
When Deep Denoising Meets Iterative Phase Retrieval

Yaotian Wang¹ Xiaohang Sun¹ Jason W. Fleischer¹

Abstract

Recovering a signal from its Fourier intensity underlies many important applications, including lensless imaging and imaging through scattering media. Conventional algorithms for retrieving the phase suffer when noise is present but display global convergence when given clean data. Neural networks have been used to improve algorithm robustness, but efforts to date are sensitive to initial conditions and give inconsistent performance. Here, we combine iterative methods from phase retrieval with image statistics from deep denoisers, via regularization-by-denoising. The resulting methods inherit the advantages of each approach and outperform other noise-robust phase retrieval algorithms. Our work paves the way for hybrid imaging methods that integrate machine-learned constraints in conventional algorithms.

1. Introduction

Phase is a fundamental component of complex signals and often carries more information than amplitude (Oppenheim & Lim, 1981). In many cases, however, only the amplitude or intensity can be measured (e.g. conventional cameras can only record the time-averaged light intensity). For such measurements, the recovery of the underlying signal requires reconstruction of the missing phase. More formally, phase retrieval (PR) is an inverse problem to recover an unknown signal $x \in \mathbb{R}^n$ or \mathbb{C}^n from the phaseless measurement

$$y^2 = |Ax|^2 + w \quad (1)$$

where A is a known linear transform and w is the noise in the measurement.

In the past decade, the general PR problem has attracted much attention from the optimization and statistics community and various algorithms with provable convergence were

developed (Candès et al., 2015a;b; Zhang & Liang, 2016; Zhang et al., 2016; Chen & Candès, 2017; Wang et al., 2018). Despite a solid theoretical foundation, generally provable algorithms have overly restrictive requirements (e.g. the statistics of measurement bases) that have limited their popularity. More progress has been made for Fourier phase retrieval (FPR), in which A is the Fourier transform, i.e. the measurements are the Fourier intensities. This is also the most common type experimentally, mostly due to the fact that diffraction in the far field is the Fourier transform of the object (Goodman, 2005). Applications range from astronomy (Dainty & Fienup, 1987) to coherent diffraction (Miao et al., 1999; Chapman & Nugent, 2010) and speckle-correlation (Bertolotti et al., 2012; Katz et al., 2014) imaging.

The most broadly used algorithms for FPR are iterative methods, pioneered by Gerchberg-Saxton (Gerchberg & Saxton, 1972) and later developed by Fienup (Fienup, 1982). Though they lack theoretical proof of convergence, empirical use of Fienup algorithms and their variants (Bauschke et al., 2003; Elser, 2003; Luke, 2004; Martin et al., 2012; Rodriguez et al., 2013) has shown the avoidance of local minima and convergence to global solutions from random initialization. Together with the simplicity of their implementation, iterative phase retrieval methods have become the workhorse of FPR (Miao et al., 2005; Bertolotti et al., 2012; Katz et al., 2014).

It has been shown that applying the natural image prior to FPR can increase robustness to noise and improve reconstruction quality (Venkatakrishnan et al., 2013; Heide et al., 2016; Metzler et al., 2018; Işil et al., 2019). However, such methods either have unsatisfying robustness when noise levels are high or are sensitive to initialization (thus relying on other algorithms to supply initial points). Both cases return us to the problem of poor reliability when the signal-to-noise ratio in measurements is low.

Our major contribution here is to combine the benefits of iterative FPR with natural image priors via Regularization-by-Denoising (RED) (Romano et al., 2017). The methods we propose deliver greater robustness to noise than other noise-robust FPR algorithms while relaxing the initialization requirements. The application of image priors also alleviates the stagnant mode issues in iterative phase retrieval (Fienup

¹Department of Electrical Engineering, Princeton University, Princeton, New Jersey, USA. Correspondence to: Jason W. Fleischer <jasonf@princeton.edu>.

& Wackerman, 1986), leading to accelerated convergence. Machine learning thus resolves long-lasting issues that have hindered traditional methods. In turn, traditional algorithms can lift the burden on deep learning by focusing it on a subset of the whole, end-to-end problem.

2. Background

We focus on two-dimension signals $x \in \mathbb{C}^{\sqrt{n} \times \sqrt{n}}$ and assume the measurement transform A in (1) to be the discrete Fourier transform (DFT)

$$\hat{x}[k_1, k_2] = \frac{1}{\sqrt{n}} \sum_{n_1, n_2=0}^{\sqrt{n}-1} x[n_1, n_2] e^{-2\pi i \frac{n_1 k_1 + n_2 k_2}{\sqrt{n}}} \quad (2)$$

denoted here as $\hat{x} = Fx$ (with F^{-1} the inverse transform). Below, we discuss the uniqueness of Fourier phase retrieval, common algorithms used, and their relation to more general optimization problems.

2.1. Uniqueness in FPR

If there is not enough sampling, the Fourier intensity may be insufficient to trace back to the input signal. For all d -dimensional signals with $d \geq 2$, except a set of measure 0 (Hayes & McClellan, 1982), it has been shown that if the Fourier intensity is oversampled by a factor no less than 2 in each dimension, then a signal is determined uniquely by its Fourier intensity up to the trivial ambiguities of translation, conjugate inversion and global phase (Hayes, 1982). Fortunately, in practice these ambiguities are often acceptable, since the geometrical transform, complex conjugate and global phase keep the characteristics of the object intact.

Oversampling in the Fourier domain is related to the so-called support constraint for FPR, a terminology that is used more often in iterative phase retrieval. For example, suppose the Fourier spectrum of $x \in \mathbb{C}^{\sqrt{n} \times \sqrt{n}}$ is oversampled twice uniformly at $k_i = \{0, 1/2, 1, \dots, \sqrt{n} - 1/2\}$ for $i = 1, 2$. For simplicity, we write this oversampled spectrum as $\hat{x}^{(2)}$. By defining $\tilde{x} \in \mathbb{C}^{\sqrt{m} \times \sqrt{m}}$ with $m = 4n$ such that $\tilde{x}[n_1, n_2] = \sqrt{\frac{m}{n}} x[n_1, n_2]$ if $n_i \in \mathbb{N} < \sqrt{n}$ and $\tilde{x}[n_1, n_2] = 0$ otherwise, we have

$$\begin{aligned} \hat{x}^{(2)}[k_1, k_2] &= \frac{1}{\sqrt{n}} \sum_{n_1, n_2=0}^{\sqrt{n}-1} x[n_1, n_2] e^{-i2\pi \frac{n_1 k_1 + n_2 k_2}{\sqrt{n}}} \\ &= \frac{1}{\sqrt{n}} \sum_{n_1, n_2=0}^{\sqrt{m}-1} \sqrt{\frac{n}{m}} \tilde{x}[n_1, n_2] e^{-i2\pi \frac{n_1 k_1 + n_2 k_2}{\sqrt{m}}} \\ &= (F\tilde{x})[\tilde{k}_1, \tilde{k}_2] \end{aligned} \quad (3)$$

where $\tilde{k}_i = \{0, 1, \dots, 2\sqrt{n} - 1\} = 2k_i$. Therefore, there exists a supported signal \tilde{x} by zero-padding P_{mn} and scaling

x by a factor of $\sqrt{m/n}$, such that its Fourier transform is the same as (uniform) oversampling in the Fourier space of x . If the vectorization order gives

$$\tilde{x}^T = \sqrt{\frac{m}{n}} [x^T \quad 0_{m-n}^T] \quad (4)$$

then

$$\hat{x}^{(2)} = F\tilde{x} = FO_{mn}x \quad (5)$$

where $O_{mn} \in \mathbb{R}^{m \times n}$ is given by

$$O_{mn} = \sqrt{\frac{m}{n}} \begin{bmatrix} I_n \\ 0 \end{bmatrix} = \sqrt{\frac{m}{n}} P_{mn} \quad (6)$$

Stated another way, oversampling FPR is equivalent to finding a supported signal \tilde{x} from its DFT intensity, with the support constraint sometimes including the support of x itself. To distinguish them, we denote the support for $x \in \mathbb{C}^n$ as $S = \{i \mid x_i \neq 0\}$ and the extended support for padded \tilde{x} as $\tilde{S} = \{j \mid \tilde{x}_j \neq 0\}$

2.2. ADMM

The Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2011; Hong et al., 2016) is a popular algorithm for solving the linearly constrained optimization problem

$$\begin{aligned} \underset{x_1, \dots, x_N}{\text{minimize}} \quad & \ell(x_1, \dots, x_N) = \sum_i^N f_i(x_i) \\ \text{subject to} \quad & \sum_i^N A_i x_i = b \end{aligned} \quad (7)$$

For each iteration, ADMM updates each x_i sequentially and dual variable u as

$$\begin{aligned} x_i^{k+1} &= \underset{x_i}{\text{argmin}} f_i(x_i) + \frac{\rho^k}{2} \left\| \sum_{j \neq i} A_j x_j^k + A_i x_i - b + u^k \right\|^2 \\ u^{k+1} &= u^k + \sum_{i=1}^N A_i x_i^{k+1} - b \end{aligned} \quad (8)$$

with $x_j^k = x_j^{k+1}$ for $j < i$ and $x_j^k = x_j^k$ otherwise for each update of x_i^{k+1} . The penalty parameter ρ^k is constant or adaptive through iterations.

Often, the minimization problem takes a form

$$\text{prox}_f(z) := \underset{v}{\text{argmin}} f(v) + \frac{1}{2} \|v - z\|_2^2 \quad (9)$$

which is defined as the proximal operator for f at z (Parikh & Boyd, 2014). The efficiency of ADMM generally depends on the complexity of evaluating the proximal operator for each f_i , while in return it is possible to minimize functions that are non-differentiable. We show below that this latter property can be quite beneficial.

2.3. Hybrid-Input-Output Method

As possibly the most used iterative method in FPR, Hybrid-Input-Output (HIO) (Fienup, 1982) is well known for its ability to converge to global minima from random initialization. HIO iterates on the padded and scaled signal \tilde{x} with the following step rules:

$$\begin{aligned} \tilde{z}^{k+1} &= F^{-1} \left(y \odot \frac{F\tilde{x}^k}{|F\tilde{x}^k|} \right) \\ \forall i, \tilde{x}_i^{k+1} &= \begin{cases} \tilde{z}_i^{k+1} & \text{if } i \in \tilde{S} \\ \tilde{x}_i^k - \beta \tilde{z}_i^{k+1} & \text{otherwise} \end{cases} \end{aligned} \quad (10)$$

where the product \odot , the modulus $|\cdot|$ and division are all performed element-by-element.

It was shown in (Bauschke et al., 2002) that HIO with $\beta = 1$ coincides with Douglas-Rachford splitting (DRS) (Douglas & Rachford, 1956; Lions & Mercier, 1979; Eckstein & Bertsekas, 1992). Since DRS is equivalent to ADMM updates on the feasibility problem with indicator functions (Boyd et al., 2011), one can find that HIO ($\beta = 1$) is equivalent to ADMM on the following minimization problem:

$$\begin{aligned} &\underset{x \in \mathbb{C}^n, z \in \mathbb{C}^m}{\text{minimize}} \quad \bar{I}_{\mathcal{M}}(z) + \bar{I}_{\mathcal{C}}(x) \\ &\text{subject to } z = O_{mn}x \end{aligned} \quad (11)$$

where the indicator function for a subset \mathcal{S} is defined as (Boyd & Vandenberghe, 2004)

$$\bar{I}_{\mathcal{S}}(x) = \begin{cases} 0 & \text{if } x \in \mathcal{S} \\ +\infty & \text{otherwise} \end{cases} \quad (12)$$

and the set \mathcal{M} is defined as the set of signals consistent with the measurement

$$\mathcal{M} := \{x \in \mathbb{C}^m \mid |Fx| = y\} \quad (13)$$

Here, \mathcal{C} is the set of signals satisfying an additional constraint, such as inset support S and nonnegativity (which result in the Hybrid-Projection-Reflection algorithm (Bauschke et al., 2003)). More details of this mapping are given in the supplementary material.

The proximal operator of the indicator function (12) is given by the projection onto the corresponding set

$$\Pi_{\mathcal{S}}(x) := \underset{z \in \mathcal{S}}{\text{argmin}} \|z - x\| = \text{prox}_{\bar{I}_{\mathcal{S}}}(x) \quad (14)$$

In particular, the projection onto \mathcal{M} can be written as

$$F^{-1} \left(y \odot \frac{Fv}{|Fv|} \right) \in \Pi_{\mathcal{M}}(v) = \text{prox}_{\bar{I}_{\mathcal{M}}}(v) \quad (15)$$

For the additional constraint \mathcal{C} in (11), nonnegativity in the real part of the signal can be used. The projection onto the

corresponding subset is

$$\Pi_{\text{Re}^+}(x)_i = \begin{cases} x_i & \text{if } \Re(x_i) \geq 0 \\ i\Im(x_i) & \text{otherwise} \end{cases} \quad (16)$$

where $\Re(\cdot)$ and $\Im(\cdot)$ are the real and imaginary parts of a complex-valued signal.

3. Related Works

In this section, we introduce the efforts to date for solving the PR problem in the presence of noise.

3.1. Iterative Phase Retrieval

Iterative phase retrieval methods commonly solve a feasibility problem, looking for a signal whose oversampled Fourier intensity is y^2 and simultaneously consistent with the other constraint \mathcal{C} . The problem occurs when noise levels increase in the measurement, resulting in oscillations and ambiguous solutions. To alleviate the degradation from corrupted data, efforts have been made to limit the effect of noise on iterative methods (Luke, 2004; Martin et al., 2012; Rodriguez et al., 2013). However, without further priors on the object space (e.g. image statistics), the denoising effect of these methods is often insufficient.

3.2. Deep Learning in PR

Deep neural networks (DNN) are well known for their capability to approximate complicated functions (given enough training data). In image processing, they have achieved significant improvements over traditional methods in areas such as denoising (Zhang et al., 2017a; 2018), deblurring (Nimisha et al., 2017), and superresolution (Dong et al., 2014; Lim et al., 2017). For solving PR, deep feedforward networks have shown some success in end-to-end predictions (Sinha et al., 2017; Rivenson et al., 2018), while network-assisted algorithms also have helped in support estimation (Kim & Chung, 2019), low-light (Goy et al., 2018) and compressive (Hand et al., 2018) situations.

However, using a feedforward neural network to approximate the inverse mapping is problematic for oversampling FPR. Such methodology relies on the assumption that forward mapping is one-to-one and well posed; this is not the case here, even with precise knowledge of the signal support, due to the existence of trivial ambiguities. Instead, the optimization method commonly adopted for solving FPR (e.g., in (Heide et al., 2016; Metzler et al., 2018)) minimizes the loss function

$$\ell(x) := f(x; y) + \alpha R(x) \quad (17)$$

where f is the data fidelity term and R is a regularizer involving prior belief, e.g. natural image statistics. This

method is effectively a maximum a posteriori estimation (Venkatakrisnan et al., 2013).

3.3. Prior by Denoisers

Using a denoiser as the prior R in (17) has been proposed to boost image inference in inverse problems. There have been two major strategies to utilize the denoiser: Plug-and-Play (PnP) regularization (Venkatakrisnan et al., 2013) and Regularization-by-Denoising (RED) (Romano et al., 2017). In PnP methods, the proximal operator for an implicit regularizer R is approximated by an image denoiser. This approach provides promising results both empirically (Venkatakrisnan et al., 2013; Heide et al., 2014; 2016; Metzler et al., 2016; Meinhardt et al., 2017; Zhang et al., 2017b) and theoretically (Chan et al., 2017). Meanwhile, RED is a framework that constructs explicit regularizers with denoisers D as the inner product between a signal and the noise it contains,

$$R(x) = \frac{\lambda}{2} \langle x, x - D(x) \rangle \quad (18)$$

It has been shown in (Romano et al., 2017) that if the denoiser D has the properties of (local) homogeneity and Jacobian symmetry, then evaluation of the proximal operator for (18) at x requires the solution to

$$z - x + \lambda(z - D(z)) = 0 \quad (19)$$

Though these properties rarely hold for common denoisers, (19) can still be adopted either as an approximation or if certain conditions hold (Reehorst & Schniter, 2019). Recent applications of RED to PR have demonstrated a significant boost in noise robustness compared with bare iterative methods (Metzler et al., 2018; Wu et al., 2019).

4. Methodology

We aim to maintain the convergence benefits of HIO while alleviating the deleterious effects of noise. To this end, we adopt ADMM as the solver but modify the loss function used in HIO. More specifically, we eliminate the inconsistency from (11) by relaxation of the loss function and include natural image priors via RED, due to its explicit form and inherent flexibility.

For relaxation of the loss function, we consider two approaches: one in the Fourier constraint and one in the over-sampling constraint. These give two algorithms, RED-ITA-F and RED-ITA-S, respectively.

In general, we refer to our algorithms as RED-ITA, and Deep-ITA for the specific choice of deep denoisers, such as DnCNN (Zhang et al., 2017a).

4.1. RED-ITA-F

We first consider substituting the indicator function on Fourier measurement with the data fidelity term. Following (Metzler et al., 2018), we seek to solve

$$\frac{1}{2} \|y - |FO_{mn}x|\|^2 + \frac{\lambda}{2} \langle x, x - D(x) \rangle \quad (20)$$

Similar to HIO, we transform (20) into a linearly constrained form as

$$\begin{aligned} & \underset{x \in \mathbb{R}^n, z \in \mathbb{C}^m}{\text{minimize}} \quad \frac{1}{2} \|y - |Fz|\|^2 + R(x) \\ & \text{subject to} \quad z = O_{mn}x \end{aligned} \quad (21)$$

where R contains RED and an additional constraint $\bar{I}_C(x)$:

$$R(x) = \bar{I}_C(x) + \frac{\lambda}{2} \langle x, x - D(x) \rangle \quad (22)$$

For $f(z) = \frac{1}{2} \|y - |Fz|\|^2$, the update rule of ADMM gives

$$\begin{aligned} x^{k+1} &= \underset{x \in \mathbb{R}^n}{\text{armin}} R(x) + \frac{\rho}{2} \|z^k - O_{mn}x + u^k\|^2 \\ z^{k+1} &= \text{prox}_{\frac{1}{\rho}f}(O_{mn}x^{k+1} - u^k) \\ u^{k+1} &= u^k + z^{k+1} - O_{mn}x^{k+1} \end{aligned} \quad (23)$$

It remains to evaluate each update step. We note that for any $x \in \mathbb{R}^n$, $v = [v_n^T \quad v_{m-n}^T]^T \in \mathbb{C}^m$, where $v_n = P_{mn}^T v = \sqrt{\frac{n}{m}} O_{mn}^T v$,

$$\begin{aligned} \|v - O_{mn}x\|^2 &= \|\Re(v) - O_{mn}x\|^2 + \|\Im(v)\|^2 \\ &= \frac{m}{n} \left\| \sqrt{\frac{n}{m}} \Re(v_n) - x \right\|^2 + \|\Re(v_{m-n})\|^2 + \|\Im(v)\|^2 \end{aligned}$$

Therefore, in terms of $v = z^k + u^k$, the x -update step in (23) can be found as

$$\begin{aligned} x^{k+1} &= \underset{x \in \mathbb{R}^n}{\text{armin}} R(x) + \frac{\rho}{2} \|v - O_{mn}x\|^2 \\ &= \underset{x \in \mathbb{R}^n}{\text{armin}} R(x) + \frac{m\rho}{2n} \left\| \sqrt{\frac{n}{m}} \Re(v_n) - x \right\|^2 \\ &= \text{prox}_{\frac{n}{m\rho}R} \left(\sqrt{\frac{n}{m}} \Re(v_n) \right) \end{aligned} \quad (24)$$

which reduces to an evaluation of the proximal operator for R . For any $\tau > 0$, if $s^+ = \text{prox}_{\tau R}(s)$, we have

$$s^+ = \Pi_C \left(\frac{s + \lambda\tau D(s^+)}{1 + \lambda\tau} \right) \quad (25)$$

(a derivation is given in the supplementary material). Similar to RED in (Romano et al., 2017), the proximal operator in (25) can be evaluated by the fixed-point approach, updating

$$s^{(k+1)} = \Pi_C \left(\frac{s + \lambda\tau D(s^{(k)})}{1 + \lambda\tau} \right) \quad (26)$$

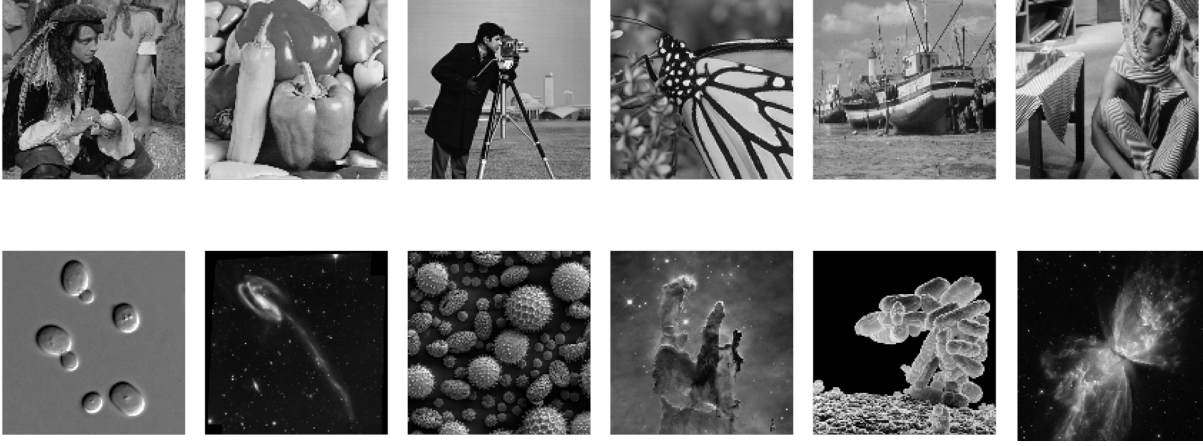


Figure 1. Test images used in the simulation. Top row: 6 commonly used “natural” test images (Zhang et al., 2017a). Bottom row: 6 “unnatural” images (Metzler et al., 2018). Images have been resized to 128×128 .

Algorithm 1 RED-ITA-F

Input: Initialization $z^0, u^0 \in \mathbb{C}^m$, $\rho, \lambda > 0$, oversampling transform O_{mn} , Fourier measurement y

for $k = 0, 1, 2, \dots$ **do**

$$v^k = z^k + u^k$$

$$\tau = (m\rho)^{-1}n$$

$$x^{k+1} = \text{prox}_{\tau R}(\frac{n}{m} \Re(O_{mn}^T v^k))$$

$$\tilde{x}^{k+1} = O_{mn} x^{k+1}$$

$$z^{k+1} = \frac{\rho}{\rho+1}(\tilde{x}^{k+1} - u^k) + \frac{1}{\rho+1}\Pi_{\mathcal{M}}(\tilde{x}^{k+1} - u^k)$$

$$u^{k+1} = u^k + z^{k+1} - \tilde{x}^{k+1}$$

end for

until convergence. In practice, the fixed point can be approximated by stopping after p iterations, denoted as $\text{prox}_{\tau R}(s) = s^{(p)}$ with p being a hyperparameter and $s^{(0)} = s$. Empirically, we found that $p = 1$ is efficient enough; therefore, p is set to 1 in all of our experiments.

For the z -update step, the proximal operator for f can be written as

$$\text{prox}_{\tau f}(s) = \frac{1}{\tau+1}s + \frac{\tau}{\tau+1}\Pi_{\mathcal{M}}(s) \quad (27)$$

This method for solving oversampling FPR is shown in Algorithm 1.

4.2. RED-ITA-S

The second approach relaxes the oversampling constraint, instead of the Fourier measurement. Rather than assuming there exists $x \in \mathbb{R}^n$ such that $O_{mn}x = z \in \mathcal{M}$, we acknowledge that the difference $\xi = z - O_{mn}x$ can be non-zero $\forall z \in \mathcal{M}$ and minimize the norm of it. That is, an

alternative to (21) is

$$\begin{aligned} & \underset{z, \xi \in \mathbb{C}^m, x \in \mathbb{R}^n}{\text{minimize}} \quad \bar{I}_{\mathcal{M}}(z) + \frac{1}{2}\|\xi\|^2 + R(x) \\ & \text{subject to: } z = O_{mn}x + \xi \end{aligned} \quad (28)$$

Note that, given $x \in \mathbb{R}^n$, the loss in (28) is an upper bound for that in (21), since $\forall z \in \mathcal{M}$, Parseval’s theorem gives

$$\begin{aligned} \|\xi\|^2 &= \|z - O_{mn}x\|^2 \\ &= \|ye^{i\phi_z} - FO_{mn}x\|^2 \\ &\geq \|y - |FO_{mn}x|\|^2 \end{aligned} \quad (29)$$

where ϕ_z is the Fourier phase of z .

A three-block ADMM is adopted to solve (28):

$$\begin{aligned} x^{k+1} &= \underset{x \in \mathbb{R}^n}{\text{argmin}} \quad R(x) + \frac{\rho}{2}\|z^k - O_{mn}x - \xi^k + u^k\|^2 \\ z^{k+1} &= \text{prox}_{\bar{I}_{\mathcal{M}}}(O_{mn}x^{k+1} + \xi^k - u^k) \\ \xi^{k+1} &= \text{prox}_{\frac{1}{2\rho}\|\cdot\|^2}(z^{k+1} - O_{mn}x^{k+1} + u^k) \\ u^{k+1} &= u^k + z^{k+1} - O_{mn}x^{k+1} - \xi^{k+1} \end{aligned} \quad (30)$$

where

$$\begin{aligned} \text{prox}_{\bar{I}_{\mathcal{M}}}(s) &= \Pi_{\mathcal{M}}(s) \\ \text{prox}_{\frac{1}{2\rho}\|\cdot\|^2}(s) &= \frac{\tau s}{1 + \tau} \end{aligned} \quad (31)$$

This yields the RED-ITA-S shown in Algorithm 2.

4.3. Connection between PR Algorithms

(Metzler et al., 2018) proposed solving (20) with FASTA (Goldstein et al., 2014), a method known as prRED (the

Algorithm 2 RED-ITA-S

Input: Initialization $\xi^0, z^0, u^0 \in \mathbb{C}^m$, $\rho, \lambda > 0$, oversampling transform O_{mn} , Fourier measurement y

for $k = 0, 1, 2, \dots$ **do**

$v^k = z^k + u^k - \xi^k$

$\tau = (m\rho)^{-1}n$

$x^{k+1} = \text{pr}\hat{\text{O}}_{\tau R}(\frac{n}{m}\Re(O_{mn}^T v^k))$

$\tilde{x}^{k+1} = O_{mn}x^{k+1}$

$z^{k+1} \in \Pi_{\mathcal{M}}(\tilde{x}^{k+1} + \xi^k - u^k)$

$\xi^{k+1} = \frac{\rho}{\rho+1}(z^{k+1} - \tilde{x}^{k+1} + u^k)$

$u^{k+1} = u^k + z^{k+1} - \tilde{x}^{k+1} - \xi^{k+1}$

end for

variant using deep denoisers like DnCNN is referred to as prDeep). Since FASTA is a forward-backward splitting method, if the stepsize μ is fixed to be n/m and $\lambda \rightarrow 0$, prDeep reduces to (sub-)gradient descent on the squared loss on Fourier amplitude, which coincides (Marchesini, 2007) with the Error Reduction algorithm (Fienup, 1982).

RED-ITA-F reduces to HIO with $\beta = 1$ when $\rho \rightarrow 0$ and $\lambda/\rho \rightarrow 0$. Similarly for DnCNN-ADMM if the denoising step is put first and the denoiser is the identity mapping.

5. Experimental Results

We compare Deep-ITA-F/S with other widely used algorithms on FPR, namely HIO (Fienup, 1982), Oversampling Smoothness (OSS) (Rodriguez et al., 2013), DnCNN-ADMM (Venkatakrisnan et al., 2013; Heide et al., 2016; Chan et al., 2017) and prDeep (Metzler et al., 2018). We did not include any post-reconstruction procedure to clean the results as in (Işil et al., 2019), which is not tested here since the algorithm performs worse than prDeep unless an additional DNN specifically trained to enhance the quality is used. Provable methods (Candès et al., 2015a;b; Zhang & Liang, 2016; Zhang et al., 2016; Chen & Candès, 2017; Wang et al., 2018) are also excluded in comparison, since their measurement models focus on Gaussian distributions and do not cover oversampled Fourier intensities. Besides, results of a typical provable method, Wirtinger Flow (Candès et al., 2015b), for oversampling FPR have been reported in (Metzler et al., 2018), where it significantly underperformed even HIO.

In principle, any denoiser can be adopted in RED. Here, we choose DnCNN (Zhang et al., 2017a), based on its competitive denoising performance and its flexibility on the input signal. DnCNN is stacked by Convolutional and Batch Normalization layers with Rectified Linear Unit (ReLU) activation functions. With a zero-padding of size 1 for 3×3 convolutional kernel size, the output dimension remains the same as that of the input. DnCNN models are trained on patches of natural images with mean-squared-error as the

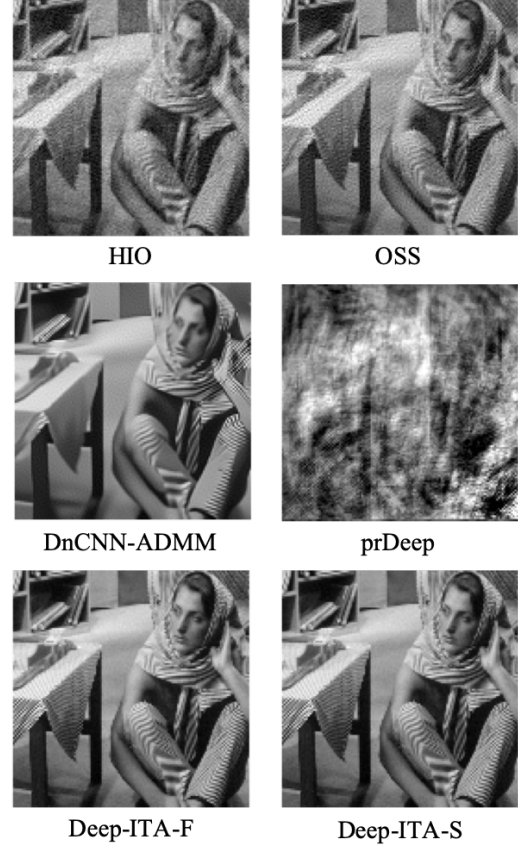


Figure 2. Reconstructions from random initialization with $\alpha = 4$. Both RED-ITA-F/S have the best reconstruction results.

loss function, using Adam as the optimizer (Kingma & Ba, 2014).

The test images used in the simulations, shown in Figure 1, consist of 6 commonly used “natural” images and 6 “unnatural” ones. The images are resized to 128×128 and their Fourier intensities are oversampled uniformly by a factor of 2 in each dimension, yielding measurements of size 256×256 . The signals used as ground truth are real-valued and have dynamic range of $[0, 255]$.

For simulation, shot noise is assumed to dominate the noise in the measurement. While this noise follows a Poisson distribution, it is commonly approximated as a Gaussian (Metzler et al., 2018; Işil et al., 2019). The noisy measurement y on the oversampled Fourier amplitude $q = \hat{x}^{(2)}$ thus has the distribution

$$y^2 = |q|^2 + w \quad w \sim \mathcal{N}(0, \text{diag}(\alpha^2 |q|^2)) \quad (32)$$

It is worth noting that the (effective) SNR in the measurements scales roughly with y/α , which is affected by α and any scaling in $|q|$. We define two metrics to characterize the SNR: $\text{MSNR}_1 = 10 \log_{10}(\| |q|^2 \|_2 / \| y^2 - |q|^2 \|_2)$ (Işil et al.,

Table 1. PSNRs and SSIMs of reconstructions initialized with random noise with varying noise level in the measurements. For $\alpha = 0$, no noise is added to the Fourier intensity. For $\alpha = 4$, averaged $\text{MSNR}_1 = 32.09\text{dB}$, $\text{MSNR}_2 = 33.36\text{dB}$.

$\alpha = 0$	AVERAGE PSNR			AVERAGE SSIM		
	NATURAL	UNNATURAL	OVERALL	NATURAL	UNNATURAL	OVERALL
HIO	48.88	56.01	52.45	0.94	0.88	0.91
OSS	24.27	44.31	34.29	0.73	0.82	0.77
PRDEEP	13.70	18.27	15.99	0.21	0.27	0.24
DNCNN-ADMM	29.11	27.94	28.52	0.87	0.74	0.80
DEEP-ITA-F	65.06	57.88	61.47	1.00	0.99	1.00
DEEP-ITA-S	64.94	57.93	61.44	1.00	0.99	1.00

$\alpha = 4$	AVERAGE PSNR			AVERAGE SSIM		
	NATURAL	UNNATURAL	OVERALL	NATURAL	UNNATURAL	OVERALL
HIO	23.29	26.36	24.83	0.69	0.67	0.68
OSS	22.02	30.70	26.36	0.64	0.70	0.67
PRDEEP	14.52	19.46	16.99	0.23	0.36	0.30
DNCNN-ADMM	28.46	27.65	28.05	0.86	0.69	0.77
DEEP-ITA-F	36.32	34.39	35.35	0.97	0.87	0.92
DEEP-ITA-S	36.47	35.65	36.06	0.97	0.91	0.94

Table 2. PSNRs and SSIMs of reconstructions initialized from HIO with varying noise level in the measurements. For $\alpha = 8$, the averaged $\text{MSNR}_1 = 29.09\text{dB}$, $\text{MSNR}_2 = 27.54\text{dB}$. For $\alpha = 12$, averaged $\text{MSNR}_1 = 27.38\text{dB}$, $\text{MSNR}_2 = 24.49\text{dB}$. For $\alpha = 16$, averaged $\text{MSNR}_1 = 25.84\text{dB}$, $\text{MSNR}_2 = 22.52\text{dB}$.

$\alpha = 8$	AVERAGE PSNR			AVERAGE SSIM		
	NATURAL	UNNATURAL	OVERALL	NATURAL	UNNATURAL	OVERALL
HIO (INIT.)	20.78	23.03	21.91	0.56	0.53	0.55
OSS	22.02	27.58	24.80	0.63	0.65	0.64
PRDEEP	28.50	30.75	29.62	0.87	0.79	0.83
DNCNN-ADMM	26.95	27.76	27.35	0.81	0.68	0.75
DEEP-ITA-F	32.90	31.36	32.13	0.94	0.83	0.89
DEEP-ITA-S	33.31	32.78	33.04	0.94	0.86	0.90

$\alpha = 12$	AVERAGE PSNR			AVERAGE SSIM		
	NATURAL	UNNATURAL	OVERALL	NATURAL	UNNATURAL	OVERALL
HIO (INIT.)	19.36	21.66	20.51	0.47	0.45	0.46
OSS	20.78	25.07	22.93	0.56	0.56	0.56
PRDEEP	28.24	27.46	27.85	0.85	0.74	0.79
DNCNN-ADMM	25.43	25.89	25.66	0.79	0.61	0.70
DEEP-ITA-F	30.09	29.11	29.60	0.91	0.79	0.85
DEEP-ITA-S	31.95	30.38	31.17	0.93	0.81	0.87

$\alpha = 16$	AVERAGE PSNR			AVERAGE SSIM		
	NATURAL	UNNATURAL	OVERALL	NATURAL	UNNATURAL	OVERALL
HIO (INIT.)	17.59	20.50	19.05	0.36	0.38	0.37
OSS	19.65	23.37	21.51	0.50	0.51	0.50
PRDEEP	26.44	24.65	25.54	0.81	0.65	0.73
DNCNN-ADMM	22.87	24.27	23.57	0.65	0.55	0.60
DEEP-ITA-F	27.63	26.79	27.20	0.86	0.75	0.81
DEEP-ITA-S	28.14	27.38	27.76	0.86	0.75	0.81

2019) and $\text{MSNR}_2 = 20 \log_{10}(\| |q| \|_2 / \|y - |q| \|_2)$ (Luke, 2004).

Results from two experimental setups are reported here. In the first, we test the convergence of the competing phase retrieval algorithms with random initialization. All algorithms are initialized with the same random point and run for the same total number of 1200 iterations. In the second, we fol-

low the initializing strategy used in (Metzler et al., 2018; Işil et al., 2019): first, make 50 runs of randomly initialized HIO (giving \hat{x}_i for $i = 1, \dots, 50$), each with 50 iterations; next, pass the one with the lowest residual $\hat{x} = \text{argmin}_i f(\hat{x}_i)$ to initialize another HIO run of 1000 iterations. The output is then used as initialization for other algorithms. For both experiments, the whole procedure is repeated three times




Figure 3. Reconstructions from random initialization with $\alpha = 0$. The stripes in the HIO reconstruction are artifacts from stagnation (Fienup & Wackerman, 1986); they are resolved in our method.

and the one most matched with the measurement is selected as the final output for each algorithm.

The parameters in the algorithms were as follows: for HIO and OSS, $\beta = 0.9$. The regularization parameter λ is found best set as $0.01\bar{\sigma}^2$ for DnCNN-ADMM, $0.025\bar{\sigma}^2$ for both Deep-ITA-S/F, and $0.05\bar{\sigma}^2$ for prDeep, where $\bar{\sigma}$ is the standard deviation of noise in the Fourier amplitude (or set to 0.1 if no noise is added). Similar to the practice in (Metzler et al., 2018), prDeep and Deep-ITAs sequentially use DnCNN models that are trained with noise standard deviations of 60, 40, 20 and 10, each with 300 iterations for a total of 1200 iterations. The penalty parameter ρ used in Deep-ITAs is set to $\frac{1}{2}\lambda$. We notice that reducing λ and ρ when using the DnCNNs for high noise levels can increase the stability of our methods. For prDeep and Deep-ITAs, we use the nonnegativity in (16) as the additional constraint in the regularizer.

For quantitative evaluation of the reconstructions, we characterize the output by its Peak Signal-to-Noise Ratio (PSNR) compared to the ground truth as well as the Structure Similarity (SSIM) Index (Wang et al., 2004). The PSNR computed for each reconstruction is capped at 80dB, in case an outlier has a high value that adversely affects the estimation of mean reconstruction quality (which could happen, e.g., in the noise-free case $\alpha = 0$.)

5.1. Random Initialization

Results of the experiments with random initialization are shown in Figure 2 and Table 1. Our methods outperform every other PR algorithm by large margins, in both PSNR and SSIM. Significantly, this includes HIO even when noise is absent (Figure 3). (This is probably due to stagnation in HIO, which is hard to overcome in a limited number of iterations (Fienup & Wackerman, 1986).) prDeep has issues with random initialization, which is not surprising considering its connection with Error Reduction, which has been shown to have slow convergence in practice (Fienup,

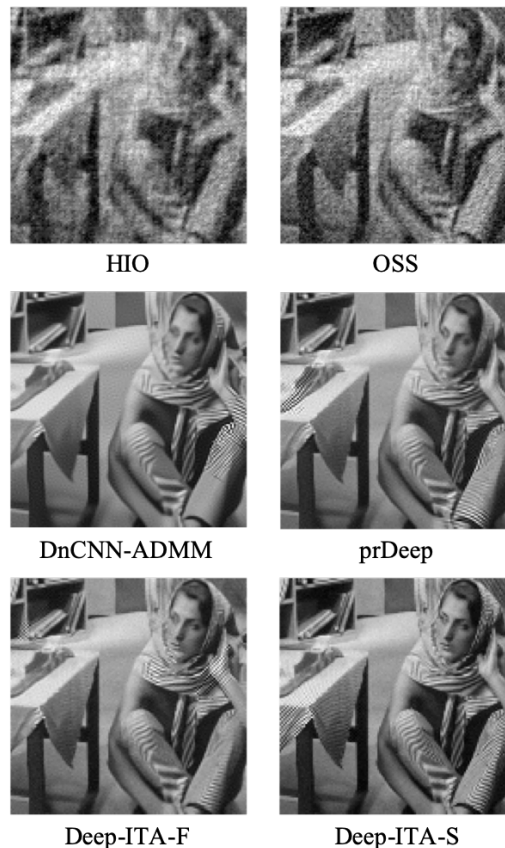


Figure 4. Reconstructions initialized from HIO with $\alpha = 12$. Deep-ITA-S has the best reconstruction results.

1982). On the contrary, DnCNN-ADMM and Deep-ITAs have the ability to work with random initial points, since all of them use ADMM as the solver. Our methods are more effective, as we integrate the denoiser in the update via RED, rather than apply it in a Plug-and-Play manner.

5.2. Initialization by HIO

Table 2 shows the performance of test algorithms with different level of noise in Fourier intensity when initialized with HIO. Deep-ITAs exhibit higher robustness to noise for every level of noise added. Figure 4 shows a visual comparison between PR algorithms for $\alpha = 12$, where Deep-ITA-S provides the best reconstruction. For the other methods, artifacts appear in the reconstructions and many details are lost.

6. Conclusion

Phase retrieval is part of a more general class of algorithms that has (to date) resisted full, end-to-end solutions from machine learning. While an admirable goal, such approaches

often apply machine learning in situations where it is ill-suited. It also neglects traditional algorithms and their corresponding strengths, viz. convergence benefits.

The approach advocated here is to build algorithms in the fashion of traditional methods but with added priors utilizing deep neural networks. In the problem of Fourier phase retrieval, we added the object-space regularizer of image statistics and improved noise robustness. More generally, the results pave the way for hybrid methods that integrate machine-learned constraints in conventional algorithms.

Acknowledgements

This material is based upon work supported by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-18-1-0219 and by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111890042.

References

- Bauschke, H. H., Combettes, P. L., and Luke, D. R. Phase retrieval, error reduction algorithm, and fienuip variants: a view from convex optimization. *J. Opt. Soc. Am. A*, 19(7):1334–1345, Jul 2002.
- Bauschke, H. H., Combettes, P. L., and Luke, D. R. Hybrid projection–reflection method for phase retrieval. *J. Opt. Soc. Am. A*, 20(6):1025–1034, Jun 2003.
- Bertolotti, J., van Putten, E. G., Blum, C., Lagendijk, A., Vos, W. L., and Mosk, A. P. Non-invasive imaging through opaque scattering layers. *Nature*, 491(7423):232–234, November 2012.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- Candès, E. J., Eldar, Y. C., Strohmer, T., and Voroninski, V. Phase retrieval via matrix completion. *SIAM Review*, 57(2):225–251, 2015a.
- Candès, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015b.
- Chan, S. H., Wang, X., and Elgendy, O. A. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2017.
- Chapman, H. N. and Nugent, K. A. Coherent lensless X-ray imaging. *Nature Photonics*, 4(12):833–839, December 2010.
- Chen, Y. and Candès, E. J. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Communications on Pure and Applied Mathematics*, 70(5):822–883, 2017.
- Dainty, J. C. and Fienup, J. R. Phase retrieval and image reconstruction for astronomy. *Image Recovery: Theory and Application*, pp. 231–275, 1987.
- Dong, C., Loy, C. C., He, K., and Tang, X. Learning a deep convolutional network for image super-resolution. In *Computer Vision – ECCV 2014*, pp. 184–199, Cham, 2014. Springer International Publishing.
- Douglas, J. and Rachford, H. H. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):421–439, 1956.
- Eckstein, J. and Bertsekas, D. P. On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1):293–318, April 1992.
- Elser, V. Phase retrieval by iterated projections. *J. Opt. Soc. Am. A*, 20(1):40–55, Jan 2003.
- Fienup, J. R. Phase retrieval algorithms: a comparison. *Appl. Opt.*, 21(15):2758–2769, Aug 1982.
- Fienup, J. R. and Wackerman, C. C. Phase-retrieval stagnation problems and solutions. *J. Opt. Soc. Am. A*, 3(11):1897–1907, Nov 1986.
- Gerchberg, R. W. and Saxton, W. O. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- Goldstein, T., Studer, C., and Baraniuk, R. A field guide to forward-backward splitting with a fast implementation. *arXiv preprint arXiv:1411.3406*, 2014.
- Goodman, J. W. *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.
- Goy, A., Arthur, K., Li, S., and Barbastathis, G. Low photon count phase retrieval using deep learning. *Phys. Rev. Lett.*, 121:243902, Dec 2018.
- Hand, P., Leong, O., and Voroninski, V. Phase retrieval under a generative prior. In *Advances in Neural Information Processing Systems*, pp. 9136–9146, 2018.

- Hayes, M. The reconstruction of a multidimensional sequence from the phase or magnitude of its fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(2):140–154, 1982.
- Hayes, M. H. and McClellan, J. H. Reducible polynomials in more than one variable. *Proceedings of the IEEE*, 70(2):197–198, 1982.
- Heide, F., Steinberger, M., Tsai, Y.-T., Rouf, M., Pająk, D., Reddy, D., Gallo, O., Liu, J., Heidrich, W., Egiazarian, K., Kautz, J., and Pulli, K. Flexisp: A flexible camera image processing framework. *ACM Trans. Graph.*, 33(6), November 2014.
- Heide, F., Diamond, S., Nießner, M., Ragan-Kelley, J., Heidrich, W., and Wetzstein, G. Proximal: Efficient image optimization using proximal algorithms. *ACM Trans. Graph.*, 35(4), July 2016.
- Hong, M., Luo, Z.-Q., and Razaviyayn, M. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.
- Işil, Ç., Oktem, F. S., and Koç, A. Deep iterative reconstruction for phase retrieval. *Appl. Opt.*, 58(20):5422–5431, Jul 2019.
- Katz, O., Heidmann, P., Fink, M., and Gigan, S. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature Photonics*, 8(10):784–790, October 2014.
- Kim, K. and Chung, S. Fourier phase retrieval with extended support estimation via deep neural network. *IEEE Signal Processing Letters*, 26(10):1506–1510, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lim, B., Son, S., Kim, H., Nah, S., and Mu Lee, K. Enhanced deep residual networks for single image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- Lions, P. L. and Mercier, B. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- Luke, D. R. Relaxed averaged alternating reflections for diffraction imaging. *Inverse Problems*, 21(1):37–50, nov 2004.
- Marchesini, S. Phase retrieval and saddle-point optimization. *J. Opt. Soc. Am. A*, 24(10):3289–3296, Oct 2007.
- Martin, A. V., Wang, F., Loh, N. D., Ekeberg, T., Maia, F. R. N. C., Hantke, M., van der Schot, G., Hampton, C. Y., Sierra, R. G., Aquila, A., Bajt, S., Barthelmess, M., Bostedt, C., Bozek, J. D., Coppola, N., Epp, S. W., Erk, B., Fleckenstein, H., Foucar, L., Frank, M., Graafsma, H., Gumprecht, L., Hartmann, A., Hartmann, R., Hauser, G., Hirsemann, H., Holl, P., Kassemeyer, S., Kimmel, N., Liang, M., Lomb, L., Marchesini, S., Nass, K., Pedersoli, E., Reich, C., Rolles, D., Rudek, B., Rudenko, A., Schulz, J., Shoeman, R. L., Soltau, H., Starodub, D., Steinbrener, J., Stellato, F., Strüder, L., Ullrich, J., Weidenspointner, G., White, T. A., Wunderer, C. B., Barty, A., Schlichting, I., Bogan, M. J., and Chapman, H. N. Noise-robust coherent diffractive imaging with a single diffraction pattern. *Opt. Express*, 20(15):16650–16661, Jul 2012.
- Meinhardt, T., Moller, M., Hazirbas, C., and Cremers, D. Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- Metzler, C., Schniter, P., Veeraraghavan, A., and richard baraniuk. prDeep: Robust phase retrieval with a flexible deep network. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pp. 3501–3510, 10–15 Jul 2018.
- Metzler, C. A., Maleki, A., and Baraniuk, R. G. From denoising to compressed sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.
- Miao, J., Charalambous, P., Kirz, J., and Sayre, D. Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*, 400(6742):342–344, July 1999.
- Miao, J., Nishino, Y., Kohmura, Y., Johnson, B., Song, C., Risbud, S. H., and Ishikawa, T. Quantitative image reconstruction of gan quantum dots from oversampled diffraction intensities alone. *Phys. Rev. Lett.*, 95:085503, Aug 2005.
- Nimisha, T. M., Kumar Singh, A., and Rajagopalan, A. N. Blur-invariant deep learning for blind-deblurring. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- Oppenheim, A. V. and Lim, J. S. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981.
- Parikh, N. and Boyd, S. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- Reehorst, E. T. and Schniter, P. Regularization by denoising: Clarifications and new interpretations. *IEEE Transactions on Computational Imaging*, 5(1):52–67, 2019.

- Rivenson, Y., Zhang, Y., Günaydin, H., Teng, D., and Ozcan, A. Phase recovery and holographic image reconstruction using deep learning in neural networks. *Light: Science & Applications*, 7(2):17141–17141, February 2018.
- Rodriguez, J. A., Xu, R., Chen, C.-C., Zou, Y., and Miao, J. Oversampling smoothness: an effective algorithm for phase retrieval of noisy diffraction intensities. *Journal of Applied Crystallography*, 46(2):312–318, Apr 2013.
- Romano, Y., Elad, M., and Milanfar, P. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- Sinha, A., Lee, J., Li, S., and Barbastathis, G. Lensless computational imaging through deep learning. *Optica*, 4(9):1117–1125, Sep 2017.
- Venkatakrishnan, S. V., Bouman, C. A., and Wohlberg, B. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pp. 945–948, 2013.
- Wang, G., Giannakis, G. B., and Eldar, Y. C. Solving systems of random quadratic equations via truncated amplitude flow. *IEEE Transactions on Information Theory*, 64(2):773–794, 2018.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- Wu, Z., Sun, Y., Liu, J., and Kamilov, U. Online regularization by denoising with applications to phase retrieval. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- Zhang, H. and Liang, Y. Reshaped wirtinger flow for solving quadratic system of equations. In *Advances in Neural Information Processing Systems*, pp. 2622–2630, 2016.
- Zhang, H., Zhou, Y., Liang, Y., and Chi, Y. Reshaped wirtinger flow and incremental algorithms for solving quadratic systems of equations. *arXiv preprint arXiv:1605.07719*, 2016.
- Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017a.
- Zhang, K., Zuo, W., Gu, S., and Zhang, L. Learning deep cnn denoiser prior for image restoration. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017b.
- Zhang, K., Zuo, W., and Zhang, L. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018.