Physics-Based Learned Design: Optimized Coded-Illumination for Quantitative Phase Imaging

Michael R. Kellman [ID], Emrah Bostan, Nicole A. Repina, and Laura Waller [ID]

Abstract—**Coded illumination can enable quantitative phase microscopy of transparent samples with minimal hardware requirements. Intensity images are captured with different source patterns, then a nonlinear phase retrieval optimization reconstructs the image. The nonlinear nature of the processing makes optimizing the illumination pattern designs complicated. The traditional techniques for the experimental design (e.g., condition number optimization, and spectral analysis) consider only linear measurement formation models and linear reconstructions. Deep neural networks (DNNs) can efficiently represent the nonlinear process and can be optimized over via training in an end-to-end framework. However, DNNs typically require a large amount of training examples and parameters to properly learn the phase retrieval process, without making use of the known physical models. In this paper, we aim to use both our knowledge of the physics and the power of machine learning together. We propose a new data-driven approach for optimizing coded-illumination patterns for an LED array microscope for a given phase reconstruction algorithm. Our method incorporates both the physics of the measurement scheme and the nonlinearity of the reconstruction algorithm into the design problem. This enables efficient parameterization, which allows us to use only a small number of training examples to learn designs that generalize well in the experimental setting without retraining. We show experimental results for both a well-characterized phase target and mouse fibroblast cells, using coded-illumination patterns optimized for a sparsity-based phase reconstruction algorithm. Our learned design results using two measurements demonstrate similar accuracy to Fourier ptychography with 69 measurements.**

Index Terms—**Phase imaging, unrolled network, physics-based, experimental design, illumination design.**

I. Introduction

QUANTITATIVE Phase Imaging (QPI) enables stain-free and label-free microscopy of transparent biological samples *in vitro* [1], [2]. Compared with coherent methods [3], [4], QPI techniques that use partially coherent light achieve higher spatial resolution, more light throughput, and reduced speckle artifacts. Phase contrast can be generated by interference [5], [6] or defocus [7]–[9]. More recently, *coded-illumination microscopy* [10]–[15] has been demonstrated as an accurate and inexpensive QPI scheme. To realize coded-illumination, we replace a commercial microscope's illumination unit with a light-emitting diode (LED) domed array (see Fig. 1) [16]. This flexible hardware platform has been used for various QPI applications including super-resolution [10], [11], [13], multi-contrast [12], [17], and 3D imaging [14], [15].

Coded-illumination microscopy uses intensity measurements with asymmetric source patterns [18] to retrieve 2D phase information. Quantitative Differential Phase Contrast [19]–[22] (qDPC) typically captures four measurements with rotated half-circle source patterns, from which the phase is computationally recovered using a partially coherent model. The performance of qDPC is predominantly determined by how the phase information is encoded in (via coded-illumination) and decoded from (via phase recovery) the intensity measurements.

The half-circle illumination designs of qDPC were derived analytically based on a Weak Object Approximation [20], [21], [23], [24] which linearizes the physics in order to make the inverse problem mathematically convenient. This linearized model enables one to derive a phase transfer function and analyze the spatial frequency coverage of any given source pattern [21], [22], [25], [26]. However, the non-linearity of the actual system makes it impossible to predict an optimal source design without knowing the sample's phase *a priori*. In addition, these types of analysis are inherently restricted to linear reconstruction algorithms and will not necessarily result in improved accuracy when the phase is retrieved via non-linear iterative methods.

Motivated by the success of deep learning [27] for image reconstruction problems [28]–[33], data-driven approaches have been adopted for learning coded-illumination patterns. For instance, researchers have used machine learning to maximize the phase contrast of each coded-illumination measurement [34], to improve accuracy on classification tasks [35], and to reconstruct phase [36]. All of these techniques learn the input-output relationship with a deep convolutional neural network (CNN) using training data. It is not straightforward to include the well-characterized system physics; hence, the CNN is required
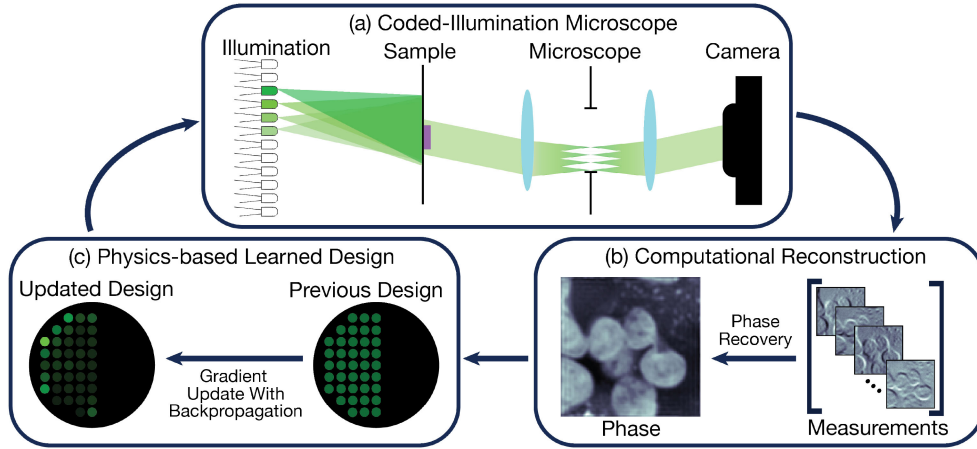
Fig. 1. Learning coded-illumination designs for quantitative phase imaging: (a) The LED-array microscope captures multiple intensity measurements with different coded-illumination source patterns. (b) The measurements are used to computationally reconstruct the sample's complex-field using an iterative phase recovery algorithm. (c) An optimization procedure for learning optimal coded-illumination patterns updates the illumination design.

to learn both the physical measurement formation and the phase reconstruction process. This task requires training of 10 s to 100 s of thousands of parameters and an immense number of training examples.

Here, we introduce a new data-driven approach to optimizing the source pattern design for coded-illumination phase retrieval by directly including both the system physics and the non-linear nature of a reconstruction algorithm in the learning process. Our approach *unrolls* the iterations of a generic non-linear reconstruction algorithm to construct an *unrolled network* [37]–[43]. Similar to CNNs, our *unrolled network* consists of several layers (one for each iteration); however, in our case each layer consists of well-specified operations to incorporate measurement formation and sparse regularization, instead of standard operations such as generic convolutions. The key benefits of our approach are:

- incorporation of the system physics and reconstruction non-linearities in the illumination design process.
- efficient parameterization of the unrolled network.
- incorporation of practical constraints.
- reduced number of training examples required.

We deploy our data-driven approach to learn improved coded-illumination patterns for phase reconstruction. Each layer of the unrolled network is parameterized by only a few variables (LED brightness values), enabling an efficient use of training data (<100 simulated training examples). We compare the QPI performance of our learned designs to previous work and demonstrate that our designs generalize well to the experimental setting with biological samples.

II. QUANTITATIVE PHASE IMAGING

qDPC recovers a sample's complex transmittance function from several coded-illumination measurements. The phase recovery optimization algorithm aims to minimize the Euclidean norm of the error between the measurements and the expected measurements formed with the current phase estimate. Using a gradient-based procedure, the phase estimate is iteratively updated until convergence. For a partially coherent source, the phase can be recovered with resolution up to twice the coherent diffraction limit. In this section, we describe the measurement formation process and phase recovery optimization.

A. System Modeling

A thin sample's transmission function can be approximated as a 2D complex function, $o(\mathbf{r}) = e^{j\phi(\mathbf{r}) - \mu(\mathbf{r})}$, characterized by its absorption, $\mu(\mathbf{r})$, and phase, $\phi(\mathbf{r}) = \frac{2\pi}{\lambda}\Delta n(\mathbf{r})d(\mathbf{r})$, where $\mathbf{r}$ are 2D spatial coordinates, λ is the wavelength of the illumination, $d(\mathbf{r})$ is the physical thickness of the sample, and $\Delta n(\mathbf{r})$ is the change in refractive index from the background. Intensity measurements, $y(\mathbf{r})$, of the sample are a non-linear function of $o(\mathbf{r})$, mathematically described by,

$$y(\mathbf{r}) = |p(\mathbf{r}) * (s(\mathbf{r}) \odot o(\mathbf{r}))|^2, \quad (1)$$

where $|\cdot|^2$ denotes squared absolute value, $*$ denotes convolution, $\odot$ denotes elementwise multiplication, $s(\mathbf{r})$ is the illumination's complex-field at the sample plane and $p(\mathbf{r})$ is the point spread function (PSF) of the microscope. The illumination from each LED is approximated as a tilted plane wave, $s(\mathbf{r}) = e^{\frac{j}{\lambda}\mathbf{u}_{pos}^T \mathbf{r}}$, with tilt angle, $\mathbf{u}_{pos}$, determined by the physical position of the LED relative the microscope [44].

Because the measured image in Eq. 1 is non-linear with respect to the sample's transmission function, recovering phase generally requires non-convex optimization. However, biological samples in closely index-matched fluid have a small *scatter-scatter* term. This means that a *weak object approximation* can be made; linearizing the measurement formation model such that phase recovery requires only a linear deconvolution of the measurements with their respective weak object transfer functions (WOTFs) [20]–[24]. Further, unstained biological samples are predominantly phase objects since they are only weakly absorbing (i.e. $\mu(\mathbf{r})$ is small). With these approximations, we can

express each intensity measurement as a linear system with contributions from the background and phase contrast. In Fourier space,

$$\widehat{y}(\mathbf{u}) \approx B\delta(\mathbf{u}) + ih(\mathbf{u})\widehat{\phi}(\mathbf{u}), \tag{2}$$

where $\widehat{\ }$ denotes Fourier transform, $\mathbf{u}$ are 2D spatial-frequency coordinates, B is the measurement's background energy concentrated at the DC and $h(\mathbf{u})$ is the phase WOTF. The phase WOTFs are a function of the source and the pupil distributions of the microscope [21]. For a single LED the WOTF is:

$$h^{(single)}(\mathbf{u}) = i(\widehat{p}(\mathbf{u}) \star \widehat{s}(\mathbf{u}) - \widehat{s}(\mathbf{u}) \star \widehat{p}(\mathbf{u})), \tag{3}$$

where $\star$ is the correlation operator, defined as $(x_1 \star x_2)(\mathbf{r}) = \int x_1(\widetilde{\mathbf{r}})x_2^*(\widetilde{\mathbf{r}} - \mathbf{r})d\widetilde{\mathbf{r}}$ for $\mathbf{r}$ in the domain of $\widehat{p}$ and $\widehat{s}$.

In [21], multiple LEDs are turned on simultaneously to increase signal-to-noise (SNR) and improve phase contrast. Because the fields generated by each LED's illumination are spatially incoherent with each other, the measurement from multiple LEDs will simply be the weighted sum of each LED's individual measurement, where the weights correspond to the LEDs' brightness values. The phase WOTF for illumination by multiple LEDs will also be the weighted sum of the single-LED phase WOTFs. Mathematically,

$$\widehat{y}^{(multi)}(\mathbf{u}) = \sum_{w \in \mathcal{W}} c_w \widehat{y}^{(single)}(\mathbf{u}) \tag{4}$$

$$h^{(multi)}(\mathbf{u}) = \sum_{w \in \mathcal{W}} c_w h_w^{(single)}(\mathbf{u}), \tag{5}$$

where $\mathcal{W}$ is the set of LEDs turned on and $c_w \geq 0$ are the LEDs' brightness values.

Following common practice [45], we discretize the 2D spatial distributions and format them as vectors (bold lower case) (e.g. $\widehat{\mathbf{h}}$ represents the transfer function's 2D spatial-frequency distribution and ϕ represents the 2D spatial phase distribution). The measurements[1] are described in Fourier space as $\widehat{\mathbf{y}} = \mathbf{A}\widehat{\phi}$ with system function $\mathbf{A} = diag(\widehat{\mathbf{h}})$.

Based on this model, we define $\mathbf{Y} \in \mathbb{R}^{M \times S}$ as the Fourier transform of S single LED measurements, $\widehat{\mathbf{y}}$, along the columns. Then, $\mathbf{C} \in \mathbb{R}^{S \times K}$ is defined as the S single-LED weights for each of K measurements, and $\mathbf{c}_k \in \mathbb{R}^S$ is the k^{th} column of $\mathbf{C}$. The product $\widehat{\mathbf{y}}_k = \mathbf{Y}\mathbf{c}_k$ simulates the k^{th} multiple-LED measurement. Similarly, we define $\mathbf{H} \in \mathbb{R}^{N \times S}$ as S single LED phase WOTFs, $\widehat{\mathbf{h}}$ along the columns, such that the product $\mathbf{A}_k = diag(\mathbf{H}\mathbf{c}_k)$ gives the corresponding multiple-LED phase WOTF for the k^{th} measurement.

B. Phase Recovery

Phase recovery using the forward model in Section II-A can be formulated as a regularized linear inverse problem,

$$\widehat{\phi}^\star = \mathcal{R}((\widehat{\mathbf{y}}_k)_{k=1}^K, \mathcal{P}(\cdot)) \tag{6}$$

$$= \arg\min_{\widehat{\phi}} \frac{1}{2K} \sum_{k=1}^K \|\widehat{\mathbf{y}}_k - \mathbf{A}_k\widehat{\phi}\|_2^2 + \mathcal{P}(\widehat{\phi}), \tag{7}$$

[1]In practice, $\mathbf{y}$ typically refers to the so-called flattened image, where the background energy in (2) is removed via background subtraction.

Algorithm 1: Accelerated Proximal Gradient Descent (APGD) for Phase Recovery.

1: **procedure** APGD($(\widehat{\mathbf{y}}_k)_{k=1}^K, N, \alpha, \mathcal{P}(\cdot)$)
2: $\quad \widehat{\phi}^{(0)} = \mathbf{0}, \widehat{\phi}^{(-1)} = \mathbf{0}$
3: $\quad$ **for** $n \in \{1 \ldots N\}$ **do**
4: $\quad\quad \mathbf{s}^{(n)} \leftarrow \mu^{(n)}\widehat{\phi}^{(n-1)} + (1 - \mu^{(n)})\widehat{\phi}^{(n-2)}$
5: $\quad\quad \mathbf{z}^{(n)} \leftarrow \mathbf{s}^{(n)} - \frac{\alpha}{K} \sum_{k=1}^K (-\mathbf{A}_k^H)$
$\quad\quad (\widehat{\mathbf{y}}_k - \mathbf{A}_k\mathbf{s}^{(n)})$
6: $\quad\quad \widehat{\phi}^{(n)} \leftarrow \text{prox}_{\alpha\mathcal{P}}(\mathbf{z}^{(n)})$
7: $\quad$ **end for**
8: $\quad$ **return** $\widehat{\phi}^{(N)}$
9: **end procedure**

---

where $\phi^\star$ is the recovered phase, K is the number of measurements acquired, $\widehat{\mathbf{y}}_k$ is the Fourier transform of the k^{th} measurement and $\mathcal{P}(\cdot)$ is a user-chosen regularizer. We solve this optimization problem efficiently using the accelerated proximal gradient descent (APGD) algorithm by iteratively applying an acceleration update, a gradient update and a proximal update [46], [47]. The algorithm is detailed in Algorithm 1, where α is the gradient step size, N is the number of iterations, $\mathbf{s}$ and $\mathbf{z}$ are intermediate variables, $\mu^{(n)}$ is the acceleration parameter derived by the recursion, $\mu^{(n)} = \frac{1+\sqrt{1+4\mu^{(n-1),2}}}{2}$ [47], and $\text{prox}_{\mathcal{P}}(\cdot)$ is the proximal operator corresponding to the user-chosen regularizer $\mathcal{P}(\cdot)$ [46].

III. PHYSICS-BASED LEARNED DESIGN

Given the phase recovery algorithm in Section II-B, we now describe our main contribution of learning the coded-illumination designs for a given reconstruction algorithm and training set.

A. Unrolled Physics-Based Network

Traditionally, DNNs contain many layers of weighted linear mixtures and non-linear activation functions [27]. Here, we consider specific linear functions which capture the system physics of measurement formation and specific non-linear activation functions which promote sparsity [37], [38]. Starting from Algorithm 1, we treat each iteration as a layer such that when unrolled they form a network of N layers, denoted $\mathcal{R}$ (Fig. 2). Each layer of $\mathcal{R}$ contains a module for each of the iterative algorithm's updates (i.e. an acceleration module, a gradient module (incorporates system physics), and a proximal module (incorporates sparsity)). The regularization and step size parameters specified for Algorithm 1 are fixed. The network's inputs comprise $(\widehat{\mathbf{y}}_k)_{k=1}^K$ and the network's output is $\widehat{\phi}^{(N)}$. The design parameters of the network, which will be learned, govern the relative brightness of the LEDs and are incorporated in the measurement formation and the system WOTFs.

B. Learning Objective

Our learning objective is to minimize the phase reconstruction error of the training data over the space of possible LED
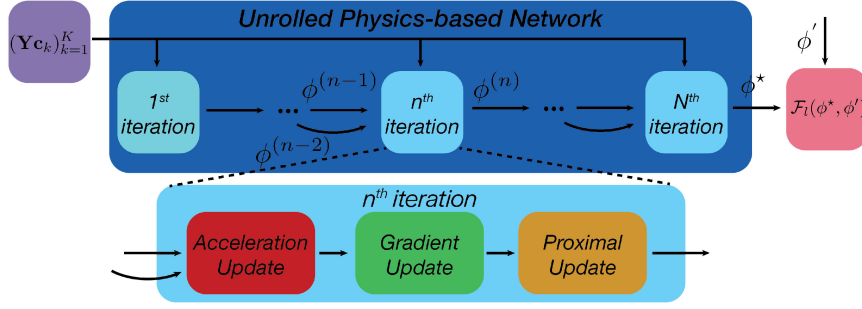
Fig. 2. Unrolled physics-based network: feed-forward schematic for the unrolled accelerated proximal gradient descent (APGD) network for N iterations (dark blue box). The network takes intensity measurements, $\mathbf{y}_k$, parameterized by the coded-illumination design, $\mathbf{c}_k$, as input and outputs the reconstructed phase, $\phi^\star$. Finally, the output is compared with the ground truth phase, ϕ', using a user-chosen loss function, $\mathcal{F}_l$ (pink box). The inset into a single (n^{th}) iteration (light blue box) shows each iteration's three steps: acceleration update, gradient update, and proximal update.

configurations, subject to constraints that enforce physical feasibility and eliminate degenerate and trivial solutions:

$$\mathbf{C}^\star = \arg\min_{\mathbf{C}} \mathcal{F}(\mathbf{C}) \qquad (8)$$

$$\text{s.t. } \mathbf{c}_k \geq 0 \qquad \text{(non-negativity)} \qquad (9)$$

$$\|\mathbf{c}_k\|_1 = 1 \qquad \text{(scale)} \qquad (10)$$

$$\mathbf{m}_k \odot \mathbf{c}_k = \mathbf{0} \qquad \text{(geometric)}$$

$$\forall k \in \{1 \ldots K\}, \qquad (11)$$

where,

$$\mathcal{F}(C) = \frac{1}{L} \sum_{l=1}^{L} \mathcal{F}_l(\mathbf{C}) \qquad (12)$$

$$= \frac{1}{2L} \sum_{l=1}^{L} \|\mathcal{R}((\mathbf{Y}_l \mathbf{c}_k)_{k=1}^K) - \widehat{\phi}_l'\|_2^2. \qquad (13)$$

Here, $(\mathbf{Y}_l, \phi_l')_{l=1}^L$ are L training pairs for which $\mathbf{Y}_l$ is a matrix of the Fourier transform of single-LED measurements for the l^{th} sample with optical phase, ϕ_l'. $\odot$ is the elementwise product operator, $\mathbf{m}_k$ is a geometric constraint mask for the k^{th} measurement, and $\mathbf{0}$ is the null vector.

The non-negativity constraint (Eq. 9) prevents non-physical solutions by enforcing the brightness of each LED to be greater than or equal to zero. This is enforced by projecting the parameters onto the set of non-negative real numbers. The scale constraint (Eq. 10) enforces that each coded-illumination design must have weights with sum equal to 1, in order to eliminate arbitrary scalings of the same design. This is enforced by scaling the parameters for each measurement such that their sum is one. The geometric constraint (Eq. 11) enforces that the coded-illumination designs do not use conjugate-symmetric LED pairs to illuminate the sample within the same measurement, since these would also result in degenerate solutions (e.g. two symmetric LEDs produce opposite phase contrast measurements that would cancel each other out). To prevent this, we force the source patterns for each measurement to reside within only one of the major semi-circle sets (e.g. top, bottom, left, right). This

---

Algorithm 2: Physics-Based Learned Design Algorithm.

1:    **procedure** PBLD($(\mathbf{Y}_l, \phi_l')_{l=0}^L, \mathbf{C}, \gamma, T$)
2:       **for** $t \in \{0 \ldots T\}$ **do**   ▷ Gradient descent loop
3:          **for** $l \in \{1 \ldots L\}$ **do**   ▷ Training data loop
4:             $r_l \leftarrow \mathcal{R}((\mathbf{Y}_l \mathbf{c}_k)_{k=1}^K) - \widehat{\phi}_l'$
5:             $\mathbf{G}_l \leftarrow BackPropagation(r_l)$
6:          **end for**
7:          $\mathbf{C}^{(t+1)} \leftarrow \mathcal{B}(\mathbf{C}^{(t)} - \frac{\gamma}{L} \sum_{l=1}^L \mathbf{G}_l)$
8:          $\mathbf{C}^{(t+1)} \leftarrow \beta^{(t)} \mathbf{C}^{(t+1)} + (1 - \beta^{(t)}) \mathbf{C}^{(t)}$
9:       **end for**
10:      **return** $\mathbf{C}^{(T)}$
11: **end procedure**

---

constraint is enforced by setting the LED brightnesses outside the allowed semi-circle to zero.

We solve Eq. 8 iteratively via accelerated projected gradient descent (Algorithm 2). At each iteration, the coded-illumination design for each measurement is updated with the analytical gradient, projected onto the constraints (denoted by $\mathcal{B}(\cdot)$) and updated again with a contribution from the previous iteration (weighted by $\beta^{(t)}$). $\mathcal{B}(\cdot)$ enforces the constraints in the following order: non-negativity, geometric, and scale.

C. Gradient Update

The gradient of the loss function (Eq. 8) with respect to the design parameters has contributions at every layer of the unrolled network through both the measurement terms, $\widehat{\mathbf{y}}_k$, and the phase WOTF terms, $\mathbf{A}_k$, for each measurement $k \in \{1 \ldots K\}$. Here, we outline our algorithm for updating the coded-illumination design weights via a two-step procedure: backpropagating the error from layer-to-layer and computing each layer's gradient contribution. For simplicity, we outline the gradient update for only a single training example, l, as the gradient for all the training examples is the sum of their individual gradients.

Unlike pure gradient descent, where each iteration's estimate only depends on the previous', accelerated methods like

---

Algorithm 3: Gradient Update for Single Training Example.

1:     **procedure** Backpropagation(BP)($\mathbf{r}^{(N)}$)
2:        **for** $n \in \{N \dots 0\}$ **do**
3:           $\mathbf{b}^{(n)} \leftarrow \frac{\partial \widehat{\phi}}{\partial \mathbf{z}} \mathbf{r}^{(n)}$
4:           $\mathbf{v}^{(n)} \leftarrow (I - \frac{\alpha}{K} \sum_{k=1}^{K} \mathbf{A}_k^H \mathbf{A}_k) \mathbf{b}^{(n)}$
5:           $\mathbf{r}^{(n-1)} \leftarrow \mu^{(n)} \mathbf{v}^{(n)} + (1 - \mu^{(n+1)}) \mathbf{v}^{(n+1)}$
6:           $\mathbf{Q}^{(n)} \leftarrow \frac{\alpha}{K} \sum_{k=1}^{K} \left(\frac{\partial \mathbf{A}_k^H \widehat{\mathbf{y}}_k}{\partial \mathbf{C}} - \frac{\partial \mathbf{A}_k^H \mathbf{A}_k}{\partial \mathbf{C}} s^{(n-1)} \right)$
              $\mathbf{b}^{(n)}$
7:        **end for**
8:        **return** $\sum_{n=0}^{N} \mathbf{Q}^{(n)}$
9:     **end procedure**

Algorithm 1 linearly combine the previous two iteration's estimates to improve convergence. As a consequence, backpropagating error from layer-to-layer requires contributions from two successive layers. Specifically, we compute the error at all N layers with the recursive relation,

$$\frac{\partial \mathcal{F}_l}{\partial \widehat{\phi}^{(n-2)}} = \frac{\partial \mathbf{s}^{(n)}}{\partial \widehat{\phi}^{(n-2)}} \frac{\partial \mathbf{z}^{(n)}}{\partial \mathbf{s}^{(n)}} \frac{\partial \widehat{\phi}^{(n)}}{\partial \mathbf{z}^{(n)}} \frac{\partial \mathcal{F}_l}{\partial \widehat{\phi}^{(n)}}$$

$$+ \frac{\partial \mathbf{s}^{(n-1)}}{\partial \widehat{\phi}^{(n-2)}} \frac{\partial \mathbf{z}^{(n-1)}}{\partial \mathbf{s}^{(n-1)}} \frac{\partial \widehat{\phi}^{(n-1)}}{\partial \mathbf{z}^{(n-1)}} \frac{\partial \mathcal{F}_l}{\partial \widehat{\phi}^{(n-1)}}, \quad (14)$$

where each partial gradient constitutes a single step in Algorithm 1 (fully derived in the supplement).

With the backpropagated error at each layer, we compute the gradient of the loss function with respect to $\mathbf{C}$ as,

$$\nabla_{\mathbf{C}} \mathcal{F}_l(\mathbf{C}) = \sum_{n=0}^{N} \mathbf{Q}^{(n)}, \quad (15)$$

for which,

$$\mathbf{Q}^{(n)} = \frac{\alpha}{K} \sum_{k=1}^{K} \left(\frac{\partial \mathbf{A}_k^H \widehat{\mathbf{y}}_k}{\partial \mathbf{C}} - \frac{\partial \mathbf{A}_k^H \mathbf{A}_k}{\partial \mathbf{C}} \mathbf{s}^{(n-1)} \right) \frac{\partial \widehat{\phi}^{(n)}}{\partial \mathbf{z}^{(n)}} \frac{\partial \mathcal{F}_l}{\partial \widehat{\phi}^{(n)}}. \quad (16)$$

Here, $(\partial \widehat{\phi}^{(n)} / \partial \mathbf{z}^{(n)})$ backpropagates the error through the proximal operator and other partials with respect to $\mathbf{C}$ relate the backpropagated error at each layer to the changes in $\mathbf{C}$. Derivations of these partial gradients are in the supplementary material. In Algorithm 3, we unite these two steps to form a recursive algorithm which efficiently computes the analytic gradient for a single training example. Alternatively, general purpose auto-differentiation included in learning libraries (e.g. PyTorch, TensorFlow) can be used to perform the gradient updates.

IV. RESULTS

Our proposed method learns the coded-illumination design for a given reconstruction and training set (Fig. 3(b)), yet up to this point we have not detailed specific parameters of our phase reconstruction. In our results, we set the parameters of our reconstruction algorithm (Algorithm 1) to have a fixed CPU time by fixing the number of iterations at $N = 40$ and the step size
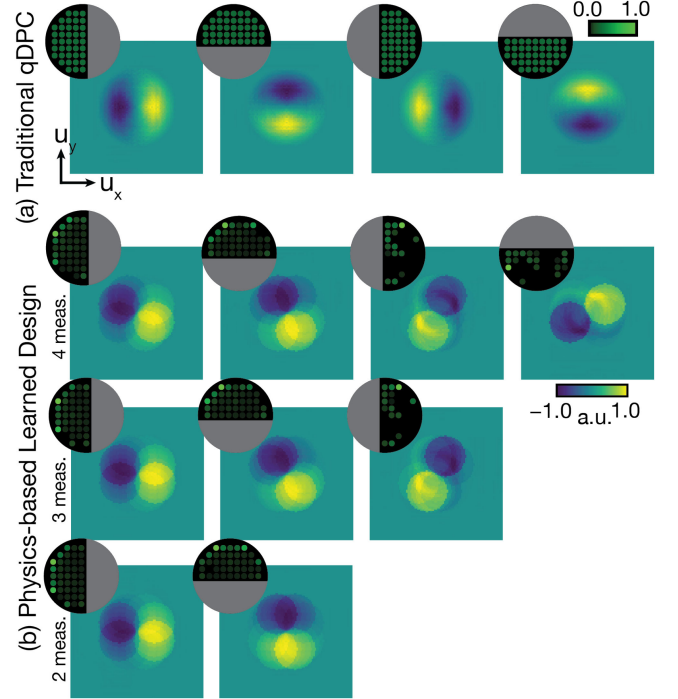

Fig. 3. Coded-illumination designs and their corresponding phase weak object transfer functions (WOTFs) for: (a) Traditional qDPC and (b) learned designs for the case where 4, 3, or 2 measurements are allowed for each phase reconstruction. The illumination source patterns are in the upper left corners, with gray semi-circles denoting where the LEDs are constrained to be "off."

to $\alpha = 0.2$ (see supplement for parameter analysis). In addition, the regularization term, $\mathcal{P}(\phi)$, has been defined generally (e.g. ℓ_1 penalty, total variation (TV) penalty [48], BM3D [49]). Here, we choose to enforce TV-based sparsity:

$$\mathcal{P}(\phi) = \tau \sum_i \| D_i \phi \|_1, \quad (17)$$

where $\tau = 1e^{-3}$ is set to trade off the TV cost with the data consistency cost and D_i is the first-order difference operator along the i^{th} image dimension. We efficiently implement the proximal operator of Eq. 17 in closed form via parallel proximal method [43], [50], [51] (details in supplement).

A. Learning

To train our coded-illumination design parameters using Algorithm 2, we generate a dataset of 100 examples (90 for training, 10 for testing). Each example contains ground truth phase from a small region (95×95 pixels) of a larger image and 69 simulated single LED measurements (using Eq. 1). The LEDs are uniformly spaced within a circle such that each single-LED intensity measurement is a brightfield measurement. The physical system parameters used to generate the phase WOTFs and simulate the training data measurements are $\lambda = 0.532 \ \mu m$, pixel pitch $= 6.5 \ \mu m$, magnification $= 20 \times$, and $NA_{obj} = 0.25$. To train, we use ℓ_2 cost between reconstructed phase and ground truth phase as our loss function and approximate the full gradient of Eq. 8 with a batch gradient from random batches of 10% of the training pairs at each iteration.
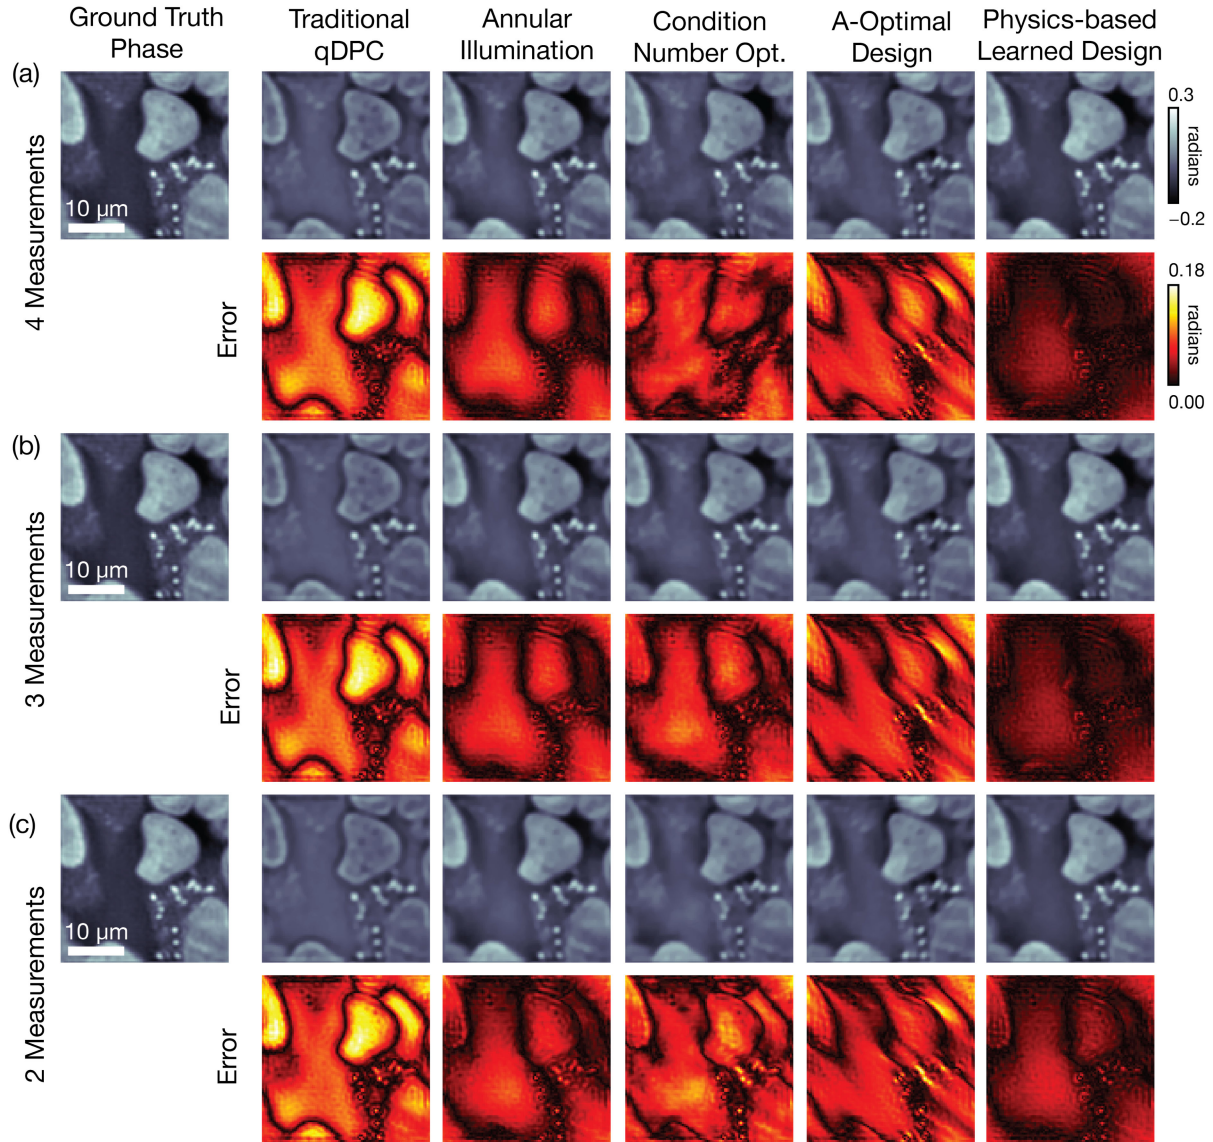
Fig. 4. Phase reconstruction results using simulated measurements with different coded-illumination designs. We compare results from: traditional qDPC (half-circles), annular illumination, condition number optimization, A-optimal design, and our proposed physics-based learned designs. We show results for the cases of (a) four, (b) three, and (c) two measurements allowed for each phase reconstruction. Absolute error maps are shown below each reconstruction.

We use a learning rate of $\gamma = 1e^{-2}$ (training and testing convergence curves are provided in the supplement). The training is performed on a multi-core CPU (Dual-socket Intel Xeon E5 Processor @ 2.1 GHz with 64 cores and 504 GB of RAM) and batch updates are computed in parallel with each training example on a single core. Each batch update takes ~6 seconds. 200 updates are performed, resulting in a total training time of 20 minutes.

B. Analysis

Traditional qDPC uses 4 measurements to adequately cover frequency space. Our learned designs are more efficient and may require fewer measurements; hence, we show learned designs for the cases of 4, 3 and 2 measurements. The designs and their corresponding phase WOTFs are shown in Fig. 3.

Comparing our learned designs with previous work, Fig. 4 shows the phase reconstruction for a single simulated test

example using 4, 3 and 2 measurements. The ground truth phase is compared with the phase reconstructed using traditional qDPC designs [21], annular illumination designs [21], condition number optimized designs [52], A-optimal designs [53], and our physics-based learned designs. Table I reports the peak SNR (PSNR) statistics (mean and standard deviation) for the phase reconstructions from $\mathcal{R}$ evaluated on our set of testing examples. Our learned designs give significant improvement over other designs, recovering both the high and low frequencies more accurately. The reduction in performance for learned design with 2 measurements (as compared to 3 and 4 measurements) is due to reduced sensitivity to low frequencies.

Comparing varying depth networks, in Table II we report the PSNR statistics for the phase reconstructions from $\mathcal{R}$ evaluated on the set of testing examples using learned designs for networks with 10, 40, and 100 unrolled iterations with fixed step size, $\alpha = 0.2$, and regularization parameter, $\tau = 1e^{-3}$. If too few iterations

TABLE I
PSNR RESULTS: AVERAGE AND STANDARD DEVIATION PSNR (DB) OF PHASE RECONSTRUCTIONS FROM THE SIMULATED TESTING EXAMPLES USING DIFFERENT
ILLUMINATION SCHEMES AND DIFFERENT NUMBERS OF MEASUREMENTS FACTOR FORMAT: MEAN $\pm$ STD

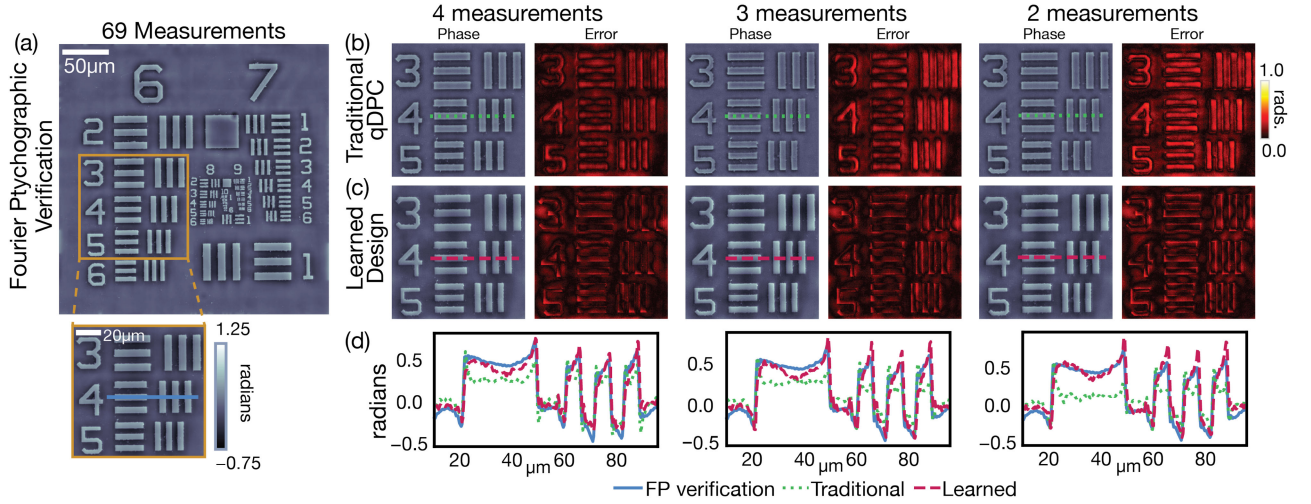| # Meas.	Random illumination	Traditional qDPC	Annular illumination	Cond. Number optimization	A-optimal design	Physics-based learned design
4	12.30 ± 2.12	15.67 ± 2.19	20.40 ± 2.09	20.37 ± 2.41	17.94 ± 2.54	$\mathbf{28.46 \pm 2.50}$
3	12.33 ± 2.12	15.28 ± 2.18	20.44 ± 2.26	19.33 ± 2.03	18.05 ± 2.59	$\mathbf{28.04 \pm 2.59}$
2	12.25 ± 2.12	14.87 ± 2.23	20.21 ± 2.24	17.19 ± 2.28	18.08 ± 2.64	$\mathbf{23.73 \pm 2.18}$

Fig. 5.   USAF phase target reconstructions: Experimental comparison between phase results with (a) Fourier Ptychography (FP) using 69 images, (b) traditional qDPC, and (c) learned designs, for the case of 4, 3, and 2 measurements. Error maps show the difference from the FP reconstruction. (d) Cross-sections show that phase from our learned designs (long-dashed red) is closer to that of FP (solid blue) than traditional qDPC (short-dashed green).

TABLE II
PSNR RESULTS: AVERAGE AND STANDARD DEVIATION PSNR (DB) OF PHASE
RECONSTRUCTIONS FROM THE SIMULATED TESTING EXAMPLES USING
LEARNED DESIGN FOR DIFFERENT NUMBERS OF UNROLLED ITERATIONS
FACTOR FORMAT: MEAN $\pm$ STD

# of unrolled iterations	10	40	100
Mean $\pm$ Std. (dB)	22.53 ± 2.29	28.48 ± 2.50	27.39 ± 1.92

are unrolled, the network cannot fully reconstruct the phase, resulting in lower mean PSNR. As more iterations are unrolled, the regularization biases the solution and slightly reduces the mean PSNR.

C. Experimental Validation

To demonstrate that our learned designs generalize well in the experimental setting, we implement our method on an LED array microscope. A commercial Nikon TE300 microscope is equipped with a custom quasi-Dome [16] illumination system (581 programmable RGB LEDs: $\lambda_R = 625$ nm, $\lambda_G = 532$ nm, $\lambda_B = 450$ nm) and a PCO.edge 5.5 monochrome camera (2560×2160, 6.5 μm pixel pitch, 16 bit). We image two samples: a USAF phase target (Benchmark Technologies) and fixed 3T3 mouse fibroblast cells (prepared as detailed in the supplement). In order to validate our method, we compare results against phase experimentally estimated via pupil-corrected Fourier Ptychography (FP) [13], [44], [54] with equivalent

resolution. FP is expected to have good accuracy, since it uses significantly more measurements (69 single-LED measurements) and a non-linear reconstruction process.

Using the USAF target, we compare phase reconstructions from FP with traditional qDPC and our learned design measurements (Fig. 5). Traditional qDPC reconstructions consistently under-estimate the phase values. However, phase reconstructions using our learned design measurements are similar to phase estimated with FP. As the number of measurements is reduced, the performance quality of the reconstruction using traditional qDPC degrades, while the reconstruction using the learned design remains accurate.

To demonstrate our method with live biological samples, we repeated the experiments with 3T3 mouse fibroblast cells. Fig. 6 shows that phase reconstructions from traditional qDPC again consistently under-estimate phase values, while phase reconstructions using learned design measurements match the phase estimated with FP well.

V. DISCUSSION

Our proposed experimental design method efficiently learns the coded-illumination designs by incorporating both the system physics and the non-linear nature of iterative phase recovery. Learned designs with only 2 measurements can efficiently reconstruct phase with quality similar to Fourier Ptychography (69 measurements) and better than qDPC (4 measurements), giving an improvement in temporal resolution by a factor of
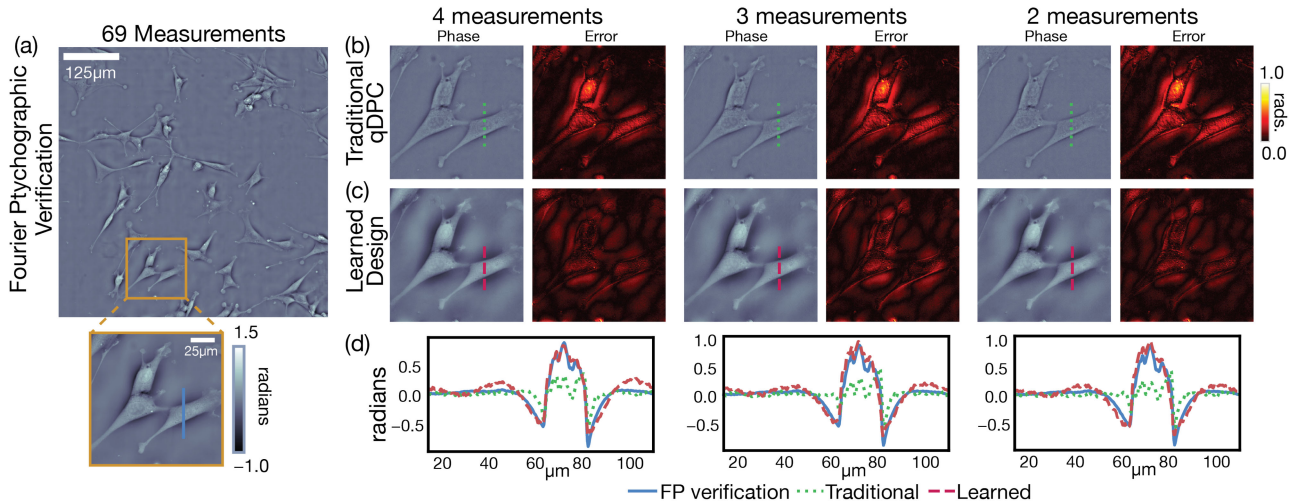
Fig. 6. 3T3 mouse fibroblast cells reconstructions: experimental comparison between phase results with (a) Fourier Ptychography (FP) using 69 measurements, (b) traditional qDPC, and (c) learned designs, for the case of 4, 3, and 2 measurements. Error maps show the difference from the FP reconstruction. (d) Cross-sections show that phase from our learned designs (long-dashed red) is closer to that of FP (solid blue) than traditional qDPC (short-dashed green).

$2\times$ over traditional qDPC and far fewer than FP. Additionally, we demonstrate (Table I) that the performance of our designs on a set of testing examples is superior to previously-proposed coded-illumination designs. Visually, our learned design reconstructions closely resemble the ground truth phase, with both low-frequency and high-frequency information accurately recovered.

By parameterizing our learning problem with only a few weights per measurement, our method can efficiently learn an experimental design with a small simulated dataset. This enables fast training and reduces computing requirements significantly. Obtaining large experimental datasets for training may be difficult in microscopy, so it is important that our method can be trained on simulated data only. Experimental results in Section IV-C show similar quality to simulated results, with both using the designs learned from simulated data only.

Phase recovery with the learned designs' measurements are trained with a given number of reconstruction iterations (e.g. determined by a CPU budget). This makes our method particularly well-suited for real-time processing. qDPC can also be implemented in real-time, but limiting the compute time for the inverse problem (by restricting the number of iterations) limits convergence and causes low-frequency artifacts. Our learned designs incorporate the number of iterations (and hence processing time) into the design process, producing high-quality phase reconstructions within a reasonable compute time. While one can reduce the number of iterations, if too few iterations are unrolled the accuracy of the model inversion decreases (Table II).

Finally, rather than use a fixed step size and regularization parameter (as outlined in Section IV), these parameters can be jointly learned with the illumination patterns to optimize the whole system. For given noise statistics the regularization parameter could be learned; however, it would perform sub optimally for different noise statistics and would require retraining. Future systems should learn regularization parameters that can be adapted post training to account for variable noise levels similar to in [55].

VI. OUTLOOK

Our method is general to the problem of experimental design. Similar to QPI, many fields (e.g. Magnetic resonance imaging (MRI), fluorescence microscopy) use physics-based non-linear iterative reconstruction techniques to achieve state-of-the-art performance. With the correct model parameterization and physically-relevant constraints, our method could be applied to learn optimal design for these applications (e.g. undersampling patterns for compressed sensing MRI [56], PSFs for fluorescence microscopy [57]).

Requirements for applying our method are simple: the reconstruction algorithm's updates must be differentiable (e.g. gradient update and proximal update) so that analytic gradients of the learning loss can be computed with respect to the design parameters. Of practical importance, the proximal operator of the regularizer should be chosen so that it has a closed form. While this is not a strict requirement, if the operator itself requires an additional iterative optimization, error will have to be backpropagated through an excessive number of iterations. Here, we choose to penalize anisotropic TV, whose proximal operator can be approximated in closed form [51]. Further, including an acceleration update improves the convergence of gradient-based reconstructions. As a result, the *unrolled network* can be constructed using fewer layers than its unaccelerated counterpart. This will reduce both computation time and training requirements.

VII. CONCLUSION

We have presented a general framework for incorporating the non-linearities of regularized reconstruction and known system physics to learn optimal experimental design. Here, we have applied this method to learn coded-illumination source designs for quantitative phase recovery. Our coded-illumination designs can improve the temporal resolution of the acquisition and enable real-time processing, while maintaining high accuracy. We demonstrated here that our learned designs achieve

high-quality reconstructions experimentally without the need for retraining.

ACKNOWLEDGMENT

REFERENCES

[1] G. Popescu, *Quantitative Phase Imaging of Cells and Tissues*. New York, NY, USA: McGraw-Hill, Mar. 2011.

[2] M. Mir, B. Bhaduri, R. Wang, R. Zhu, and G. Popescu, *Quantitative Phase Imaging*. Amsterdam, The Netherlands: Elsevier, Jul. 2012.

[3] B. Rappaz, B. Breton, E. Shaffer, and G. Turcatti, "Digital holographic microscopy: A quantitative label-free microscopy technique for phenotypic screening," *Combinatorial Chem. High Throughput Screening*, vol. 17, no. 1, pp. 80–88, Jan. 2014.

[4] E. Cuche, P. Marquet, and C. Depeursinge, "Simultaneous amplitude-contrast and quantitative phase-contrast microscopy by numerical reconstruction of Fresnel off-axis holograms," *Appl. Opt.*, vol. 38, no. 34, pp. 6994–7001, Dec. 1999.

[5] B. Bhaduri, H. V. Pham, M. Mir, and G. Popescu, "Diffraction phase microscopy with white light," vol. 37, no. 6, pp. 1094–1096, Mar. 2012.

[6] Z. Wang *et al.*, "Spatial light interference microscopy (SLIM)," *Opt. Express*, vol. 19, no. 2, pp. 1016–1026, Jan. 2011.

[7] T. E. Gureyev, A. Roberts, and K. A. Nugent, "Partially coherent fields, the transport-of-intensity equation, and phase uniqueness," *J. Opt. Soc. Amer. A*, vol. 12, no. 9, pp. 1942–1946, Sep. 1995.

[8] N. Streibl, "Phase imaging by the transport equations of intensity," *Opt. Commun.*, vol. 49, no. 1, pp. 6–10, Feb. 1984.

[9] L. Waller, L. Tian, and G. Barbastathis, "Transport of intensity phase-amplitude imaging with higher order intensity derivatives," *Opt. Express*, vol. 18, no. 12, pp. 12 552–12 561, Jun. 2010.

[10] L. Tian, Z. Liu, L.-H. Yeh, M. Chen, J. Zhong, and L. Waller, "Computational illumination for high-speed in vitro Fourier ptychographic microscopy," *Optica*, vol. 2, no. 10, pp. 904–908, Oct. 2015.

[11] G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier ptychographic microscopy," *Nature Photon.*, vol. 7, no. 9, pp. 739–745, Jul. 2013.

[12] G. Zheng, C. Kolner, and C. Yang, "Microscopy refocusing and dark-field imaging by using a simple LED array," *Opt. Lett.*, vol. 36, no. 20, pp. 3987–3989, Oct. 2011.

[13] L. Tian, X. Li, K. Ramchandran, and L. Waller, "Multiplexed coded illumination for Fourier ptychography with an LED array microscope," *Biomed. Opt. Express*, vol. 5, no. 7, pp. 1–14, Jun. 2014.

[14] L. Tian and L. Waller, "3D intensity and phase imaging from light field measurements in an LED array microscope," *Optica*, vol. 2, no. 2, pp. 104–111, Feb. 2015.

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

[16] Z. Phillips, R. Eckert, and L. Waller, "Quasi-dome: A self-calibrated high-NA LED illuminator for Fourier ptychography," presented at the Imaging Systems and Applications, San Francisco, CA, USA, Jun. 2017, paper IW4E–5.

[17] Z. Liu, L. Tian, S. Liu, and L. Waller, "Real-time brightfield, darkfield, and phase contrast imaging in a light-emitting diode array microscope," *J. Biomed. Opt.*, vol. 19, no. 10, p. 106002, Oct. 2014.

[18] B. Kachar, "Asymmetric illumination contrast: A method of image formation for video light microscopy," *Science*, vol. 227, no. 4688, pp. 766–768, Feb. 1985.

[19] D. K. Hamilton and C. J. R. Sheppard, "Differential phase contrast in scanning optical microscopy," *J. Microsc.*, vol. 133, no. 1, pp. 27–39, Jan. 1984.

[20] S. B. Mehta and C. J. R. Sheppard, "Quantitative phase-gradient imaging at high resolution with asymmetric illumination-based differential phase contrast," *Opt. Lett.*, vol. 34, no. 13, pp. 1924–1926, Jul. 2009.

[21] L. Tian and L. Waller, "Quantitative differential phase contrast imaging in an LED array microscope," *Opt. Express*, vol. 23, no. 9, pp. 11 394–11 403, May 2015.

[22] R. A. Claus, P. P. Naulleau, A. R. Neureuther, and L. Waller, "Quantitative phase retrieval with arbitrary pupil and illumination," *Opt. Express*, vol. 23, no. 20, pp. 26 672–26 682, Oct. 2015.

[23] D. K. Hamilton, C. J. R. Sheppard, and T. Wilson, "Improved imaging of phase gradients in scanning optical microscopy," *J. Microsc.*, vol. 135, no. 3, pp. 275–286, Sep. 1984.

[24] N. Streibl, "Three-dimensional imaging by a microscope," *J. Opt. Soc. Amer. A*, vol. 2, no. 2, pp. 121–127, Feb. 1985.

[25] J. Li, Q. Chen, J. Zhang, Y. Zhang, L. Lu, and C. Zuo, "Efficient quantitative phase microscopy using programmable annular LED illumination," *Biomed. Opt. Express*, vol. 8, no. 10, pp. 4687–4705, Oct. 2017.

[26] Y.-Z. Lin, K.-Y. Huang, and Y. Luo, "Quantitative differential phase contrast imaging at high resolution with radially asymmetric illumination," *Opt. Lett.*, vol. 43, no. 12, pp. 2973–2976, Jun. 2018.

[27] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

[28] 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.

[29] S. Wang *et al.*, "Accelerating magnetic resonance imaging via deep learning," in *Proc. IEEE 13th Int. Symp. Biomed. Imag.*, Apr. 2016, pp. 514–517.

[30] Y. Rivenson, Y. Zhang, H. Günaydın, D. Teng, and A. Ozcan, "Phase recovery and holographic image reconstruction using deep learning in neural networks," *Nature Light Sci. Appl.*, vol. 7, no. 2, Feb. 2018.

[31] A. Sinha, J. Lee, S. Li, and G. Barbastathis, "Lensless computational imaging through deep learning," *Optica*, vol. 4, no. 9, pp. 1117–1125, Sep. 2017.

[32] Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan, "Deep learning microscopy," *Optica*, vol. 4, no. 11, pp. 1437–1443, Nov. 2017.

[33] T. Nguyen, Y. Xue, Y. Li, L. Tian, and G. Nehmetallah, "Deep learning approach for Fourier ptychography microscopy," *Opt. Express*, vol. 26, no. 20, pp. 26470–26484, Oct. 2018.

[34] B. Diederich, R. Wartmann, H. Schadwinkel, and R. Heintzmann, "Using machine-learning to optimize phase contrast in a low-cost cellphone microscope," *PLoS ONE*, vol. 13, no. 3, pp. 1–20, Mar. 2018.

[35] R. Horstmeyer, R. Chen, B. Kappes, and B. Judkewitz, "Convolutional neural networks that teach microscopes how to image," Sep. 2017, arXiv:1709.07223.

[36] A. Robey and V. Ganapati, "Optimal physical preprocessing for example-based super-resolution," *Opt. Express*, vol. 26, no. 24, pp. 31 333–31 350, Nov. 2018.

[37] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. 27th Int. Conf. Int. Conf. Mach. Learn.*, Jun. 2010, pp. 399–406.

[38] B. Xin, Y. Wang, W. Gao, D. Wipf, and B. Wang, "Maximal sparsity with deep networks?" in *Proc. Advances Neural Inf. Process. Syst.*, 2016, pp. 4340–4348.

[39] J. Sun *et al.*, "Deep ADMM-Net for compressive sensing MRI," in *Proc. Advances Neural Inf. Process. Syst.*, 2016, pp. 10–18.

[40] K. Hammernik *et al.*, "Learning a variational network for reconstruction of accelerated MRI data," *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3055–3071, Nov. 2017.

[41] S. Diamond, V. Sitzmann, F. Heide, and G. Wetzstein, "Unrolled optimization with deep priors," May 2017, arXiv:1705.08041.

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

[43] 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.

[44] G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier ptychographic microscopy," *Nature Photon.*, vol. 7, no. 9, pp. 739–745, Jul. 2013.

[45] E. Bostan, U. S. Kamilov, M. Nilchian, and M. Unser, "Sparse stochastic processes and discretization of linear inverse problems," *IEEE Trans. Image Process.*, vol. 22, no. 7, pp. 2699–2710, Jul. 2013.

[46] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, Aug. 2014.

[47] 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, Jan. 2009.

[48] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization method for total variation-based image restoration," *Multiscale Model. Simul.*, vol. 4, no. 2, pp. 460–489, 2005.

[49] 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. 8, pp. 2080–2095, Aug. 2007.

[50] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-Point Algorithms For Inverse Problems in Science and Engineering*. New York, NY, USA: Springer, 2011, pp. 185–212.

[51] U. S. Kamilov, "A parallel proximal algorithm for anisotropic total variation minimization," *IEEE Trans. Image Process.*, vol. 26, no. 2, pp. 539–548, Feb. 2017.

[52] P. Marechal and J. Ye, "Optimizing condition numbers," *SIAM J. Optim.*, vol. 20, no. 2, pp. 935–947, 2009.

[53] C. S. Wong and J. C. Masaro, "A-optimal design matrices," *Discrete Math.*, vol. 50, pp. 295–318, 1984.

[54] X. Ou, G. Zheng, and C. Yang, "Embedded pupil function recovery for Fourier ptychographic microscopy," *Opt. Express*, vol. 22, no. 5, pp. 4960–4972, Mar. 2014.

[55] H. Q. Nguyen, E. Bostan, and M. Unser, "Learning convex regularizers for optimal Bayesian denoising," *IEEE Trans. Signal Process.*, vol. 66, no. 4, pp. 1093–1105, Feb. 2018.

[56] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, Dec. 2007.

[57] S. R. P. Pavani *et al.*, "Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function," *Proc. Nat. Acad. Sci.*, vol. 106, no. 9, pp. 2995–2999, Mar. 2009.

Emrah Bostan received the M.Sc. and Ph.D. degrees in electrical engineering from École Polytechnique Fédérale de Lausanne, Switzerland, in 2011 and 2016, respectively. He is currently a Postdoctoral researcher with the Computational Imaging Lab, University of California Berkeley, Berkeley, CA, USA. His research interest focuses on designing advanced signal/image processing algorithms for optical imaging applications.

Nicole A. Repina received the B.A. degree in biochemistry and molecular biology from Boston University, Boston, MA, USA, in 2013. She is currently working toward the Ph.D. degree in bioengineering with the University of California Berkeley, Berkeley, CA, USA and is funded by the National Science Foundation Graduate Research Fellowship and the Siebel Scholars foundation. Her research focuses on optogenetic control of stem cell signaling and optical methods for optogenetic stimulation.

Laura Waller received the B.S., M.Eng., and Ph.D. degrees in electrical engineering and computer sciences from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2004, 2005 and 2010, respectively. She is the Ted Van Duzer Associate Professor of electrical engineering and computer sciences with the University of California Berkeley, Berkeley, CA, USA, a Senior Fellow with the Berkeley Institute of Data Science, and affiliated with the UCB/UCSF Bioengineering Graduate Group. She was a Postdoctoral Researcher and a Lecturer of physics with Princeton University from 2010 to 2012. She is a Packard Fellow for Science & Engineering, Moore Foundation Data-driven Investigator, Bakar Fellow, and Chan-Zuckerberg Biohub Investigator. She has received the Carol D. Soc Distinguished Graduate Mentoring Award, Agilent Early Career Professor Award (Finalist), NSF CAREER Award, and the SPIE Early Career Achievement Award.

Michael R. Kellman received the B.S. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 2015, and the M.S. degree in electrical engineering and computer sciences from the University of California Berkeley, Berkeley, CA, USA, in 2017, where he is currently working toward the Ph.D. degree with the Department of Electrical Engineering and Computer Sciences, where his research is focused on the design of computational imaging systems for application in optical microscopy and medical imaging.