$ controls the number of gradient components at each step of SGD-Net. Note that (8) directly implies that $\mathbb{E}[\widehat{\nabla}g(\mathbf{x})] = \nabla g(\mathbf{x})$, where the expectation is taken over the random indices i_1, \dots, i_B .

For notational convenience, we will use the vector ϕ to denote the physical parameters within the data-consistency layers of the unfolded network. In the context of the batch network in Fig. 1(b), which uses all the measurements at every step, the physical parameters ϕ relate to the gradients $\{\nabla g_i\}$ at every step of the unfolded network. SGD-Net seeks to minimize the complexity of the unfolded network by replacing ϕ with its minibatch approximation $\widehat{\phi}$, obtained by applying (8) to every step. More specifically, for (7) we can denote the physical parameters as $\phi = (\phi_1, \phi_2)$ with

$$\phi_1 = \text{vec} \left(\frac{1}{I} \sum_{i=1}^I \mathbf{A}_i^\top \mathbf{A}_i \right) \quad \text{and} \quad \phi_2 = \frac{1}{I} \sum_{i=1}^I \mathbf{A}_i^\top \mathbf{y}_i, \quad (9)$$

where $\text{vec}(\cdot)$ vectorizes the matrix into a vector. Similarly, for the SGD-Net data-consistency layer in (8), we have $\widehat{\phi} = (\widehat{\phi}_1, \widehat{\phi}_2)$ with

$$\widehat{\phi}_1 = \text{vec} \left(\frac{1}{B} \sum_{i=1}^B \mathbf{A}_{i_b}^\top \mathbf{A}_{i_b} \right) \quad \text{and} \quad \widehat{\phi}_2 = \frac{1}{B} \sum_{i=1}^B \mathbf{A}_{i_b}^\top \mathbf{y}_{i_b}. \quad (10)$$

Note that we directly have that $\mathbb{E}[\widehat{\phi}] = \phi$, where the expectation is taken over the random indices i_1, \dots, i_B . It is important to note that ϕ and $\widehat{\phi}$ do not need to be physically stored as vectors, since they are simply used for notational convenience in our theoretical analysis.

B. Stochastic Deep Unfolding Network

Given an initial solution $\tilde{\mathbf{x}} = \mathbf{A}^H \mathbf{y}$, SGD-Net iteratively refines it by infusing information from both the minibatch gradient of the data-fidelity term $\hat{\nabla} g$ and the learned operator D_θ defined as

$$D_\theta(\mathbf{x}) = (I - R_\theta)(\mathbf{x}) = \mathbf{x} - R_\theta(\mathbf{x}), \quad (11)$$

where R_θ is the artifact-removal CNN. Unlike in SIMBA [26], the prior in SGD-Net is optimized end-to-end using the training data to maximally reduce the artifacts. We fix the total number of SGD-Net steps to $Q \geq 1$, with each step given by

$$\mathbf{x}^{q+1} = \mathbf{x}^q - \gamma(\hat{\nabla} g(\mathbf{x}^q) + \tau D_\theta(\mathbf{x}^q)), \quad (12)$$

where $\gamma > 0$ is a step-size parameter. Fig. 1(a) illustrates the algorithmic details of SGD-Net, which can, in principle, be implemented with or without weight-sharing across the Q steps. In our implementation, we opted to share the weights of R_θ across different steps to make it more suitable for large-scale imaging applications. While there are no trainable parameters within the stochastic data-consistency layers, one can still use backpropagation to compute the gradient of SGD-Net with respect to the trainable parameters θ . Since SGD-Net randomly processes a subset of measurements at each step, the prediction of SGD-Net is in fact randomized. Our theoretical analysis in Section IV precisely characterizes the training of SGD-Net relative to that of the batch network that uses all the measurement in every step.

SGD-Net is compatible with any CNN architecture used to implement D_θ . In our implementation, we adopted U-Net [10], which has been extensively used in the context of imaging inverse problems [11], [69]. The CNN consists of four scales, each with a skip connection between downsampling and upsampling. These connections increase the effective receptive field of the CNN. The number of channels in each layer are 32, 64, 128, and 256. We make two additional modifications to the U-Net. First, the activation function in our setting is PReLU (parametric ReLU, $f(x) = \max(0, x) + a * \min(0, x)$, where a is a trainable parameter). Second, since we adopt small minibatch training, we use group normalization (GN) [70] as an alternative to batch normalization (BN). The computation in GN is independent of the minibatch dimension, which makes its accuracy stable for a wide range of minibatch sizes.

C. End-to-End Training of SGD-Net

Let $\mathcal{T}_{\theta; \phi}(\tilde{\mathbf{x}}, \mathbf{y})$ be the batch *unfolded RED (U-RED)* network, where θ are the *learnable* parameters within D_θ , ϕ are *physical* parameters within the data consistency layers, \mathbf{y} is the measurement vector, and $\tilde{\mathbf{x}}$ is the input to the network. End-to-end training seeks to compute the learnable parameters θ of $\mathcal{T}_{\theta; \phi}$ by minimizing a loss function F over M training samples $\{\mathbf{x}_j, \mathbf{y}_j\}$. Let F_j be a loss functions over the training sample $(\mathbf{x}_j, \mathbf{y}_j)$, then the training is formulated as the following optimization problem

over learnable parameters

$$\arg \min_{\theta} F(\theta; \phi) \quad \text{with} \quad F(\theta; \phi) := \frac{1}{M} \sum_{j=1}^M F_j(\theta; \phi). \quad (13)$$

For example, a popular choice for the loss function F_j as the *mean square error (MSE)* between the image predicted by $\mathcal{T}_{\theta, \phi}$ and the desired image \mathbf{x}_j

$$F_j(\theta; \phi) = \|\mathcal{T}_{\theta, \phi}(\tilde{\mathbf{x}}_j, \mathbf{y}_j) - \mathbf{x}_j\|_2^2 \quad \text{with} \quad \tilde{\mathbf{x}}_j = \mathbf{A}^H \mathbf{y}_j. \quad (14)$$

One can train both U-RED and SGD-Net using gradient-based optimizers, such as the *stochastic gradient descent (SGD)* [71]. When training SGD-Net, iteration k of training is performed by generating two sets of independent random variables. First, an index in j_k is randomly selected from $\{1, \dots, M\}$, then, a stochastic approximation $\hat{\phi}^k$ of ϕ is generated by replacing the batch gradients by their randomized minibatch approximations (8) at every step of the network. The trainable weights can then be updated as

$$\theta^{k+1} = \theta^k - \eta_k \hat{\nabla} F(\theta^k; \hat{\phi}^k), \quad (15)$$

where $\hat{\nabla} F(\theta^k; \hat{\phi}^k) = \nabla F_{j_k}(\theta^k; \hat{\phi}^k)$ and $\eta_k > 0$ is the step-size at the training iteration $k \geq 0$. As illustrated in Fig. 1, SGD-Net can significantly reduce the complexity of the data-consistency layers by using $\hat{\phi}^k$ instead of the full ϕ . In Section IV, we present a theoretical analysis of training SGD-Net using the SGD iteration (15).

We use a pre-training strategy to accelerate the training of the weights θ within SGD-Net. The training of the artifact-removal CNN R_θ is considerably faster than the training of the full SGD-Net, since the latter consist of both R_θ and data-consistency blocks. In order to accelerate the training, we first pre-train R_θ separately from a random initialization of θ using the MSE loss in (4). Since the weights θ are shared across steps of SGD-Net, we can then initialize them with those obtained from pre-training. We observed that this pre-training strategy is considerably more efficient than initializing the entire SGD-Net network with the random weights.

IV. THEORETICAL ANALYSIS

We now present the theoretical analysis of the training of SGD-Net by using the SGD algorithm in (14). Note that our analysis does not explicitly assume that the unfolded architecture corresponds to RED, which means that it is also applicable to other architectures, including those based on PnP.

Our theoretical analysis requires a number of technical assumptions that act as sufficient conditions for the main theorem below. Our first assumption is on Lipschitz continuity of the gradient of F over both ϕ and θ .

Assumption 1: Function F has a global minimizer θ^* and satisfies the following continuity assumption

$$\begin{aligned} & \|\nabla F(\theta_1; \phi_1) - \nabla F(\theta_2; \phi_2)\|_2 \\ & \leq L_\theta \|\theta_1 - \theta_2\|_2 + L_\phi \|\phi_1 - \phi_2\|_2, \end{aligned}$$

with $L_\theta, L_\phi > 0$ for every (θ_1, ϕ_1) and (θ_2, ϕ_2) .

Assumption 1 is a standard assumption in traditional optimization that holds for many widely-used cost functions F [37], [72], [73]. The only difference in our statement of this assumption is in the explicit definition of two constants L_θ and L_ϕ , which can be avoided by defining $L = \max\{L_\theta, L_\phi\}$. Note that we do *not* assume that F is convex.

A common assumption in stochastic optimization is that the minibatch gradients are unbiased estimators of the full gradient and have bounded variances [74], [75]. Our second assumption states this condition for the physical parameters of the unfolded network.

Assumption 2: The physical parameters of the network satisfy

$$\mathbb{E} [\hat{\phi}^k] = \phi, \quad \mathbb{E} [\|\hat{\phi}^k - \phi\|_2^2] \leq \frac{\sigma^2}{B},$$

for all iterations $k \geq 0$, where $\sigma > 0$ is a constant and $B \geq 1$ is the minibatch size.

This is a mild assumption since SGD-Net obtains $\hat{\phi}$ by relating ∇g at every step to $\hat{\nabla} g$ via (9) and (10). This automatically ensures that $\hat{\phi}$ is an unbiased estimator of ϕ . Our final assumption is related to Assumption 2, but considers the selection of the gradients $\hat{\nabla} F$ during SGD training.

Assumption 3: The stochastic gradients in (15) satisfy the following two conditions for any fixed vectors θ and ϕ .

(a) The stochastic gradient is unbiased:

$$\mathbb{E} [\hat{\nabla} F(\theta; \phi)] = \nabla F(\theta; \phi).$$

(b) The variance of the stochastic gradient is bounded:

$$\mathbb{E} [\|\hat{\nabla} F(\theta; \phi) - \nabla F(\theta; \phi)\|_2^2] \leq \epsilon^2.$$

The expectations are taken with respect to the random index $j \in \{1, \dots, M\}$ used to select the training sample.

Theorem 1: Run the SGD learning in (15) for $K \geq 1$ iterations under Assumptions 1-3 using the step-size parameters $0 < \eta_k \leq 1/L$ with $L = \max\{L_\theta, L_\phi\}$ and the minibatch size $B \geq 1$. Then, the iterates generated by (15) satisfy the bound

$$\begin{aligned} \sum_{k=0}^{K-1} \eta_k \mathbb{E} [\|\nabla F(\theta^k; \phi)\|_2^2] &\leq 2(F(\theta^0; \phi) - F(\theta^*; \phi)) \\ &+ \frac{L^2 \sigma^2}{B} \left(\sum_{k=0}^{K-1} \eta_k \right) + L \epsilon^2 \left(\sum_{k=0}^{K-1} \eta_k^2 \right). \end{aligned}$$

Proof: See the appendix. ■

Theorem 1 allows us to establish various forms of convergence results by controlling the step-sizes η_k . For example, when $\eta_k = 1/(L\sqrt{K})$, one obtains

$$\min_{k \in \{0, \dots, K-1\}} \mathbb{E} [\|\nabla F(\theta^k; \phi)\|] \leq \frac{C}{\sqrt{K}} + \frac{L^2 \sigma^2}{B},$$

where $C = 2(F(\theta^0; \phi) - F(\theta^*; \phi))L + \epsilon^2$ is a constant. This implies that SGD in (15), which relies only on the minibatch approximation $\hat{\phi}^k$ of ϕ , achieves, in expectation, the *first-order necessary conditions* of optimality for (13) up to an error term $L^2 \sigma^2 / B$. This error term can be made as small as possible by

controlling B within SGD-Net. Theorem 1 shows that SGD-Net can closely approximate the performance of the full-batch U-RED for a small enough step size η_k and a large-enough minibatch size B . However, this theoretical result is not intended to be used as a recipe for finding good parameters η_k and B for training SGD-Net. In practice, it is better to tune those parameters empirically, as is commonly done in the literature on optimization for deep learning. Section V presents simulations showing the ability of SGD-Net to match the performance of U-RED for different values of B while also substantially reducing the training and testing complexity.

V. NUMERICAL VALIDATION

We now empirically validate SGD-Net in the context of two computational imaging modalities, IDT and sparse-view CT. Our first goal is to validate the proposed theorems in Section IV and the second one is to highlight the effectiveness and efficiency of our method for processing a large number of measurements. All the experiments were performed on a machine equipped with an Intel Xeon Gold 6130 Processor and eight NVIDIA GeForce RTX 2080 Ti GPUs.

Several image reconstruction methods were used as references, including TV [1], U-Net [10], RED-DnCNN [17] and ISTA-Net⁺ [28]. TV is formulated in (2), and was implemented using the accelerated proximal gradient descent method (APGM) [76]. U-Net corresponds to our own implementation of the architecture used in [14]. The network was trained in the usual supervised fashion using the ℓ_2 -loss [11]. We adopted the DnCNN architecture used in [67] as the AWGN denoiser for RED. The network has seventeen layers, including 15 hidden layers, an input layer, and an output layer. We have also experimented using RED with the U-Net like architecture from SGD-Net, but observed that this does not improve the image reconstruction quality achieved by RED. ISTA-Net⁺ is a widely-used deep unfolding architecture based on the feed-forward network obtained by unfolding and truncating ISTA. We unfolded ISTA-Net⁺¹ for 12 steps and set the number of feature maps in each convolution layers equal to 64 for the best SNR performance in our experimental settings. Note that U-RED uses the complete set of measurements I corresponds to the traditional deep unfolding of the batch RED algorithm. In all the experiments, we train SGD-Net and U-RED using the same training strategy and parameter initialization settings. The only difference between those two algorithms is in the number of measurements used in their data-consistency layers. All methods were implemented in Pytorch with a GPU backend.

We used the following *signal-to-noise ratio (SNR)* [26], [77] in dB for quantitatively comparing different algorithms

$$\text{SNR}(\hat{x}, x) = \max_{a, b \in \mathbb{R}} \left\{ 20 \log_{10} \left(\frac{\|x\|_2}{\|x - a\hat{x} + b\|_2} \right) \right\}, \quad (16)$$

where \hat{x} and x represents the noisy vector and ground truth respectively, while the purpose of a and b is to adjust for contrast

¹The code for ISTA-Net⁺ is publicly available at <https://github.com/jianzhangs/ISTA-Net-PyTorch>

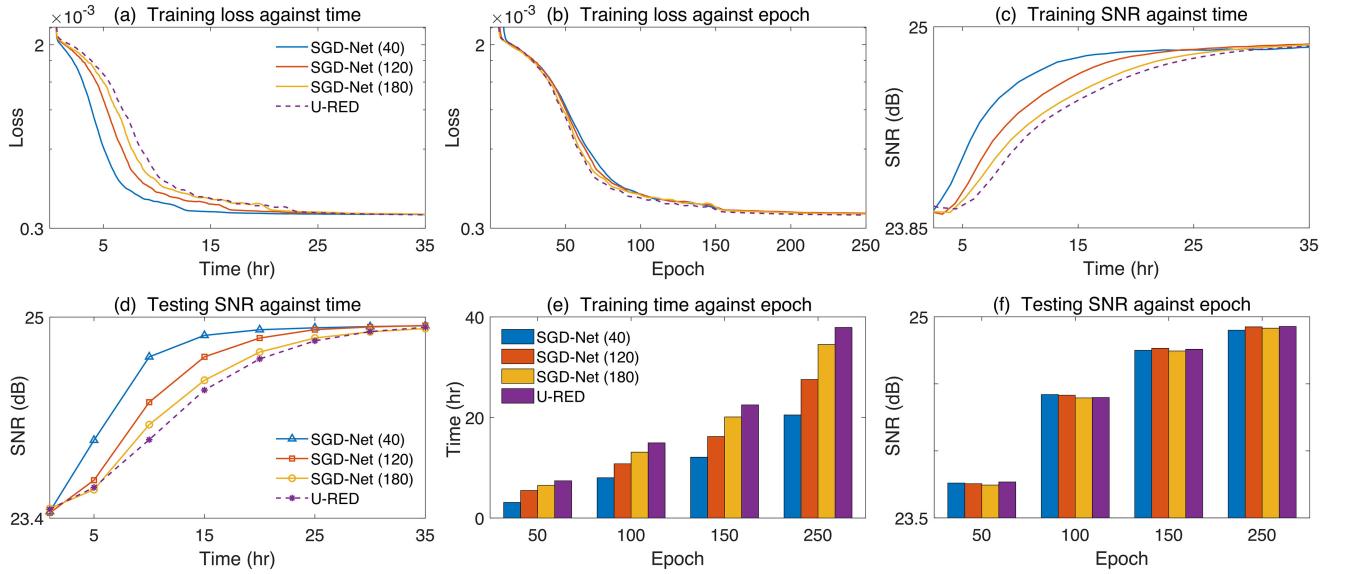


Fig. 2. Quantitative evaluation of SGD-Net on IDT for different minibatch sizes $B \in \{40, 120, 180\}$ used at each step of the network against U-RED using the full batch of $I = 240$ measurements. (a) Illustration of the loss against time in hours for different values of B evaluated on the training set. (b) Illustration of the loss against the epoch number evaluated on the training set. (c) Illustration of the SNR (dB) against time evaluated on the training set. (d) Illustration of the SNR (dB) against time in hours evaluated on the testing set. (e) Illustration of the amount of time required to reach a certain epoch for different values of B . (f) Illustration of the SNR (dB) achieved at different epochs for different values of B evaluated over the testing set. The figure highlights that by using minibatches of size $1 \leq B \ll I$ one can achieve nearly $2 \times$ improvement in training time over U-RED for the same final imaging quality.

and offset. We also used the *structural similarity index measure (SSIM)* [78] as an alternative metric.

A. Intensity Diffraction Tomography

IDT [38] is a data intensive computational imaging modality that seeks to recover the spatial distribution of the complex permittivity contrast of an object given a set of its intensity-only measurements. In this problem, \mathbf{A} consists of a set of I complex measurement operators $[\mathbf{A}_1, \dots, \mathbf{A}_I]^T$, where each \mathbf{A}_i is a convolution corresponding to the i th measurement \mathbf{y}_i . In the simulation, we extracted a random subset of 350 slices of 320×320 images for training, 10 images for validation, and 35 images for testing from the NYU fastMRI Initiative database [79]. Followed by the experimental setup in [26], [38], the simulated images are assumed to be on the focal plane $z = 0\mu\text{m}$ with LEDs located at $z_{\text{LED}} = -70\text{mm}$. The wavelength of the illumination was set to $\lambda = 630\text{nm}$ and the background medium index was assumed to be water with $\epsilon_b = 1.33$. We generated $I = 240$ intensity measurements with $40 \times$ microscope objectives (MO) and 0.65 numerical aperture (NA). All simulated measurements were additionally corrupted by AWGN corresponding to $\{15, 20, 25\}$ dB of input SNR.

We trained SGD-Net and U-RED with the initialization $\mathbf{x}^0 = \mathbf{A}^H \mathbf{y}$, where \mathbf{A}^H denotes the conjugate transpose. The proposed recursive model SGD-Net and U-RED were unrolled for $Q = 8$ steps and trained using SGD with minibatch size 1. While the step-size parameter of the data-consistency block in each step of SGD-Net and U-RED was fixed to $\gamma = 5 \times 10^{-3}$, the regularization parameter was set as a learnable parameter, initialized with $\tau = 2$. The learning rate of SGD was set in two

stages. In the first 150 epochs, we adopt the cyclic learning rate policy [80], where the policy cycles the learning rate between 8×10^{-3} and 4×10^{-3} with 2000 training iterations in the decreasing half of a cycle. In stage 2, the learning rate was gradually reduced by a factor of 0.5 every 50 epochs. The number of total training epochs was 300. It is worth to note that each SGD-Net was trained with the noise corresponding to 20 dB of input SNR in order to test the stability of the proposed method with respect to changes in amount of measurements noise. Both TV and RED-DnCNN were initialized with $\mathbf{x} = 0$, and we used `fminbound` in the `scipy.optimize` toolbox to identify the optimal regularization parameters. We trained the DnCNN denoiser in RED for AWGN removal at four noise levels corresponding to $\sigma \in \{5, 10, 15, 20\}$. For each experiment, we selected the denoiser achieving the highest SNR value. Several instances of U-Net and ISTA-Net⁺ were trained by mapping the *backprojection* (BP) $\mathbf{A}^H \mathbf{y}$ to the ground truth for each input SNR levels. We initialized the step-size and the regularization parameters of ISTA-Net⁺ to the same values as SGD-Net, subsequently learned all these parameters during training, as done in the original paper.

Fig. 2 highlights the ability of SGD-Net to reduce the computational complexity of training compared to the full batch network U-RED. The three top plots compares the average loss and SNR achieved by SGD-Net when evaluated on the training set using different values for B at each step. Note that U-RED uses a fixed set of full ($I = 240$) measurements, while SGD-Net selects a random subset of B measurements at every step. Fig. 2(d) presents the average SNR achieved by SGD-Net against the training time evaluated on the testing set. Fig. 2(e) highlights the time necessary to run a fixed number of epochs

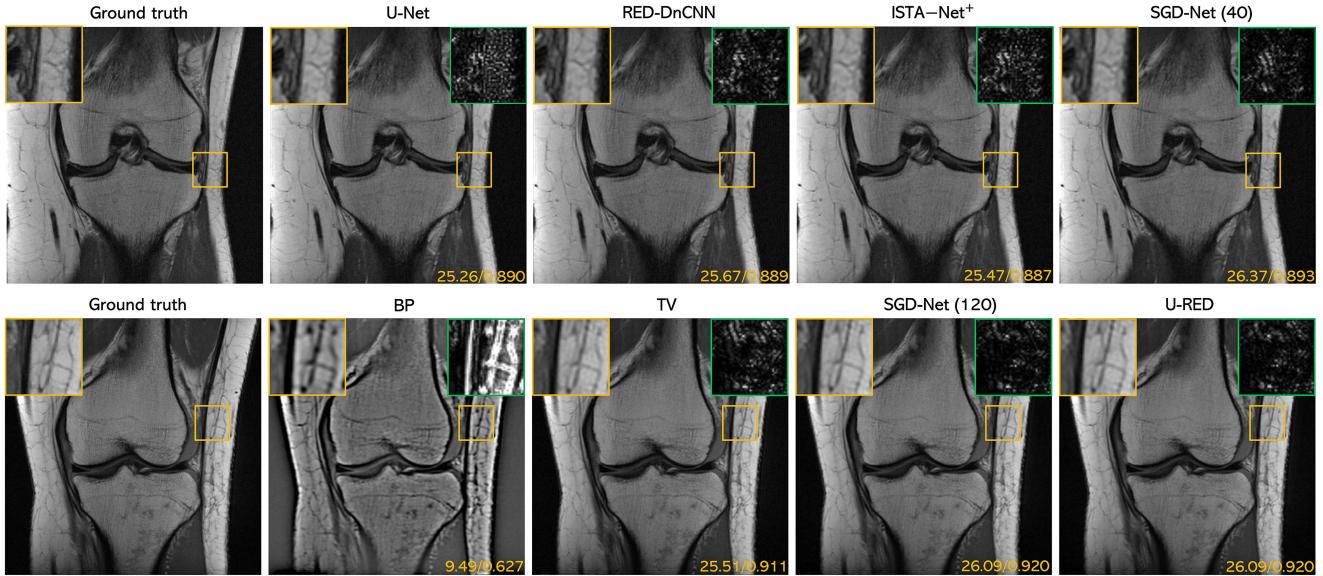


Fig. 3. Quantitative evaluation of several well-known methods on IDT under noise corresponding to input SNR of 20 dB. The total number of IDT measurements in this experiment is $I = 240$. All the baseline methods use the full set of measurements I during reconstruction. U-RED corresponds to the full batch architecture that uses all the measurements at every step. SGD-Net (40) and SGD-Net (120) use minibatches of size $B = 40$ and $B = 120$, respectively, at every step. Each image is labeled with its SNR (dB) and SSIM values with respect to the original image. The yellow box provides a close-up with a corresponding error map provided on its right. The results highlight the competitive performance of SGD-Net relative to several well-known methods, while also showing its ability to match the imaging quality achieved by the batch U-RED network. Best viewed by zooming in the display.

TABLE I
SNR AND SSIM VALUES OBTAINED BY SEVERAL REFERENCE METHODS FOR IDT. NOTE THAT THE LAST TWO COLUMNS PROVIDE THE GPU MEMORY USAGE AND THE RUN-TIMES FOR ALL THE METHODS FOR RECONSTRUCTING A 320×320 IMAGE. THE TOTAL NUMBER OF MEASUREMENT IS $I = 240$. NOTE THE EXCELLENT BALANCE BETWEEN QUALITY AND COMPLEXITY ACHIEVED BY SGD-NET

Metric	SNR			SSIM			#Iterations	Size Model/Measurement	Time CPU/GPU
Method	Input-SNR (dB)	20-5	20	20+5	20-5	20	20+5		
TV		24.26	24.31	24.39	0.887	0.890	0.891	250	—/1.01 GB
U-Net		24.27	24.33	24.35	0.887	0.889	0.889	—	118.2 MB/—
ISTA-Net ⁺		24.39	24.41	24.47	0.889	0.890	0.890	12	6.90 MB/1.01 GB
RED-DnCNN		24.54	24.61	24.67	0.890	0.892	0.893	220	2.29 MB/1.01 GB
SGD-Net (40)		24.84	24.94	24.96	0.896	0.899	0.901	8	29.6 MB/0.17 GB
SGD-Net (120)		24.87	24.93	24.94	0.898	0.899	0.900	8	29.6 MB/0.51 GB
U-RED		24.89	24.93	24.94	0.898	0.899	0.900	8	29.6 MB/1.01 GB

for different values of B . Fig. 2(f) shows the SNR achieved by SGD-Net for different epochs. Note that the average time of training 250 epochs of SGD-Net using $B \in \{40, 120, 180\}$ and U-RED was 20.50 hours, 27.55 hours, 34.51, hours and 37.90 hours, respectively. We did not observe significant SNR differences when using smaller minibatches for training compared with the usage of the full measurements. This highlights the ability of SGD-Net to reduce the complexity in deep unfolding by maintaining excellent imaging quality.

Table I provides the final SNR and SSIM values achieved by SGD-Net and several baseline methods when applied to IDT at three noise levels. Overall, model-based deep learning methods, such as RED-DnCNN, ISTA-Net⁺ and SGD-Net, achieve the best performances. Moreover, SGD-Net using $B = 40$ and $B = 120$ matches the performance of the batch algorithms in terms of the final reconstruction quality. The runtime in the

table corresponds to the average inference time that excludes the model loading. Specifically, SGD-Net with $B = 40$ is around $2.9 \times$ faster than U-RED on GPU and around $4.2 \times$ faster than U-RED on CPU with parallel processing. The memory column in the table corresponds to the usage of GPU memory. Specifically, the memory considerations in image reconstruction must take into account the size of all variables related to the image volume \mathbf{x} , the measured data $\{\mathbf{y}_i\}$, and the measurement operators $\{\mathbf{A}_i\}$. SGD-Net addresses the problem of storing and processing the measurements and the measurement operators on the GPU during end-to-end training. Our implementation stores each \mathbf{A}_i as two separate arrays for phase and absorption. In addition, each matrix is stored in the Fourier space to reduce the computational complexity of evaluating convolutions. This results in the storage of complex valued arrays for each, consisting of pairs of single precision floats for every element. The real and

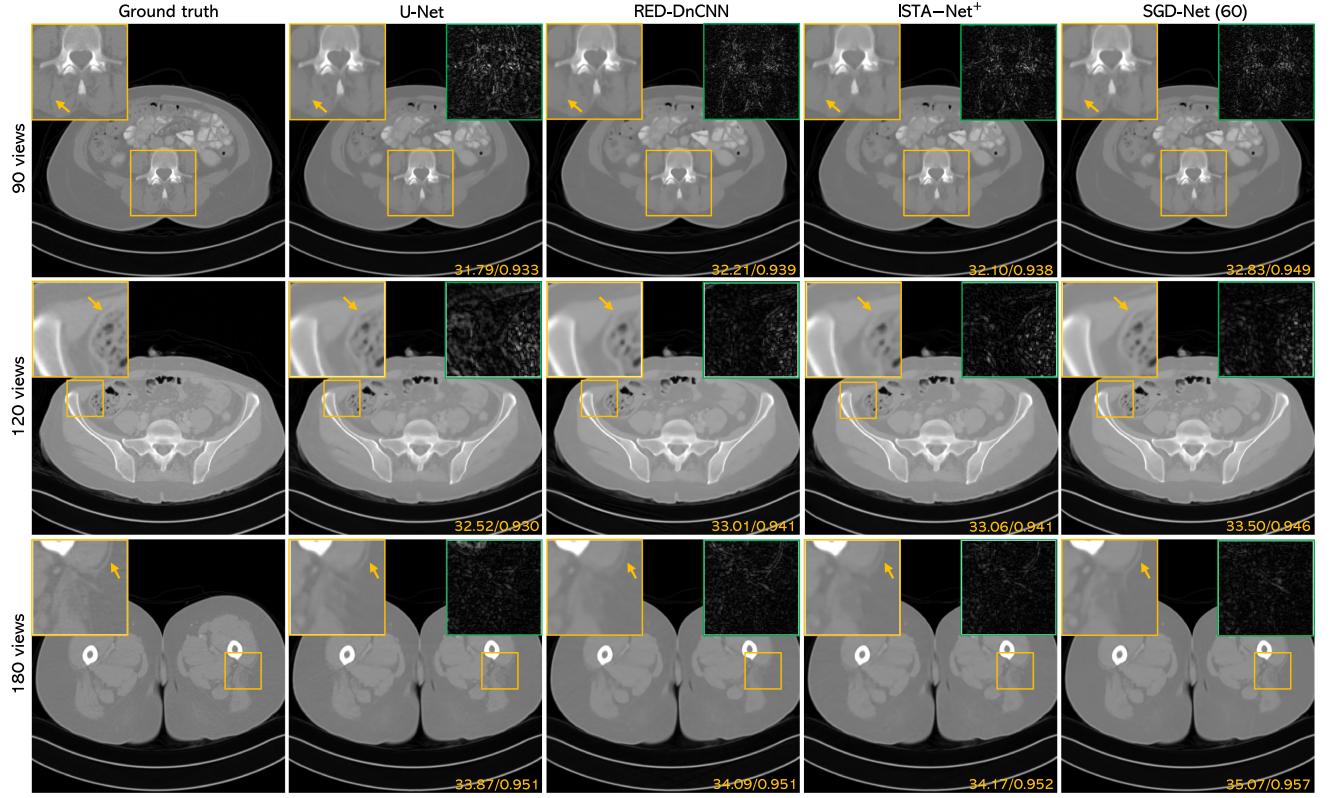


Fig. 4. Visual illustration of SGD-Net with $B = 60$ relative to several well-known deep-learning baseline methods on sparse-view CT with $\{90, 120, 180\}$ views and noise of 50 dB input SNR. Note that RED-DnCNN and ISTA-Net⁺ use the full set of measurements at each step. Each image is labeled with the corresponding SNR (dB) and SSIM values. This figure highlights that SGD-Net can achieve competitive performance relative to several well-known methods, while also providing a mechanism to reduce the complexity of data-consistency layers. Best viewed by zooming in the display.

imaginary parts of each array were then separated into two input channels of SGD-Net. Thus, the shape of each measurements and measurement operators in U-RED for reconstructing one slice is $1 \times 320 \times 320 \times 240 \times 2$. A detailed discussion on the IDT forward model is available in [26], [38]. While U-RED requires 1.01 GB of GPU memory due to its processing of all measurements in every iteration, SGD-Net with $B = 40$ requires only 0.17 GB, which is about 1/6 th of the full volume. This highlights the potential of applying SGD-Net to large scale image reconstruction.

Fig. 3 visually highlights the imaging quality of SGD-Net relative to several baseline methods. The top row of Fig. 3 presents the results obtained by several learning-based methods. As shown in the zoomed regions and the corresponding error maps, SGD-Net with $B = 40$ outperforms all other learning based methods. The bottom row of Fig. 3 highlights the comparable quality obtained by SGD-Net with $B = 120$ and that of U-RED using all $I = 240$ measurements.

B. Sparse-View CT

Conventional CT requires many views for high-quality image reconstruction. In the following experiments, we explore the possibility of high-quality imaging when reducing the number of views in CT imaging. We consider reconstruction of simulated data obtained from the clinically realistic CT images provided by

Mayo Clinic for the AAPM Low Dose CT Grand Challenge [81]. The data from 7 patients was used for training, one patient data for validation, and two patients' data for testing. This provides us with 2070 slices of 512×512 images for training, 150 slices of 512×512 images for validation. The testing data consists of 275 slices of 512×512 images. We implemented \mathbf{A} and \mathbf{A}^H with RayTransform in Operator Discretization Library (ODL).² [82], which uses GPU accelerated astra-gpu backend [83]. In particular, the scanning geometry is a fan-beam source with $I \in \{90, 120, 180\}$ projection views equally distributed around 360° and 1447 detector pixels. The sinograms were generated by slightly perturbing the view angles by a zero-mean AWGN with standard deviation of 0.003 degrees so as to make the experiments more realistic [77]. We added Gaussian noise to the sinograms to make the input SNR equal to 50 dB.

We trained the SGD-Net by using the filtered backprojection (FBP) initialization $\mathbf{x}^0 = \mathbf{A}^H \mathbf{Fy}$. The FBP was performed with a Hann filter by using the method `fbp-op` in ODL. We set the number of steps in SGD-Net to $Q = 8$. In these experiments, we numerically found that using Adam solver [84] with adaptive learning rate is around $2 \times$ faster than applying for SGD with cyclic learning rate when training both SGD-Net and U-RED. Hence, we trained SGD-Net and U-RED using the Adam solver,

²The code for ODL is publicly available at <https://github.com/odlgroup/odl>

TABLE II

SNR AND SSIM VALUES OBTAINED BY SEVERAL REFERENCE METHODS FOR THE RECONSTRUCTION OF A 512×512 IMAGE IN SPARSE-VIEW CT WITH NOISE OF 50 DB INPUT SNR. THE HIGHEST SNR AND SSIM VALUES ARE IN BOLD. NOTE THAT THE LAST TWO ROWS PROVIDE THE AVERAGE TEST-TIMES FOR ALL THE COMPETING METHODS FOR 180 VIEWS ON GPU AND CPU. SGD-NET ENABLES ONE TO BALANCE THE TIME COMPLEXITY OF RECONSTRUCTION AGAINST THE FINAL IMAGING QUALITY

Views	Metric	Method							
		FBP	TV	U-Net	RED-DnCNN	ISTA-Net ⁺	SGD-Net (30)	SGD-Net (60)	U-RED
90	SNR	17.56	30.09	31.17	31.93	32.01	32.76	32.88	32.87
	SSIM	0.362	0.924	0.930	0.935	0.934	0.942	0.943	0.943
120	SNR	20.03	31.23	32.54	33.13	33.17	33.91	33.95	34.01
	SSIM	0.449	0.929	0.936	0.941	0.940	0.948	0.949	0.950
180	SNR	23.19	32.97	34.04	34.49	34.61	35.44	35.46	35.46
	SSIM	0.582	0.940	0.948	0.950	0.951	0.957	0.958	0.958
Time (views=180)	CPU	0.859s	304.1s	2.061s	460.3s	15.95s	11.56s	13.31s	20.72s
	GPU	0.147s	13.58s	0.217s	5.177s	0.331s	0.269s	0.278s	0.325s

with minibatch size 2 and weight decay 2×10^{-8} . As can be seen below the training with Adam seems to behave similarly to that with SGD, with SGD-Net nearly matching the performance of U-RED. For physical parameters, we fixed the step-size in each step to $\gamma = 5 \times 10^{-3}$ and initialized the trainable regularization parameter to $\tau = 4$. The learning rate starts from 1×10^{-3} and is halved at epoch 20, then gradually reduced by a factor of 0.7 every 10 epochs. The number of total training epochs is 100. We adopt gradient clipping [85] in order to accelerate and stabilize the training. We also applied the same training schemes for U-Net and ISTA-Net⁺. We set the number of iterations for TV and RED to 240. For DnCNN in RED, we trained it for AWGN removal at four noise levels corresponding to $\sigma \in \{2, 5, 10, 15\}$. Several instances of U-Net and ISTA-Net⁺, corresponding to different numbers of views, were trained by mapping the FBP $\mathbf{A}^H \mathbf{F} \mathbf{y}$ to the ground truth for each sparse view.

Table II provides the SNR and SSIM values obtained by SGD-Net and the baseline methods when applied to sparse-view CT with different projection views. Overall, all methods offer significant gains over FBP, with model-based deep learning methods (RED-DnCNN, ISTA-Net⁺ and SGD-Net) achieving the best performance. Moreover, SGD-Net using $B \in \{30, 60\}$ achieves comparable image reconstruction quality to U-RED, which highlights the potential of using minibatches within deep unfolding networks. Similar conclusions can be drawn by Fig. 5, which shows the statistical evaluation of SGD-Net and other deep-learning methods. Visual inspection of the results highlight the excellent performance of SGD-Net. Specifically, Fig. 4 presents the reconstruction results by all the learning-based methods. SGD-Net with $B = 60$ performs better than other methods on all individual slices, highlighted by the yellow arrow. In Fig. 6, FBP is dominated by streaking artifacts, while TV reduces those artifacts, but blurs the fine structures by producing cartoon-like features. The zoomed regions suggest that SGD-Net with $B = 30$ can accurately reconstruct the fine details as good as its batch version using all the measurements.

Fig. 7 provides several additional evaluations highlighting the influence of the number of steps and minibatch sizes. Fig. 7(a) shows the SNR performance versus number of unrolling steps. It can be observed that the average SNR values improve as we increase the number of model steps, and SGD-Net with $B = 60$ consistently achieves a similar performance with the full

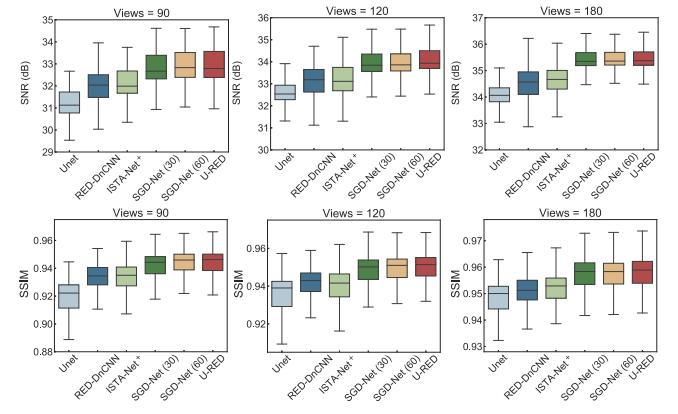


Fig. 5. Quantitative evaluation of several reference methods on sparse-view CT with noise corresponding to 50 dB of input SNR. The top row quantifies the image quality in terms of SNR (dB), while the bottom row in terms of SSIM. Columns from left to right provide results for 90, 120, and 180 views. This figure highlights that the usage of minibatches within SGD-Net does not reduce its ability to achieve high imaging quality.

batch unrolled RED algorithm. Fig. 7(b) shows the performance of SGD-Net trained for $B \in \{30, 60, 90\}$ and U-RED when tested with different minibatch sizes. This results highlights the robustness of pre-trained SGD-Net to the changes in minibatch sizes. One can observe that the performance of U-RED degrades much faster than SGD-Net with $B = 30$ when using smaller minibatches. Fig. 7(c) shows the computational time of using different minibatch sizes for reconstructing all the testing slices. We conducted 40 trials and 20 trials to calculate the average performance for each minibatch on GPU and CPU, respectively.

VI. CONCLUSION

The proposed SGD-Net method introduces stochastic approximation to the data-consistency layers of deep unfolding networks. Such approximations lead to complexity reductions during both training and testing of the complete network, making the network potentially applicable for problems with a large number of measurements. We provided extensive numerical results motivating the practical relevance of SGD-Net. Our results indicate that SGD-Net provides competitive imaging quality compared to the traditional and learning-based methods due to

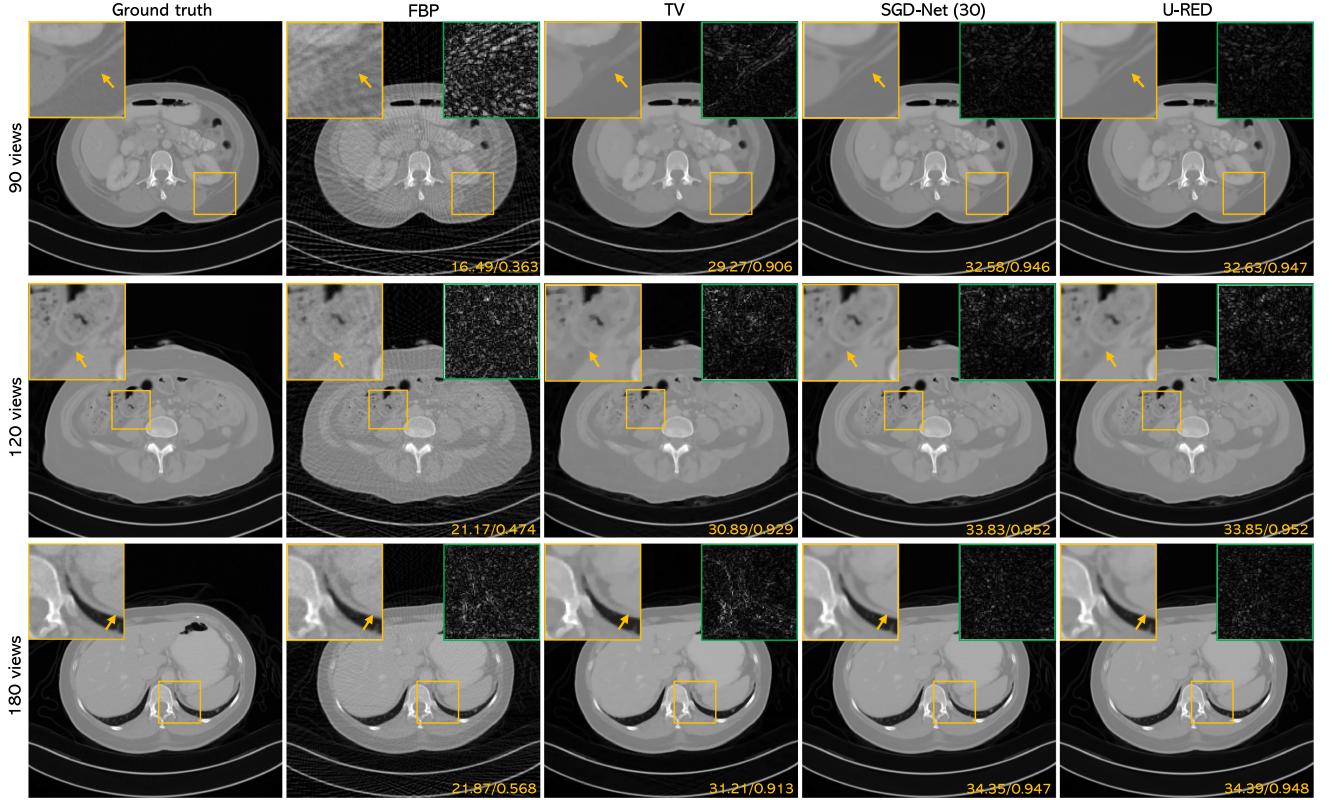


Fig. 6. Visual illustration of SGD-Net with $B = 30$ relative to FBP, TV, and U-RED on sparse-view CT with $\{90, 120, 180\}$ views and noise of 50 dB input SNR. See Table II for quantitative and Fig. 4 for visual comparisons with additional reference methods. Note that SGD-Net provides substantial improvements over traditional image reconstruction methods, matching the performance of U-RED that uses all the measurements in each step of the network. Best viewed by zooming in the display.

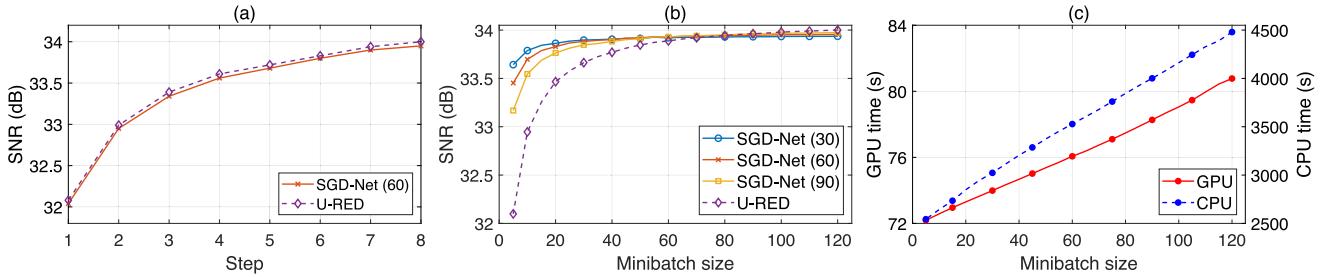


Fig. 7. Quantifying the influence of the number of steps and minibatch sizes. (a) SNR versus number of steps for the full batch network and SGD-Net with $B = 60$. Note how SGD-Net matches the full batch network in all settings. (b) SGD-Net architectures trained with certain minibatch sizes are tested on different minibatch sizes. This plot shows that networks trained on smaller minibatches achieve better generalization. (c) Runtime of SGD-Net for different minibatch sizes on GPU and CPU at test time. One can achieve significant savings for both GPU and CPU implementations with small minibatches.

the training of the priors in conjunction with the forward model. In particular, SGD-Net using small minibatches achieves the SNR performance of the network using all the available measurements at a fraction of complexity. While our experiments focused on IDT and CT, the method is broadly applicable to many other imaging modalities such as optical diffraction tomography (ODT) [86] and photoacoustic tomography (PAT) [29], where the evaluation of the measurement model is computationally intensive. Additionally, while we implemented SGD-Net based on the RED framework, the idea can be used for other model-based deep learning architectures.

APPENDIX

This appendix presents the proof of Theorem 1. Our analysis is based on the traditional analysis of SGD in the context of machine learning [37], [72], [73]. Note that while our theoretical bounds explicitly rely on the Lipschitz constant $L = \max\{L_\theta, L_\phi\}$, this constant does not provide a practical recipe for setting parameters η and B . Instead, those parameters should be empirically fine-tuned as is commonly done in the context of optimization for deep learning.

Assumption 1 leads to two bounds. By setting $\phi_1 = \phi_2 = \phi$ in Assumption 1, we obtain the traditional Lipschitz continuity bound on the gradient

$$\|\nabla F(\boldsymbol{\theta}_1; \boldsymbol{\phi}) - \nabla F(\boldsymbol{\theta}_2; \boldsymbol{\phi})\|_2^2 \leq L\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2,$$

for all $\boldsymbol{\theta}_1$, $\boldsymbol{\theta}_2$, and $\boldsymbol{\phi}$, which directly leads to the traditional quadratic upper bound (see Lemma 1.2.3 in [72])

$$\begin{aligned} F(\boldsymbol{\theta}_1; \boldsymbol{\phi}) &\leq F(\boldsymbol{\theta}_2; \boldsymbol{\phi}) + \nabla F(\boldsymbol{\theta}_2; \boldsymbol{\phi})^\top(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) \\ &\quad + \frac{L}{2}\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2. \end{aligned} \quad (17)$$

By setting $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_2 = \boldsymbol{\theta}$ in Assumption 1, we obtain the following useful bound for our proof

$$\begin{aligned} \frac{1}{2}\|\nabla F(\boldsymbol{\theta}; \boldsymbol{\phi}) - \nabla F(\boldsymbol{\theta}; \hat{\boldsymbol{\phi}})\|_2^2 &\leq \frac{L^2}{2}\|\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}\|_2^2 \\ \Leftrightarrow -\nabla F(\boldsymbol{\theta}; \boldsymbol{\phi})^\top \nabla F(\boldsymbol{\theta}; \hat{\boldsymbol{\phi}}) &+ \frac{1}{2}\|\nabla F(\boldsymbol{\theta}; \hat{\boldsymbol{\phi}})\|_2^2 \\ &\leq \frac{L^2}{2}\|\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}\|_2^2 - \frac{1}{2}\|\nabla F(\boldsymbol{\theta}; \boldsymbol{\phi})\|_2^2 \end{aligned} \quad (18)$$

The unbiasedness and boundedness of the variance of the stochastic gradient in Assumption 3 implies that for any fixed vector $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, we have that

$$\begin{aligned} \mathbb{E} \left[\|\hat{\nabla} F(\boldsymbol{\theta}; \boldsymbol{\phi}) - \nabla F(\boldsymbol{\theta}; \boldsymbol{\phi})\|_2^2 \mid \boldsymbol{\theta}, \boldsymbol{\phi} \right] &\leq \epsilon^2 \\ \Leftrightarrow \mathbb{E} \left[\|\hat{\nabla} F(\boldsymbol{\theta}; \boldsymbol{\phi})\|_2^2 \mid \boldsymbol{\theta}, \boldsymbol{\phi} \right] &\leq \|\nabla F(\boldsymbol{\theta}; \boldsymbol{\phi})\|_2^2 + \epsilon^2, \end{aligned} \quad (19)$$

where the expectation is taken over the index $j \in \{1, \dots, M\}$ of the stochastic gradient.

Now, we are ready to establish the result in Theorem 1. Consider a single iteration of optimizing SGD-Net with SGD

$$\boldsymbol{\theta}^{k+1} = \boldsymbol{\theta}^k - \eta_k \hat{\nabla} F(\boldsymbol{\theta}^k; \hat{\boldsymbol{\phi}}^k).$$

From the quadratic upper bound (17), we get

$$\begin{aligned} F(\boldsymbol{\theta}^{k+1}; \boldsymbol{\phi}) - F(\boldsymbol{\theta}^k; \boldsymbol{\phi}) &\leq \nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})^\top(\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k) + \frac{L}{2}\|\boldsymbol{\theta}^{k+1} - \boldsymbol{\theta}^k\|_2^2 \\ &= -\eta_k \nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})^\top \hat{\nabla} F(\boldsymbol{\theta}^k; \hat{\boldsymbol{\phi}}^k) + \frac{\eta_k^2 L}{2}\|\hat{\nabla} F(\boldsymbol{\theta}^k; \hat{\boldsymbol{\phi}}^k)\|_2^2, \end{aligned} \quad (20)$$

where $\boldsymbol{\phi}$ represents the true physical parameters and $\hat{\boldsymbol{\phi}}^k$ their stochastic approximation. By taking the conditional expectation with respect to the previous trainable parameters $\boldsymbol{\theta}^k$ and physical parameters $\hat{\boldsymbol{\phi}}^k$, using (18) and (19), we obtain

$$\begin{aligned} \mathbb{E} \left[F(\boldsymbol{\theta}^{k+1}; \boldsymbol{\phi}) \mid \boldsymbol{\theta}^k, \hat{\boldsymbol{\phi}}^k \right] - F(\boldsymbol{\theta}^k; \boldsymbol{\phi}) &\leq -\eta_k \nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})^\top \mathbb{E} \left[\hat{\nabla} F(\boldsymbol{\theta}^k; \hat{\boldsymbol{\phi}}^k) \mid \boldsymbol{\theta}^k, \hat{\boldsymbol{\phi}}^k \right] \\ &\quad + \frac{\eta_k}{2}\|\nabla F(\boldsymbol{\theta}^k; \hat{\boldsymbol{\phi}}^k)\|_2^2 + \frac{\eta_k^2 L \epsilon^2}{2} \\ &\leq -\frac{\eta_k}{2}\|\nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})\|_2^2 + \frac{\eta_k L^2}{2}\|\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}^k\|_2^2 + \frac{\eta_k^2 L \epsilon^2}{2}, \end{aligned}$$

where we also used the unbiasedness of the stochastic gradient $\hat{\nabla} F$ and the fact that $0 < \eta_k \leq 1/L$. By rearranging the terms, using Assumption 2, and taking the law of total expectation

$$\begin{aligned} \mathbb{E} \left[\eta_k \|\nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})\|_2^2 \right] &\leq 2(\mathbb{E} [F(\boldsymbol{\theta}^k; \boldsymbol{\phi})] - \mathbb{E} [F(\boldsymbol{\theta}^{k+1}; \boldsymbol{\phi})]) \\ &\quad + \left(\frac{\eta_k L^2 \sigma^2}{B} + \eta_k^2 L \epsilon^2 \right). \end{aligned}$$

By summing this bound over $0 \leq k \leq K-1$, we get

$$\begin{aligned} \sum_{k=0}^{K-1} \eta_k \mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}^k; \boldsymbol{\phi})\|_2^2 \right] &\leq 2(F(\boldsymbol{\theta}^0; \boldsymbol{\phi}) - \mathbb{E} [F(\boldsymbol{\theta}^K; \boldsymbol{\phi})]) + \sum_{k=0}^{K-1} \left(\frac{\eta_k L^2 \sigma^2}{B} + \eta_k^2 L \epsilon^2 \right) \\ &\leq 2(F(\boldsymbol{\theta}^0; \boldsymbol{\phi}) - F(\boldsymbol{\theta}^*; \boldsymbol{\phi})) + \sum_{k=0}^{K-1} \left(\frac{\eta_k L^2 \sigma^2}{B} + \eta_k^2 L \epsilon^2 \right), \end{aligned}$$

where we used the fact that $F(\boldsymbol{\theta}^*; \boldsymbol{\phi}) \leq \mathbb{E}[F(\boldsymbol{\theta}^K; \boldsymbol{\phi})]$, with $\boldsymbol{\theta}^*$ denoting a global minimizer of F .

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 Jeffrey’s 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] Y. Hu, S. G. Lingala, and M. Jacob, “A fast majorize-minimize algorithm for the recovery of sparse and low-rank matrices,” *IEEE Trans. Image Process.*, vol. 21, no. 2, pp. 742–753, Feb. 2012.
- [6] S. Lefkimiatis, J. P. Ward, and M. Unser, “Hessian Schatten-norm regularization for linear inverse problems,” *IEEE Trans. Image Process.*, vol. 22, no. 5, pp. 1873–1888, May 2013.
- [7] 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, 2017.
- [8] 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.
- [9] 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.
- [10] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Proc. Med. Image. Comput. Comput. Assist. Intervent.*, 2015, pp. 234–241.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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. 14 678–14 688, May 2018.
- [15] 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.

[16] 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.*, Austin, TX, USA, Dec. 3-5, 2013, pp. 945–948.

[17] 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.

[18] 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.

[19] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comp. Imag.*, vol. 3, no. 1, pp. 84–98, Mar. 2017.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] E. K. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, "Plug-and-play methods provably converge with properly trained denoisers," in *Proc. 36th Int. Conf. Mach. Learn.*, vol. 97, Long Beach, CA, USA, Jun. 09–15 2019, pp. 5546–5557.

[25] G. Mataev, P. Milanfar, and M. Elad, "DeepRED: Deep image prior powered by RED," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Oct. 2019, pp. 1–10.

[26] 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. Topics Signal Process.*, vol. 14, no. 6, pp. 1163–1175, Oct. 2020, doi: [10.1109/JSTSP.2020.2999820](https://doi.org/10.1109/JSTSP.2020.2999820).

[27] 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. Topics Signal Process.*, vol. 14, no. 6, pp. 1–1, Oct. 2020, doi: [10.1109/JSTSP.2020.2998402](https://doi.org/10.1109/JSTSP.2020.2998402).

[28] J. Zhang and B. Ghanem, "ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 1828–1837.

[29] 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.

[30] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018.

[31] 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.

[32] S. A. Hosseini, B. Yaman, S. Moeller, M. Hong, and M. Akçakaya, "Dense recurrent neural networks for accelerated MRI: History-cognizant unrolling of optimization algorithms," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 6, pp. 1280–1291, Oct. 2020.

[33] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya, "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magn. Reson. Med.*, Wiley Online Library, vol. 84, no. 6, pp. 3172–3191, Jul. 2020.

[34] L. Bottou and O. Bousquet, "The tradeoffs of large scale learning," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 3–6, 2007, pp. 161–168.

[35] D. P. Bertsekas, "Incremental proximal methods for large scale convex optimization," *Math. Program. Ser. B*, vol. 129, pp. 163–195, 2011.

[36] D. Kim, D. Pal, J. Thibault, and J. A. Fessler, "Accelerating ordered subsets image reconstruction for X-ray CT using spatially nonuniform optimization transfer," *IEEE Trans. Med. Imag.*, vol. 32, no. 11, pp. 1965–1978, Nov. 2013.

[37] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Rev.*, vol. 60, no. 2, pp. 223–311, 2018.

[38] R. Ling, W. Tahir, H.-Y. Lin, H. Lee, and L. Tian, "High-throughput intensity diffraction tomography with a computational microscope," *Biomed. Opt. Exp.*, vol. 9, no. 5, pp. 2130–2141, May 2018.

[39] A. C. Kak and M. Slaney, "Principles of computerized tomographic imaging," *Med. Physics*, vol. 29, no. 1, pp. 107–107, 2002.

[40] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 123–231, 2014.

[41] J. Eckstein and D. P. Bertsekas, "On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Math. Program.*, vol. 55, pp. 293–318, 1992.

[42] M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.

[43] J. Bect, L. Blanc-Féraud, G. Aubert, and A. Chambolle, "A ℓ_1 -unified variational framework for image restoration," in *Proc. Euro. Conf. Comp. Vis.*, vol. 3024, New York, 2004, pp. 1–13.

[44] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, Nov. 2004.

[45] A. Beck and M. Teboulle, *Convex Optimization in Signal Processing and Communications*, Cambridge, Ch. Gradient-Based Algorithms With Applications to Signal Recovery Problems, 2009, pp. 42–88.

[46] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Trans. Image Process.*, vol. 19, no. 9, pp. 2345–2356, Sep. 2010.

[47] M. K. Ng, P. Weiss, and X. Yuan, "Solving constrained total-variation image restoration and reconstruction problems via alternating direction methods," *SIAM J. Sci. Comput.*, vol. 32, no. 5, pp. 2710–2736, Aug. 2010.

[48] 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.

[49] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 47–57, Mar. 2017.

[50] K. Gregor and Y. LeCun, "Learning fast approximation of sparse coding," in *Proc. 27th Int. Conf. Mach. Learn.*, Haifa, Israel, Jun. 21–24, 2010, pp. 399–406.

[51] U. Schmidt and S. Roth, "Shrinkage fields for effective image restoration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Columbus, OH, USA, Jun. 23–28, 2014, pp. 2774–2781.

[52] Y. Chen, W. Yu, and T. Pock, "On learning optimized reaction diffraction processes for effective image restoration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Boston, MA, USA, Jun. 8–10, 2015, pp. 5261–5269.

[53] U. S. Kamilov and H. Mansour, "Learning optimal nonlinearities for iterative thresholding algorithms," *IEEE Signal Process. Lett.*, vol. 23, no. 5, pp. 747–751, May 2016.

[54] E. Bostan, U. S. Kamilov, and L. Waller, "Learning-based image reconstruction via parallel proximal algorithm," *IEEE Signal Process. Lett.*, vol. 25, no. 7, pp. 989–993, Jul. 2018.

[55] C. Zhang, Y. Liu, F. Shang, Y. Li, and H. Liu, "A novel learned primal-dual network for image compressive sensing," *IEEE Access*, vol. 9, pp. 26041–26050, 2021.

[56] Y. Yang, J. Sun, H. Li, and Z. Xu, "ADMM-CSnet: A deep learning approach for image compressive sensing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 3, pp. 521–538, Mar. 2020.

[57] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 2, pp. 491–503, Feb. 2018.

[58] 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.

[59] M. Kellman *et al.*, "Memory-efficient learning for large-scale computational imaging," *IEEE Trans. Comp. Imag.*, vol. 6, pp. 1403–1414, 2020.

[60] I. Y. Chun, Z. Huang, H. Lim, and J. A. Fessler, "Momentum-net: Fast and convergent iterative neural network for inverse problems," *IEEE Trans. Pattern Anal. Mach. Intell.*, pp. 1–1, 2020.

[61] I. Y. Chun, X. Zheng, Y. Long, and J. A. Fessler, "BCD-net for low-dose CT reconstruction: Acceleration, convergence, and generalization," in *Proc. Med. Image Comput. Comput. Assist. Intervention*, 2019, pp. 31–40.

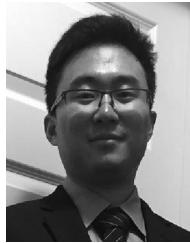
[62] S. Ravishankar, I. Y. Chun, and J. A. Fessler, "Physics-driven deep training of dictionary-based algorithms for mr image reconstruction," in *Proc. 51st Asilomar Conf. Signals, Syst. Comput.*, 2017, pp. 1859–1863.

[63] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 16, pp. 2080–2095, Aug. 2007.

- [64] A. Brifman, Y. Romano, and M. Elad, "Turning a denoiser into a super-resolver using plug and play priors," in *Proc. IEEE Int. Conf. Image Proc.*, Phoenix, AZ, USA, Sep. 25–28, 2016, pp. 1404–1408.
- [65] 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.
- [66] 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, 2020.
- [67] C. Metzler, P. Schniter, A. Veeraraghavan, and R. Baraniuk, "prDeep: Robust phase retrieval with a flexible deep network," in *Proc. 36th Int. Conf. Mach. Learn., Stockholmsmässan*, Stockholm Sweden, Jul. 10–15 2018, pp. 3501–3510.
- [68] Y. Sun, J. Liu, and U. S. Kamilov, "Block coordinate regularization by denoising," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 8–14, 2019, pp. 380–390.
- [69] K. Zhang, L. V. Gool, and R. Timofte, "Deep unfolding network for image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2020, pp. 3217–3226.
- [70] Y. Wu and K. He, "Group normalization," in *Proc. Eur. Conf. Comp. Vis.*, Sep. 2018, pp. 3–19.
- [71] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, no. 3, pp. 400–407, Sep. 1951.
- [72] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*. Norwell, MA, USA: Kluwer Academic Publishers, 2004.
- [73] P. Jain and P. Kar, "Non-convex optimization for machine learning," *Found. Trends Mach. Learn.*, vol. 10, no. 3–4, pp. 142–363, 2017.
- [74] S. Ghadimi and G. Lan, "Accelerated gradient methods for nonconvex nonlinear and stochastic programming," *Math. Program. Ser. A*, vol. 156, no. 1, pp. 59–99, Mar. 2016.
- [75] Z. Wu, Y. Sun, J. Liu, and U. S. Kamilov, "Online regularization by denoising with applications to phase retrieval," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Oct. 2019, pp. 1–9.
- [76] A. Beck and M. Teboulle, "Fast gradient-based algorithm for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [77] 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.
- [78] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [79] F. Knoll *et al.*, "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Radiol. Artif. Intell.*, vol. 2, no. 1, p. e190007, 2020.
- [80] L. N. Smith, "Cyclical learning rates for training neural networks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, Mar. 2017, pp. 464–472.
- [81] C. McCollough, "TU-FG-207A-04: Overview of the low dose CT grand challenge," *Med. Phys.*, vol. 43, no. 6Part35, pp. 3759–3760, 2016.
- [82] J. Adler and O. Öktem, "Solving ill-posed inverse problems using iterative deep neural networks," *Inverse Problems*, vol. 33, no. 12, p. 124007, 2017.
- [83] W. van Aarle *et al.*, "Fast and flexible x-ray tomography using the ASTRA toolbox," *Opt. Exp.*, vol. 24, no. 22, pp. 25 129–25 147, Oct. 2016.
- [84] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. on Learn. Representation*, 2015.
- [85] J. Zhang, T. He, S. Sra, and A. Jadbabaie, "Why gradient clipping accelerates training: A theoretical justification for adaptivity," in *Proc. Int. Conf. on Learn. Representation*, 2020.
- [86] U. S. Kamilov *et al.*, "Optical tomographic image reconstruction based on beam propagation and sparse regularization," *IEEE Trans. Comp. Imag.*, vol. 2, no. 1, pp. 59–70, Mar. 2016.



Jiaming Liu (Graduate 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 2018 in electrical and engineering 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. His research interests include machine learning, image processing, and optimization.



Yu Sun (Graduate 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 2017 in data analytics and statistics 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. His research interests include computational imaging, machine learning, deep learning, and optimization.



Weijie Gan (Graduate Student Member, IEEE) received the B.Eng. degree in automation from the South China University of Technology, Guangzhou, China, in 2018, and the M.S. degree in 2020 in computer science 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. His research interests include machine learning, image processing, and medical imaging.



Xiaojian Xu (Graduate Student Member, IEEE) received the B.Eng. degree in communication engineering, and the M.Eng. degree in communication and information system from the University of Electronic Science and Technology of China, Chengdu, China, in 2014 and 2017, respectively. She is currently working toward the Ph.D. degree with the Computational Imaging Group. Her research interests include computational imaging, optimization, inverse problems, computer vision, deep learning, machine learning, and signal processing.



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 are in 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 (currently 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 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 École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2008, 2011, and 2015, respectively. He is currently an Assistant Professor and the Director of Computational Imaging Group, Washington University in St. Louis, St. Louis, MO, USA. His primary research interests include computational imaging, signal and image processing, machine learning, and large-scale optimization. From 2015 to 2017, he was a Research Scientist at Mitsubishi Electric Research Laboratories, Cambridge, MA, USA. He was the recipient of the NSF CAREER Award and the IEEE Signal Processing Society's 2017 Best Paper Award. His Ph.D. thesis was selected as a finalist for the EPFL Doctorate Award in 2016. He was an Associate Editor for the IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING (2019–present), *Biological Imaging* (2020–present), and on IEEE Signal Processing Society's Computational Imaging Technical Committee (2016–present). He was a plenary Speaker at iTWIST 2018 and the Program Co-Chair for SampTA 2021 and BASP 2022.