

# UC Riverside

## UC Riverside Previously Published Works

### Title

Provably Convergent Algorithms for Solving Inverse Problems Using Generative Models

### Permalink

<https://escholarship.org/uc/item/4w30q917>

### Authors

Shah, Viraj  
Hyder, Rakib  
Asif, M Salman  
[et al.](#)

### Publication Date

2021-05-13

Peer reviewed

---

# PROVABLY CONVERGENT ALGORITHMS FOR SOLVING INVERSE PROBLEMS USING GENERATIVE MODELS

---

Viraj Shah, Rakib Hyder, M. Salman Asif, and Chinmay Hegde<sup>‡§</sup>

## ABSTRACT

The traditional approach of hand-crafting priors (such as sparsity) for solving inverse problems is slowly being replaced by the use of richer learned priors (such as those modeled by deep generative networks). In this work, we study the algorithmic aspects of such a learning-based approach from a theoretical perspective. For certain generative network architectures, we establish a simple non-convex algorithmic approach that (a) theoretically enjoys linear convergence guarantees for certain linear and nonlinear inverse problems, and (b) empirically improves upon conventional techniques such as back-propagation. We support our claims with the experimental results for solving various inverse problems. We also propose an extension of our approach that can handle model mismatch (i.e., situations where the generative network prior is not exactly applicable). Together, our contributions serve as building blocks towards a principled use of generative models in inverse problems with more complete algorithmic understanding.

## 1 Introduction

### 1.1 Motivation

Inverse problems arise in a diverse range of application domains including computational imaging, optics, astrophysics, and seismic geo-exploration. In each of these applications, there is a target signal or image (or some other quantity of interest) to be obtained; a device (or some other physical process) records measurements of the target; and the goal is to reconstruct an estimate of the signal from the observations.

Let us suppose that  $x^* \in \mathbb{R}^n$  denotes the signal of interest and  $y = \mathcal{A}(x^*) \in \mathbb{R}^m$  denotes the observed measurements. The aim is to recover (an estimate of) the unknown signal  $x^*$  given  $y$  and  $\mathcal{A}$ . Based on the forward measurement operator  $\mathcal{A}$ , the inverse problem can be defined in two broad categories of linear and nonlinear problems. Many important problems in signal and image processing can be modeled with a *linear* measurement operator  $\mathcal{A}$ ; examples include *compressive sensing*, the classical problem of *super-resolution* or the problem of *image inpainting*. In case of nonlinear inverse problems, the operator  $\mathcal{A}$  exhibits a nonlinearity; examples include *phase retrieval*, *blind deconvolution*, and *de-quantization*.

When  $m < n$ , the inverse problem is ill-posed, and some kind of prior (or regularizer) is necessary to obtain a meaningful solution. A common technique used to solve ill-posed inverse problems is to seek the minimum of a

---

\*V. Shah was with Iowa State University. He is now with the ECE Department at the University of Illinois, Urbana-Champaign. (e-mail: vjshah3@illinois.edu).

†C. Hegde was with Iowa State University. He is now with the Tandon School of Engineering at New York University. (e-mail: chinmay.h@nyu.edu).

‡R. Hyder and M. Asif are with the ECE Department at the University of California Riverside. (e-mail: sasif@ece.ucr.edu).

§This work was completed when VS and CH were at Iowa State University, and were supported in part by grants CAREER CCF-1750920/2005804 and CCF-1815101, a faculty fellowship from the Black and Veatch Foundation, and an equipment donation from the NVIDIA Corporation. RH and MA were supported in part by grants CAREER CCF-2046293 and ONR N00014-19-1-2264, and equipment donation from NVIDIA. Parts of this manuscript appeared in short conference papers [1], [2], and [3].

constrained optimization problem:

$$\begin{aligned} \hat{x} &= \arg \min F(x), \\ \text{s.t. } x &\in \mathcal{S}, \end{aligned} \tag{1}$$

where  $F$  is an objective function that typically depends on  $y$  and  $\mathcal{A}$ , and  $\mathcal{S} \subseteq \mathbb{R}^n$  captures some sort of *structure* that  $x^*$  is assumed to obey.

Sparsity is a common modeling assumption, particularly in signal and image processing applications, where  $\mathcal{S}$  becomes a set of sparse vectors in some (known) basis representation. The popular framework of *compressive sensing* studies the special case where the forward measurement operator  $\mathcal{A}$  can be modeled as a linear operator that satisfies certain (restricted) stability properties; when this is the case, accurate estimation of  $x^*$  can be performed, assuming that the signal  $x^*$  is sufficiently sparse [4].

Parallel to the development of algorithms that leverage sparsity priors, the last decade has witnessed analogous approaches for other families of structural constraints. These include structured sparsity [5, 6], unions-of-subspaces [7], dictionary models [8, 9], total variation models [10], analytical transforms [11], among many others.

Lately, there has been renewed interest in prior models that are parametrically defined in terms of a *deep neural network*. We call these *generative network* models. Specifically, we define

$$\mathcal{S} = \{x \in \mathbb{R}^n \mid x = G(z), z \in \mathbb{R}^k\}$$

where  $z$  is a  $k$ -dimensional latent parameter vector and  $G$  is parameterized by the weights and biases of a  $d$ -layer neural network. One way to obtain such a model is to train a generative adversarial network [12]. Generative models have found remarkable applications in image analysis [13–16], and a well-trained generative model closely captures the notion of a signal (or image) being ‘natural’ [17]. Indeed, generative neural network learning algorithms have been successfully employed to solve inverse problems such as image super-resolution and inpainting [18, 19]. However, most of these approaches are heuristic and provable characterization of such algorithms are not readily available.

## 1.2 Contributions

Our goal in this paper is to take some initial steps towards a principled use of generative priors for inverse problems by a) proposing and analyzing the well known projected gradient descent (PGD) algorithm for solving (1) for both linear and nonlinear inverse problems; b) building a general theoretical framework for analyzing performance of such approaches from an algorithmic standpoint. Specifically, apart from providing algorithms to solve inverse problems using generative network models, we also wish to understand the algorithmic costs involved with such algorithms: how computationally challenging they are, whether they provably succeed, and how to make such models robust.

The starting point of our work is the seminal paper by [20], who study the benefits of using generative models in the context of compressive sensing. In this paper, the authors pose the estimated target as the solution to a non-convex optimization problem and establish upper bounds on the *statistical* complexity of obtaining a ‘‘good enough’’ solution. Specifically, they prove that if the generative network is a mapping  $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$  simulated by a  $d$ -layer neural network with width  $\leq n$  and with activation functions obeying certain properties, then  $m = O(kd \log n)$  random observations are sufficient to obtain a good enough reconstruction estimate. However, they do not explicitly discuss an *algorithm* to perform such non-convex optimization. Moreover, the authors do not study the *algorithmic* costs of solving the optimization problem, and standard results in non-convex optimization are sufficient to only obtain sub-linear convergence rates. In this work, we make several advances towards understanding the convergence properties of gradient descent (and related algorithms).

First: we establish a projected gradient descent (PGD) algorithm with linear convergence rates for the compressive sensing setup identical to [20], and demonstrate its empirical benefits over previous work. This constitutes **Contribution I** of this paper.

Second: we generalize this to a much wider range of *nonlinear* inverse problems. Using standard techniques, we propose a generic version of our PGD algorithm named  $\varepsilon$ -PGD for solving (1) where  $F$  is a smooth and strongly-convex objective function. We analyze this algorithm and prove its linear convergence under certain regimes of its smoothness and strong-convexity parameters. We also provide empirical results for solving nonlinear inverse problems. This forms **Contribution II** of this paper.

Third: we address the challenging inverse problem of *phase retrieval* [21, 22]. This is similar to the compressed sensing setup described above, except with the additional difficulty that only the magnitude of the observations are available. We prove that a natural variation of PGD coupled with an intermediate phase estimation step converges (locally) linearly to the true solution. This forms **Contribution III** of this paper.

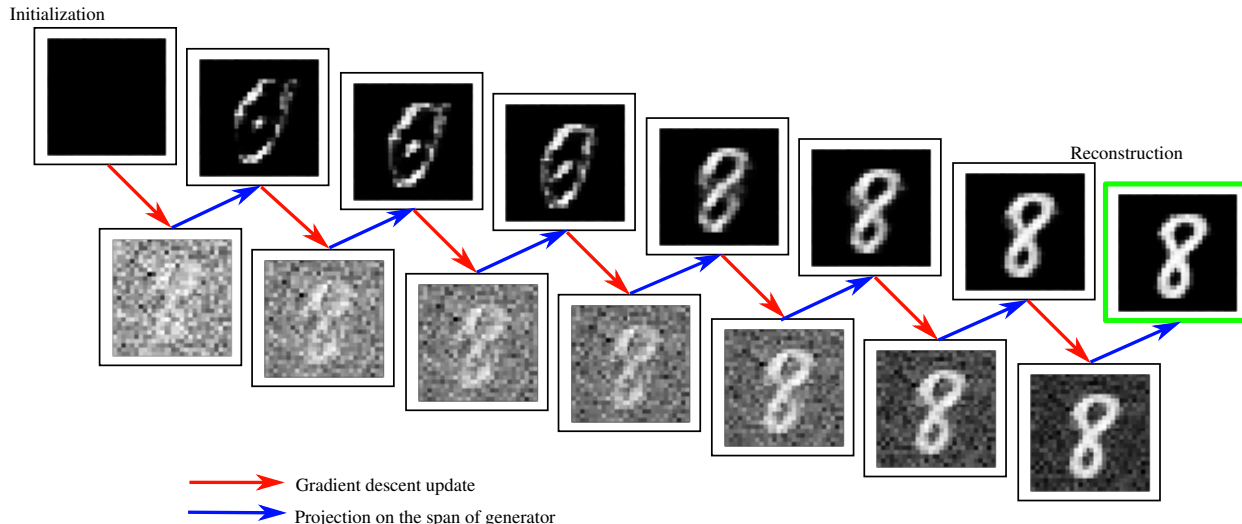


Figure 1: *Illustration of our framework. Starting from a zero vector, we perform a gradient descent update step (red arrow) and projection step (blue arrow) alternately.*

Fourth: a drawback of [20] (and our contribution I) is the inability to deal with targets that are outside the range of the generative network model. This is not merely an artifact of their analysis; generative networks are rigid in the sense that once they are learned, they are incapable of reproducing any target outside their range. (This is in contrast with other popular parametric models such as sparsity models that exhibit a “graceful decay” property in the sense that if the sparsity parameter  $s$  is large enough, such models capture all possible points in the target space.)

We address this gap, and propose an alternative algorithm using our general framework. We call it Myopic  $\varepsilon$ -PGD algorithm. It is novel, nonlinear extension of our previous work [23, 24]. Under (fairly) standard assumptions, this algorithm also can be shown to demonstrate linear convergence. This constitutes **Contribution IV** of this paper.

In summary: we complement the work of [20] and [25] by providing PGD based algorithms for solving inverse problems, and algorithmic upper bounds for the corresponding problems studied in those works. Together, our contributions serve as further building blocks towards an algorithmic theory of generative models in both linear and nonlinear inverse problems.

### 1.3 Techniques

At a high level, our algorithms are standard. The primary novelty is in their applications to generative network models, and some aspects of their theoretical analysis.

Suppose that  $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$  is the generative network model under consideration. The cornerstone of our analysis is the assumption of an  $\varepsilon$ -approximate (Euclidean) projection oracle onto the range of  $G$ . We pre-suppose the availability of a computational routine  $P_G$  that, given any vector  $x \in \mathbb{R}^n$ , can return a vector  $x' \in \text{Range}(G)$  that approximately minimizes  $\|x - x'\|_2^2$ . The availability of this oracle, of course, depends on the nature of  $G$ . Some further comments on how to heuristically approximate this oracle are in Section 5.

For a special case of linear inverse problems (and compressive sensing in particular), we assume such oracle to be simply a gradient descent routine minimizing the  $\|x - x'\|_2^2$  over the latent variable  $z$  with  $x' = G(z)$ . Though this loss function is highly non-convex due to the presence of  $G$ , we find empirically that the gradient descent (implemented via back-propagation) works very well, and can be used as a projection oracle. Our procedure is depicted in Fig. 1. We choose a zero vector as our initial estimate ( $x_0$ ), and in each iteration, we update our estimate by following the standard gradient descent update rule (red arrow in Fig. 1), followed by projection of the output onto the span of generator ( $G$ ) (blue arrow in Fig. 1).

We support this specific PGD algorithm via a rigorous theoretical analysis. We show that the final estimate at the end of  $T$  iterations is an approximate reconstruction of the original signal  $x^*$ , with very small reconstruction error; moreover, under certain sufficiency conditions on the linear operator  $\mathcal{A}$ , PGD demonstrates linear convergence, meaning that  $T = \log(1/\delta)$  is sufficient to achieve  $\delta$ -accuracy. Further, we present a series of numerical results as validation of our approach.

We also provide a direct generalization of the above approach for nonlinear inverse problems, that we call  $\varepsilon$ -PGD. We analyze this generic algorithm to show a linear convergence by assuming that the objective function in (1) obeys the Restricted Strong Convexity/Smoothness assumptions [26]. With this assumption, the proof of convergence follows from a straightforward modification of the proof given in [27]. Through our analysis, it indeed can be seen that the PGD algorithm for linear inverse problems is in fact a special case of  $\varepsilon$ -PGD.

The fourth algorithm (Myopic  $\varepsilon$ -PGD) is a novel approach for handling model mismatch in the target. The main idea (following the lead of [25]) is to pose the target  $x^*$  as the superposition of two components:  $x^* = G(z) + \nu$ , where  $\nu$  can be viewed as an ‘‘innovation’’ term that is  $s$ -sparse in some fixed, known basis  $B$ . The goal is now to recover both  $G(z)$  and  $\nu$ . This is reminiscent of the problem of source separation or signal demixing [28], and in our previous work [24, 29] we proposed greedy iterative algorithms for solving such demixing problems. We extend this work by proving a nonlinear extension, together with a new analysis, of the algorithm proposed in [24].

## 2 Background and Related Work

### 2.1 Inverse problems

The study of solving inverse problems has a long history. As discussed above, the general approach to solve an ill-posed inverse problem is to assume that the target signal/image obeys a *prior*. Classical methods mainly used hand-crafted signal priors to distinguish ‘natural’ signals from the infinite set of feasible solutions. The prior can be encoded in the form of either a constraint set (as in Eq. (1)) or an extra regularization penalty. Several methods (including [30–32]) employ sparsity priors to solve *linear* inverse problems such as denoising, super-resolution, and inpainting. Despite their successful practical and theoretical results, all such hand-designed priors often fail to restrict the solution space only to natural images, and it is easily possible to generate signals satisfying the prior but do not resemble natural data.

### 2.2 Neural network models

The last few years have witnessed the emergence of trained *neural networks* for solving such problems. The main idea is to eschew hand-crafting any priors, and instead *learn* an end-to-end mapping from the measurement space to the image space. This mapping is simulated via a deep neural network, whose weights are learned from a large dataset of input-output training examples [33]. The works [34–41] have used this approach to solve several types of inverse problems, and has met with considerable success. However, the major limitations are that a new network has to be trained for each new linear inverse problem; moreover, most of these methods lack concrete theoretical guarantees. An exception of this line of work is the powerful framework of [42], which does *not* require retraining for each new problem; however, this too is not accompanied by theoretical analysis of statistical and computational costs.

### 2.3 Generative networks

A special class of neural networks that attempt to directly model the distribution of the input training samples are known as generative adversarial training networks, or GANs [12]. GANs have been shown to provide visually striking results [17, 43–45]. The use of GANs to solve linear inverse problems was advocated in [20]. Specifically, given (noisy) linear observations  $y = Ax^* + e$  of a signal  $x^* \in \mathbb{R}^n$ , assuming that  $x^*$  belongs to the range of a generative network  $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ , this approach constructs the reconstructed estimate  $\hat{x}$  as follows:

$$\hat{z} = \arg \min_{z \in \mathbb{R}^k} \|y - AG(z)\|_2^2, \quad \hat{x} = G(\hat{z})$$

If the observation matrix  $A \in \mathbb{R}^{m \times n}$  comprises  $m = O(kd \log n)$  i.i.d. Gaussian measurements, then together with regularity assumptions on the generative network, they prove that the solution  $\hat{x}$  satisfies:

$$\|x^* - \hat{x}\|_2 \leq C\|e\|_2.$$

for some constant  $C$  that can be reliably upper-bounded. In particular, in the absence of noise the recovery of  $x^*$  is exact. However, there is no discussion of how computationally expensive this procedure is. Observe that the above minimization is highly non-convex (since for any reasonable neural network,  $G$  is a non-convex function) and possibly also non-smooth (if the neural network contains non-smooth activation functions, such as rectified linear units, or ReLUs). More recently, [46] improve upon the approach in [20] for solving more general nonlinear inverse problems (in particular, any inverse problem that has a computable derivative). Also, [47] have analyzed the above problem for the *untrained* case where the weights of the generative model  $G$  obeys certain randomness assumptions. See, also, [48].

Under similar statistical assumptions as [20], the work of [1] provably establishes a linear convergence rate, provided that a projection oracle (on to the range of  $G$ ) is available, but only for the special case of compressive sensing. Our generalized result (Contribution II) extends this algorithm (and analysis) to more general nonlinear inverse problems.

More recently, [49] proposed a method that learns the a network-based projector for use in the PGD algorithm, making the projection step faster computationally. However, their theoretical result assumes the learned projector to be  $\delta$ -approximate, indicating that the effective training of the projector is crucial for the success of their method posing an additional challenge.

## 2.4 Phase retrieval

The phase retrieval problem has been extensively studied over the last few decades [50–52] and it appears in several applications, including optical imaging [51, 53], microscopy [54], and X-ray crystallography [55]. Phase retrieval is a non-convex problem and classical solution methods rely on alternating projection heuristics; examples include Gerchberg-Saxton [50] and Fienup [51]. In recent years, lifting-based methods were introduced that reformulate phase retrieval as a semidefinite program. [52]. Subsequently, non-convex methods have been proposed for solving phase retrieval problem with theoretical performance guarantees [56–62]. Most of the non-convex methods rely on estimating a good initial solution via the so-called *spectral initialization* method. A number of methods for solving phase retrieval using trained neural networks have been recently proposed [41, 63, 64].

## 2.5 Model mismatch

A limitation of most generative network models is that they can only reproduce estimates that are within their range; adding more observations or tweaking algorithmic parameters are completely ineffective if a generative network model is presented with a target that is far away from the range of the model. To resolve this type of model mismatch, the authors of [25] propose to model the signal  $x^*$  as the superposition of two components: a “base” signal  $u = G(z)$ , and an “innovation” signal  $v = B\nu$ , where  $B$  is a known ortho-basis and  $\nu$  is an  $l$ -sparse vector. In the context of compressive sensing, the authors of [25] solve a sparsity-regularized loss minimization problem:

$$(\hat{z}, \hat{\nu}) = \arg \min_{z, \nu} \|B^T v\|_1 + \lambda \|y - A(G(z) + v)\|_2^2.$$

and prove that the reconstructed estimate  $\hat{x} = G(\hat{z}) + \hat{\nu}$  is close enough to  $x$  provided  $m = O((k + l)d \log n)$  measurements are sufficient. However, as before, the algorithmic costs of solving the above problem are not discussed. Our third main result (Contribution III) proposes a new algorithm for dealing with model mismatches in generative network modeling, together with an analysis of its convergence and iteration complexity.

# 3 Main Algorithms and Analysis

Let us first establish some notational conventions. Below,  $\|\cdot\|$  will denote the Euclidean norm unless explicitly specified. We use  $O(\cdot)$ -notation in several places in order to avoid duplication of constants. We use  $F(\cdot)$  to denote a (scalar) objective function.

## 3.1 Contribution I: Solving linear inverse problems

Let  $\mathcal{S} \subseteq \mathbb{R}^n$  be the set of ‘natural’ images in data space with a vector  $x^* \in \mathcal{S}$ . We consider an ill-posed linear inverse problem (2) with the linear operator  $\mathcal{A}(x) = Ax$ , where  $A$  is a Gaussian random matrix. For simplicity, we do not consider the additive noise term.

$$y = Ax^*, \tag{2}$$

To solve for  $\hat{x}$  (estimate of  $x^*$ ), we choose Euclidean measurement error as the loss function  $F(\cdot)$  in Eqn. (1). Therefore, given  $y$  and  $A$ , we seek

$$\hat{x} = \arg \min_{x \in \mathcal{S}} \|y - Ax\|^2. \tag{3}$$

### 3.1.1 Algorithm

Our algorithm is described in Alg. 1. We assume that our trained generator network ( $G$ ) well approximates the high-dimensional probability distribution of the set  $\mathcal{S}$ . With this assumption, we limit our search for  $\hat{x}$  only to the range of

---

**Algorithm 1** PGD

---

1: **Inputs:**  $y, A, G, T$ , **Output:**  $\hat{x}$   
2:  $x_0 \leftarrow \mathbf{0}$   
3: **while**  $t < T$  **do**  
4:      $w_t \leftarrow x_t + \eta A^T (y - Ax_t)$   
5:      $x_{t+1} \leftarrow G(\arg \min_z \|w_t - G(z)\|)$   
6:      $t \leftarrow t + 1$   
7:  $\hat{x} \leftarrow x_T$

---

the generator function ( $G(z)$ ). The function  $G$  is assumed to be differentiable, and hence we use back-propagation for calculating the gradients of the loss functions involving  $G$  for gradient descent updates.

The optimization problem in Eqn. 3 is similar to a least squares estimation problem, and a typical approach to solve such problems is to use gradient descent. However, the candidate solutions obtained after each gradient descent update need not represent a ‘natural’ image and may not belong to set  $\mathcal{S}$ . We solve this limitation by projecting the candidate solution on the range of the generator function after each gradient descent update.

Thus, in each iteration of our proposed algorithm 1, two steps are performed in alternation: a gradient descent update step and a projection step. The first step is simply an application of a gradient descent update rule on the loss function  $F(\cdot)$  with the learning rate  $\eta$ . In projection step, we minimize the projection loss by gradient descent updates with learning rate  $\eta_{in}$ :

$$P_G(w_t) := G\left(\arg \min_z \|w_t - G(z)\|\right),$$

Though the projection loss function is highly non-convex due to the presence of  $G$ , we find empirically that the gradient descent (implemented via back-propagation) works very well. Thus, the gradient descent based minimization serves as a projection oracle in this case. In each of the  $T$  iterations, we run  $T_{in}$  gradient descent updates for calculating the projection. Therefore,  $T \times T_{in}$  is the total number of gradient descent updates required in our approach.

### 3.1.2 Analysis

Drawing parallels with standard compressive sensing theory, in our case, we need to ensure that the difference vector of any two signals in the set  $\mathcal{S}$  lies away from the nullspace of the matrix  $A$ . This condition is encoded via the S-REC (Set Restricted Eigenvalue Condition) as defined and established in [20]. We slightly modify this condition and present it in the form of squared  $\ell_2$ -norm :

**Definition 1** Let  $\mathcal{S} \in \mathbb{R}^n$ .  $A$  is  $m \times n$  matrix. For parameters  $\gamma > 0$ ,  $\delta \geq 0$ , matrix  $A$  is said to satisfy the S-REC( $\mathcal{S}, \gamma, \delta$ ) if,

$$\|A(x_1 - x_2)\|^2 \geq \gamma \|x_1 - x_2\|^2 - \delta,$$

for  $\forall x_1, x_2 \in \mathcal{S}$ .

Further, based on [65, 66], we propose the following theorem about the convergence of our algorithm:

**Theorem 1** Let  $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$  be a differentiable generator with  $d$  layers and range  $\mathcal{S}$ . Let  $A$  be a random Gaussian matrix with  $A_{i,j} \sim N(0, 1/m)$  with  $m \geq C(kd \log n)$  for some positive constant  $C$ . Then, for every vector  $x^* \in \mathcal{S}$ , the sequence  $(x_t)$  defined by the algorithm PGD [1] exhibits linear convergence for a carefully chosen range of stepsizes  $\eta$ .

**Proof:** Suppose  $F(\cdot)$  is the squared error loss function as defined above. Then, we have:

$$\begin{aligned} & F(x_{t+1}) - F(x_t) \\ &= \|Ax_{t+1}\|^2 - 2\langle y, Ax_{t+1} \rangle + 2\langle y, Ax_t \rangle - \|Ax_t\|^2, \\ &= \|Ax_{t+1} - Ax_t\|^2 + 2\langle x_t - x_{t+1}, A^T A(x^* - x_t) \rangle. \end{aligned}$$

Substituting  $y = Ax^*$  and rearranging yields,

$$\begin{aligned} 2\langle x_t - x_{t+1}, A^T (y - Ax_t) \rangle &= F(x_{t+1}) - F(x_t) \\ &\quad - \|Ax_{t+1} - Ax_t\|^2. \end{aligned} \tag{4}$$

Define:

$$w_t := x_t + \eta A^T(y - Ax_t) = x_t + \eta A^T A(x^* - x_t)$$

Then, by definition of the projection operator  $P_G$ , the vector  $x_{t+1}$  is a better (or equally good) approximation to  $w$  as the true image  $x^*$ . Therefore, we have:

$$\|x_{t+1} - w_t\|^2 \leq \|x^* - w_t\|^2.$$

Substituting for  $w_t$  and expanding both sides, we get:

$$\begin{aligned} & \|x_{t+1} - x_t\|^2 - 2\eta \langle x_{t+1} - x_t, A^T(y - Ax_t) \rangle \\ & \leq \|x^* - x_t\|^2 - 2\eta \langle x^* - x_t, A^T(y - Ax_t) \rangle. \end{aligned}$$

Substituting  $y = Ax^*$  and rearranging yields,

$$\begin{aligned} & 2\langle x_t - x_{t+1}, A^T(y - Ax_t) \rangle \\ & \leq \frac{1}{\eta} \|x^* - x_t\|^2 - \frac{1}{\eta} \|x_{t+1} - x_t\|^2 - 2F(x_t). \end{aligned} \quad (5)$$

We now use 4 and 5 to obtain,

$$\begin{aligned} F(x_{t+1}) & \leq \frac{1}{\eta} \|x^* - x_t\|^2 - F(x_t) \\ & \quad - \left( \frac{1}{\eta} \|x_{t+1} - x_t\|^2 - \|Ax_{t+1} - Ax_t\|^2 \right). \end{aligned} \quad (6)$$

Now, since  $A$  is a random Gaussian with sufficiently many rows, it satisfies the S-REC [20]:

$$\|A(x_1 - x_2)\|^2 \geq \gamma \|x_1 - x_2\|^2 - \delta.$$

As  $x^*$ ,  $x_t$  and  $x_{t+1}$  are ‘natural’ vectors,

$$\frac{1}{\eta} \|x^* - x_t\|^2 \leq \frac{1}{\eta\gamma} \|y - Ax_t\|^2 + \frac{\delta}{\eta\gamma}. \quad (7)$$

Substituting 7 in 6,

$$\begin{aligned} F(x_{t+1}) & \leq \left( \frac{1}{\eta\gamma} - 1 \right) F(x_t) \\ & \quad - \left( \frac{1}{\eta} \|x_{t+1} - x_t\|^2 - \|Ax_{t+1} - Ax_t\|^2 \right) + \frac{\delta}{\eta\gamma}. \end{aligned}$$

Moreover, as shown in [20], we have  $\|Av\| \leq \rho \|v\|$  with high probability for any  $v$ . Therefore, we write:

$$\begin{aligned} & \|Ax_{t+1} - Ax_t\|^2 \leq \rho^2 \|x_{t+1} - x_t\|^2, \\ & \|Ax_{t+1} - Ax_t\|^2 - \frac{1}{\eta} \|x_{t+1} - x_t\|^2 \leq \left( \rho^2 - \frac{1}{\eta} \right) \|x_{t+1} - x_t\|^2. \end{aligned}$$

Let us choose learning rate  $(\eta)$  such that  $\frac{1}{2\gamma} < \eta < \frac{1}{\gamma}$ . We also have  $\rho^2 \leq \gamma$ . Combining both, we get  $\rho^2 < \frac{1}{\eta}$ , which makes the L.H.S. in the above equation negative. Therefore,

$$F(x_{t+1}) \leq \left( \frac{1}{\eta\gamma} - 1 \right) F(x_t) + \frac{\delta}{\eta\gamma},$$

where  $\delta$  is inversely proportional to the number of measurements  $m$  [20]. If the slack parameter  $\delta$  is small enough then it can be ignored. Also,  $\frac{1}{2\gamma} < \eta < \frac{1}{\gamma}$  yields,

$$0 < \left( \frac{1}{\eta\gamma} - 1 \right) < 1.$$

Hence,

$$F(x_{t+1}) \leq \alpha F(x_t) + \delta; \quad 0 < \alpha < 1, \quad (8)$$

i.e., the sequence  $F(x_t)$  converges linearly upto a neighborhood of radius  $O(\delta)$ . In the noiseless case, this also implies that  $x_{t+1}$  converges to a small neighborhood around  $x^*$ .  $\square$



### 3.2 Contribution II: Solving nonlinear inverse problems

We now present the generic version of PGD algorithm suitable for large class of nonlinear inverse problems by generalizing for the loss function  $F(\cdot)$  and the projection oracle  $P_G$ .

We denote the  $F(\cdot)$  to be a scalar function with continuous gradient, and assume that  $F$  has a continuous gradient  $\nabla F = \left(\frac{\partial F}{\partial x_i}\right)_{i=1}^n$  which can be evaluated at any point  $x \in \mathbb{R}^n$ .

Recall that we wish to solve the problem

$$\begin{aligned} \hat{x} &= \arg \min F(x), \\ \text{s. t. } &x \in \text{Range}(G), \end{aligned} \tag{9}$$

where  $G$  is a generative network. To do so, we now employ a generalized version of *projected gradient descent* algorithm using the  $\varepsilon$ -approximate projection oracle for  $G$ . The algorithm is described in Alg. 2.

We define the  $\varepsilon$ -approximate projection oracle  $P_G$  as,

**Definition 2 (Approximate projection)** A function  $P_G : \mathbb{R}^n \rightarrow \text{Range}(G)$  is an  $\varepsilon$ -approximate projection oracle if for all  $x \in \mathbb{R}^n$ ,  $P_G(x)$  obeys:

$$\|x - P_G(x)\|_2^2 \leq \min_{z \in \mathbb{R}^k} \|x - G(z)\|_2^2 + \varepsilon.$$

We will assume that for any given generative network  $G$  of interest, such a function  $P_G$  exists and is computationally tractable<sup>5</sup>. Here,  $\varepsilon > 0$  is a parameter that is known a priori.

In contrast to our previous analysis, here we introduce more general restriction conditions on the  $F(\cdot)$ :

**Definition 3 (Restricted Strong Convexity/Smoothness)** Assume that  $F$  satisfies  $\forall x, y \in S$ :

$$\frac{\alpha}{2} \|x - y\|_2^2 \leq F(y) - F(x) - \langle \nabla F(x), y - x \rangle \leq \frac{\beta}{2} \|x - y\|_2^2.$$

for positive constants  $\alpha, \beta$ .

This assumption is by now standard; see [26, 27] for in-depth discussions. This means that the objective function is strongly convex / strongly smooth along certain directions in the parameter space (in particular, those restricted to the set  $S$  of interest). The parameter  $\alpha > 0$  is called the restricted strong convexity (RSC) constant, while the parameter  $\beta > 0$  is called the restricted strong smoothness (RSS) constant. Clearly,  $\beta \geq \alpha$ . In fact, throughout the paper, we assume that  $1 \leq \frac{\beta}{\alpha} < 2$ , which is a fairly stringent assumption but again, one that we do not know at the moment how to relax.

**Definition 4 (Incoherence)** A basis  $B$  and  $\text{Range}(G)$  are called  $\mu$ -incoherent if for all  $u, u' \in \text{Range}(G)$  and all  $v, v' \in \text{Span}(B)$ , we have:

$$|\langle u - u', v - v' \rangle| \leq \mu \|u - u'\|_2 \|v - v'\|_2.$$

for some parameter  $0 < \mu < 1$ .

**Remark 1** In addition to the above, we will make the following assumptions in order to aid the analysis. Below,  $\gamma$  and  $\Delta$  are positive constants.

- $\|\nabla F(x^*)\|_2 \leq \gamma$ .
- $\text{diam}(\text{Range}(G)) = \Delta$ .
- $\gamma\Delta \leq O(\varepsilon)$ .

Some comments about these assumptions may be warranted. The first says that the gradient of the loss at at the minimizer is small. The second says that the range of  $G$  is compact. The third links the parameters from the first two assumptions.

We obtain the following theoretical result:

<sup>5</sup>This may be a very strong assumption, but at the moment we do not know how to relax this in the very general case. Indeed, the computational complexity of our proposed algorithms is proportional to the complexity of such a projection oracle. For preliminary advances, see [67]

---

**Algorithm 2**  $\varepsilon$ -PGD

---

1: **Inputs:**  $y, T, \nabla$ ; **Output:**  $\hat{x}$   
2:  $x_0 \leftarrow \mathbf{0}$   
3: **while**  $t < T$  **do**  
4:      $w_t \leftarrow x_t - \eta \nabla F(x_t)$   
5:      $x_{t+1} \leftarrow P_G(w_t)$   
6:      $t \leftarrow t + 1$   
7:  $\hat{x} \leftarrow x_T$

---

**Theorem 2** *If  $F$  satisfies RSC/RSS over  $\text{Range}(G)$  with constants  $\alpha$  and  $\beta$ , then  $\varepsilon$ -PGD (Alg. 2) converges linearly up to a ball of radius  $O(\gamma\Delta) \approx O(\varepsilon)$ .*

$$F(x_{t+1}) - F(x^*) \leq \left(\frac{\beta}{\alpha} - 1\right) (F(x_t) - F(x^*)) + O(\varepsilon).$$

**Proof:** The proof is a minor modification of that in [27]. For simplicity we will assume that  $\|\cdot\|$  refers to the Euclidean norm. Let us suppose that the step size  $\eta = \frac{1}{\beta}$ . Define

$$w_t = x_t - \eta \nabla F(x_t).$$

By invoking RSS, we get:

$$\begin{aligned} & F(x_{t+1}) - F(x_t) \\ & \leq \langle \nabla F(x_t), x_{t+1} - x_t \rangle + \frac{\beta}{2} \|x_{t+1} - x_t\|^2 \\ & = \frac{1}{\eta} \langle x_t - w_t, x_{t+1} - x_t \rangle + \frac{\beta}{2} \|x_{t+1} - x_t\|^2 \\ & = \frac{\beta}{2} \left( \|x_{t+1} - x_t\|^2 + 2 \langle x_t - w_t, x_{t+1} - x_t \rangle + \|x_t - w_t\|^2 \right) \\ & \quad - \frac{\beta}{2} \|x_t - w_t\|^2 \\ & = \frac{\beta}{2} \left( \|x_{t+1} - w_t\|^2 - \|x_t - w_t\|^2 \right), \end{aligned}$$

where the last few steps are consequences of straightforward algebraic manipulation.

Now, since  $x_{t+1}$  is an  $\varepsilon$ -approximate projection of  $w_t$  onto  $\text{Range}(G)$  and  $x^* \in \text{Range}(G)$ , we have:

$$\|x_{t+1} - w_t\|^2 \leq \|x^* - w_t\|^2 + \varepsilon.$$

Therefore, we get:

$$\begin{aligned} & F(x_{t+1}) - F(x_t) \\ & \leq \frac{\beta}{2} \left( \|x^* - w_t\|^2 - \|x_t - w_t\|^2 \right) + \frac{\beta\varepsilon}{2} \\ & = \frac{\beta}{2} \left( \|x^* - x_t + \eta \nabla F(x_t)\|^2 - \|\eta \nabla F(x_t)\|^2 \right) + \frac{\beta\varepsilon}{2} \\ & = \frac{\beta}{2} \left( \|x^* - x_t\|^2 + 2\eta \langle x^* - x_t, \nabla F(x_t) \rangle \right) + \frac{\beta\varepsilon}{2} \\ & = \frac{\beta}{2} \|x^* - x_t\|^2 + \langle x^* - x_t, \nabla F(x_t) \rangle + \frac{\beta\varepsilon}{2}. \end{aligned}$$

However, due to RSC, we have:

$$\begin{aligned} \frac{\alpha}{2} \|x^* - x_t\|^2 & \leq F(x^*) - F(x_t) - \langle x^* - x_t, \nabla F(x_t) \rangle, \\ \langle x^* - x_t, \nabla F(x_t) \rangle & \leq F(x^*) - F(x_t) - \frac{\alpha}{2} \|x^* - x_t\|^2. \end{aligned}$$

Therefore,

$$\begin{aligned}
& F(x_{t+1}) - F(x_t) \\
& \leq \frac{\beta - \alpha}{2} \|x^* - x_t\|^2 + F(x^*) - F(x_t) + \frac{\beta\varepsilon}{2} \\
& \leq \frac{\beta - \alpha}{2} \cdot \frac{2}{\alpha} (F(x_t) - F(x^*) - \langle x_t - x^*, \nabla F(x^*) \rangle) \\
& \quad + F(x^*) - F(x_t) + \frac{\beta\varepsilon}{2} \\
& \leq \left(2 - \frac{\beta}{\alpha}\right) (F(x^*) - F(x_t)) + \frac{\beta - \alpha}{\alpha} \gamma \Delta + \frac{\beta\varepsilon}{2},
\end{aligned}$$

where the last inequality follows from Cauchy-Schwartz and the assumptions on  $\|\nabla F(x^*)\|$  and the diameter of  $\text{Range}(G)$ . Further, by assumption,  $\gamma\Delta \leq O(\varepsilon)$ . Rearranging terms, we get:

$$F(x_{t+1}) - F(x^*) \leq \left(\frac{\beta}{\alpha} - 1\right) (F(x_t) - F(x^*)) + C\varepsilon.$$

for some constant  $C > 0$ . □

This theorem asserts that the distance between the objective function at any iteration to the optimum *decreases by a constant factor* in every iteration. (The decay factor is  $\frac{\beta}{\alpha} - 1$ , which by assumption is a number between 0 and 1). Therefore, we immediately obtain linear convergence of  $\varepsilon$ -PGD up to a ball of radius  $O(\varepsilon)$ :

**Corollary 1** *After  $T = O(\log \frac{F(x_0) - F(x^*)}{\varepsilon})$  iterations,  $F(x_T) \leq F(x^*) + O(\varepsilon)$ .*

Therefore, the overall running time can be bounded as follows:

$$\text{Runtime} \leq (T_{\varepsilon\text{-PROJ}} + T_{\nabla}) \times \log(1/\varepsilon).$$

It is noticeable that the analysis of PGD algorithm for linear problem is a special case of the generalized analysis given by Theorem 2. That is because once we set the  $F(\cdot)$  as defined in Eq. (3), the RSC for  $F(\cdot)$  can be obtained through the S-REC condition from Eq. (4). Similarly, we can use the upper bound of the spectral norm for the Gaussian matrix  $A$  to obtain RSS for  $F(\cdot)$ .

In Sec. 4, we provide empirical results for solving nonlinear inverse problems using Alg. 2. Specifically, we consider two nonlinear forward models: a sinusoidal model with  $\mathcal{A}(x^*) = Ax^* + \sin(Ax^*)$ ; and a sigmoid model with  $\mathcal{A}(x^*) = \text{sigmoid}(Ax^*) = \frac{1}{1 + \exp(-Ax^*)}$ . While we use the  $L_2$ -loss as a loss function in the case of sinusoidal model, for the sigmoid nonlinearity, we use a loss function specified as:

$$F(x) = \frac{1}{m} \sum_{i=1}^m (\Theta(a_i^T x) - y_i a_i^T x),$$

where,  $\Theta(\cdot)$  is integral of  $\mathcal{A}(\cdot)$ , and  $a_i$  represents the rows of the measurement matrix  $A$ . The gradient of the loss can be calculated in closed form:

$$\nabla F(x^*) = \frac{1}{m} A^T (\text{sigmoid}(Ax) - y). \quad (10)$$

Such choice of the loss function is inspired by the problem of single-index model (SIM) estimation [68]. [69] also advocates the usage of such a loss function.

### 3.3 Contribution III: Phase retrieval

We now propose a method to solve a different category of nonlinear inverse problems using our overall broad approach. Specifically, we will solve inverse problems where the observations are of the form:

$$y = |Ax| + \text{noise},$$

i.e., only magnitude information of the phaseless measurements are retained. This is the well-known phase retrieval setting [21, 22]; here the added twist is that  $x$  is assumed to obey a generative prior.

One approach to solving this problem is to cast the recovery as a nonlinear optimization problem based on a suitable loss function  $F(x)$  that measures the (squared) difference between the measurements and the purported phaseless

---

**Algorithm 3** PHASE-PGD

---

1: **Inputs:**  $y, A, G, T$ , **Output:**  $\hat{x}$   
2: Choose an initial point  $x_0 \in \mathbb{R}^n$   
3: **for**  $t = 1, \dots, T$  **do**  
4:    $p_{t-1} \leftarrow \text{sign}(Ax_{t-1})$   
5:    $w_{t-1} \leftarrow x_{t-1} + \eta A^T(y \odot p_{t-1} - Ax_{t-1})$   
6:    $x_t \leftarrow \mathcal{P}_G(w_{t-1}) = G(\arg \min_z \|w_{t-1} - G(z)\|)$   
7:  $\hat{x} \leftarrow x_T$

---

observations. We will slightly depart from this and instead propose a new algorithm (Phase-PGD) for solving the above problem. See Alg. 3. In each iteration of the Phase-PGD algorithm, three steps are performed: a phase update step, a gradient descent update step, and a projection step.

The first step is to calculate the phase of  $Ax$ . For real  $A$  and  $x$ , at the  $t^{\text{th}}$  iteration, we update the phase estimate:

$$p_t = \text{phase}(Ax_t) := \text{sign}(Ax_t).$$

After calculating the phase vector  $p$ , we can use an element-wise product between  $p$  and  $y$  as an estimate of linear measurements and convert the phase retrieval problem into a linear inverse problem.

The second step is simply an application of a gradient descent update rule on the loss function  $f(\cdot)$  which is given as:

$$f(x) := \|y \odot p - Ax\|^2.$$

Thus, the gradient descent update at the  $t^{\text{th}}$  iteration is given by:

$$w_t \leftarrow x_t + \eta A^T(y \odot p_t - Ax_t),$$

where  $\eta$  is the learning rate.

The third step is the projection step, in which we aim to find an image from the span of the generator,  $\mathcal{M}$  which is closest to our current estimate  $w_t$ . We define the projection operator  $\mathcal{P}_G$  as follows:

$$\mathcal{P}_G(w_t) := G\left(\arg \min_z L_{in}(z)\right),$$

where  $L_{in}$  is the inner loss function defined as,

$$L_{in}(z) := \|w_t - G(z)\|^2.$$

We solve the inner optimization problem by running gradient descent with  $T_{in}$  number of updates on  $L_{in}(z)$ .

We now analyze this algorithm. This part of the algorithm is described in Lines 3-7 of Algorithm 3. We can prove that provided a good initial estimate ( $x_0$ ), Phase-PGD provably converges to  $x^*$ . The high level intuition is as follows. Ignoring the noise, the observation model for phase retrieval can be restated as follows:

$$\text{sign}(\langle a_i, x^* \rangle) \odot y_i = \langle a_i, x^* \rangle,$$

for all  $i = \{1, 2, \dots, m\}$ . To ease notation, denote the *phase vector*  $p \in \mathbb{R}^m$  as a vector that contains the unknown signs of the measurements, i.e.,  $p_i = \text{sign}(\langle a_i, x \rangle)$  for all  $i = \{1, 2, \dots, m\}$ . Let  $p^*$  denote the true phase vector and let  $\mathcal{P}$  denote the set of all phase vectors, i.e.  $\mathcal{P} = \{p : p_i = \pm 1, \forall i\}$ . Then our measurement model gets modified as:

$$p^* \odot y = Ax^*.$$

Therefore, the recovery of  $x^*$  can be posed as a (non-convex) optimization problem:

$$\min_{x \in \mathcal{M}, p \in \mathcal{P}} \|Ax - p \odot y\|^2 \tag{11}$$

To solve this problem, we alternate between estimating  $p$  and  $x$ . We perform two estimation steps:

- (a) if we fix the signal estimate  $x$ , then the minimizer  $p \in \mathcal{P}$  is given in closed form as:

$$p = \text{sign}(Ax), \tag{12}$$

(b) and if we fix the phase vector  $\mathbf{p}$ , the signal vector  $x \in \mathcal{M}$  can be obtained by solving:

$$\min_{x \in \mathcal{M}} \|Ax - p \odot y\|_2. \quad (13)$$

We now analyze our proposed descent scheme. We obtain:

**Theorem 3** *Suppose we have an initialization  $x_0 \in \mathcal{M}$  satisfying  $\text{dist}(x_0, x^*) \leq \delta_0 \|x^*\|_2$ , for  $0 < \delta_0 < 1$ , and suppose the number of (Gaussian) measurements,*

$$m > C(kd \log n),$$

*for some large enough constant  $C$ . Then with high probability the iterates  $x_{t+1}$  of Algorithm 3, satisfy:*

$$\text{dist}(x_{t+1}, x^*) \leq \rho \text{dist}(x_t, x^*), \quad (14)$$

*where  $x_t, x_{t+1}, x^* \in \mathcal{M}$ , and  $0 < \rho < 1$  is a constant.*

**Proof sketch:** The high level idea behind the proof is that with a  $\delta$ -ball around the true signal  $x^*$ , the ‘‘phase noise’’ can be suitably bounded in terms of a constant times the signal estimation error. To be more precise, suppose that  $z^* = Ax^* = p^* \odot y$ . Then, at any iteration  $t$ , we have:

$$\begin{aligned} z_t &= p_t \odot y \\ &= p^* \odot y + (p_t - p^*) \odot y \\ &= z^* + e_t, \end{aligned}$$

where  $e_t$  can be viewed as the ‘‘phase noise’’. Now, examining Line 6 of the above algorithm, we have that  $x_t$  is the output of Phase-PGD after  $t$  iterations. An ‘‘unpacking’’ argument similar to the one in [70, 71] indicates that:

$$\|x_t - x^*\| \leq \alpha \|x_{t-1} - x^*\| + \beta \|e_t\|,$$

where  $\alpha$  is a small enough constant. We will show that  $\|e_t\|$  can be also bounded in terms of  $\|x_{t-1} - x^*\|$ , via Lemma 4 below. Consequently:

$$\|x_t - x^*\| \leq \rho \|x_{t-1} - x^*\|,$$

where  $\rho$  is a small enough constant.

We therefore achieve a per-step error reduction scheme if the initial estimate  $x_0$  satisfies  $\|x_0 - x^*\| \leq \delta_0 \|x^*\|$ . This result can be trivially extended to the case where the initial estimate  $x_0$  satisfies  $\|x_0 + x^*\| \leq \delta_0 \|x^*\|$ , hence giving the convergence criterion of the form (for  $\rho < 1$ ):

$$\text{dist}(x_t, x^*) \leq \rho \text{dist}(x_{t-1}, x^*).$$

We now state Lemma 4 without proof; this is a straightforward adaptation of the covering number analysis of analysis of [20].

**Theorem 4** *Suppose that the generator network model  $G(\cdot)$  is comprised of  $d$  layers of neurons with ReLU activation functions and weight matrices with bounded operator norms. As long as the initial estimate is a small distance away from the true signal  $x^* \in \mathcal{M}$  ( i.e.  $\text{dist}(x_0, x^*) \leq \delta \|x^*\|$ ) and subsequently,  $\text{dist}(x_t, x^*) \leq \delta \|x^*\|$ , where  $x_t$  is the  $t^{\text{th}}$  update of Algorithm 3, then the following bound holds for any  $t \geq 0$ :*

$$\|e_{t+1}\| \leq \rho_1 \|x_t - x^*\|,$$

*with high probability, as long as  $m > C(kd \log n)$  and  $\rho_1 < 1$  is a constant.*

### 3.4 Contribution IV: Addressing signal model mismatch

A key assumption in all the above discussion has been that the underlying signal is well-modeled by the generator  $G$ . This, of course, is idealistic; natural signals (and images) exhibit a wide variety of features, some of which may not be modeled by (even) a well-trained generator.

We now generalize the  $\varepsilon$ -PGD algorithm to handle situations involving signal model mismatch. Assume that the target signal can be decomposed as:

$$x^* = G(z) + v,$$

where  $\|B^T v\|_0 \leq l \ll n$  for some ortho-basis  $B$ .

---

**Algorithm 4** MYOPIC  $\varepsilon$ -PGD
 

---

```

1: Inputs:  $y, T, \nabla$ ; Output:  $\hat{x}$ 
2:  $x_0, u_0, v_0 \leftarrow \mathbf{0}$ 
3: while  $t < T$  do
4:    $u_{t+1} = P_G(u_t - \eta \nabla_x F(x_t))$ 
5:    $v_{t+1} = \text{Thresh}_{B,l}(v_t - \eta \nabla_x F(x_t))$ 
6:    $x_{t+1} = u_{t+1} + v_{t+1}$ 
7:    $t \leftarrow t + 1$ 
8:  $\hat{x} \leftarrow x_T$ 

```

---

For this model, we attempt to solve a (slightly) different optimization problem:

$$\hat{x} = \arg \min F(x), \quad (15)$$

$$\text{s. t. } x = G(z) + v,$$

$$\|B^T v\|_0 \leq l. \quad (16)$$

We propose a new algorithm to solve this problem that we call *Myoptic  $\varepsilon$ -PGD*. This algorithm is given in Alg. 4<sup>6</sup>.

**Theorem 5** *Let  $\oplus$  denote the Minkowski sum. If  $F$  satisfies RSC/RSS over  $\text{Range}(G) \oplus \text{Span}(B)$  with constants  $\alpha$  and  $\beta$ , and if we assume  $\mu$ -incoherence between  $B$  and  $\text{Range}(G)$ , we have:*

$$F(x_{t+1}) - F(x^*) \leq \left( \frac{2 - \frac{\beta}{\alpha} \frac{1-2.5\mu}{1-\mu}}{1 - \frac{\beta}{2\alpha} \frac{\mu}{1-\mu}} \right) (F(x_t) - F(x^*)) + O(\varepsilon).$$

**Proof:** We will generalize the proof technique of [24]. We first define some auxiliary variables that help us with the proof. Let:

$$w_t = x_t - \eta \nabla F(x_t),$$

$$w_t^u = u_t - \eta \nabla F(x_t),$$

$$w_t^v = v_t - \eta \nabla F(x_t).$$

and let  $x^* = u^* + v^*$  be the minimizer that we seek. As above, by invoking RSS and with some algebra, we obtain:

$$F(x_{t+1}) - F(x^*) \leq \frac{\beta}{2} \left( \|x_{t+1} - w_t\|^2 - \|x_t - w_t\|^2 \right), \quad (17)$$

However, by definition,

$$x_{t+1} = u_{t+1} + v_{t+1},$$

$$x_t = u_t + v_t.$$

Therefore,

$$\begin{aligned} & \|x_{t+1} - w_t\|^2 \\ &= \|u_{t+1} - (u_t - \eta \nabla F(x_t)) + \\ & \quad v_{t+1} - (v_t - \eta \nabla F(x_t)) + \eta \nabla F(x_t)\|^2 \\ &= \|u_{t+1} - (u_t - \eta \nabla F(x_t))\|^2 + \|\eta \nabla F(x_t)\|^2 + \\ & \quad \|v_{t+1} - (v_t - \eta \nabla F(x_t))\|^2 \\ & \quad + 2\langle u_{t+1} - (u_t - \eta \nabla F(x_t)), \eta \nabla F(x_t) \rangle \\ & \quad + 2\langle v_{t+1} - (v_t - \eta \nabla F(x_t)), \eta \nabla F(x_t) \rangle \\ & \quad + 2\langle u_{t+1} - (u_t - \eta \nabla F(x_t)), v_{t+1} - (v_t - \eta \nabla F(x_t)) \rangle. \end{aligned}$$

---

<sup>6</sup>The algorithm is a variant of block-coordinate descent, except that the block updates share the same gradient term.

But  $u_{t+1}$  is an  $\varepsilon$ -projection of  $w_t^u$  and  $u^*$  is in the range of  $G$ , we have:

$$\|u_{t+1} - w_t^u\|^2 \leq \|u^* - w_t^u\|^2 + \varepsilon.$$

Similarly, since  $v_{t+1}$  is an  $l$ -sparse thresholded version of  $w_t^v$ , we have:

$$\|v_{t+1} - w_t^v\|^2 \leq \|v^* - w_t^v\|^2.$$

Plugging in these two upper bounds, we get:

$$\begin{aligned} & \|x_{t+1} - w_t\|^2 \\ & \leq \|u^* - (u_t - \eta \nabla F(x_t))\|^2 + \varepsilon \\ & \quad + \|\eta \nabla F(x_t)\|^2 + \|v^* - (v_t - \eta \nabla F(x_t))\|^2 \\ & \quad + 2\langle u_{t+1} - (u_t - \eta \nabla F(x_t)), \eta \nabla F(x_t) \rangle \\ & \quad + 2\langle v_{t+1} - (v_t - \eta \nabla F(x_t)), \eta \nabla F(x_t) \rangle \\ & \quad + 2\langle u_{t+1} - (u_t - \eta \nabla F(x_t)), v_{t+1} - (v_t - \eta \nabla F(x_t)) \rangle. \end{aligned}$$

Expanding squares and cancelling (several) terms, the right hand side of the above inequality can be simplified to obtain:

$$\begin{aligned} & \|x_{t+1} - w_t\|^2 \\ & \leq \|u^* + v^* - w_t\|^2 + \varepsilon \\ & \quad + 2\langle u_{t+1} - u_t, v_{t+1} - v_t \rangle - 2\langle u^* - u_t, v^* - v_t \rangle \\ & = \|x^* - w_t\|^2 + \varepsilon + 2\langle u_{t+1} - u_t, v_{t+1} - v_t \rangle \\ & \quad - 2\langle u^* - u_t, v^* - v_t \rangle. \end{aligned}$$

Plugging this into (17), we get:

$$\begin{aligned} & F(x_{t+1}) - F(x^*) \\ & \leq \underbrace{\frac{\beta}{2} \left( \|x^* - w_t\|^2 - \|x_t - w_t\|^2 \right)}_{\mathbb{T}_1} \\ & \quad + \underbrace{\beta \left( \langle u_{t+1} - u_t, v_{t+1} - v_t \rangle - 2\langle u^* - u_t, v^* - v_t \rangle \right)}_{\mathbb{T}_2} \\ & \quad + \frac{\beta \varepsilon}{2}. \end{aligned}$$

We already know how to bound the first term  $\mathbb{T}_1$ , using an identical argument as in the proof of Theorem 2. We get:

$$\mathbb{T}_1 \leq \left( 2 - \frac{\beta}{\alpha} \right) (F(x^*) - F(x_t)) + \frac{\beta - \alpha}{\alpha} \gamma \Delta.$$

The second term  $\mathbb{T}_2$  can be bounded as follows. First, observe that

$$\begin{aligned} & |\langle u_{t+1} - u_t, v_{t+1} \rangle| \\ & \leq \mu \|u_{t+1} - u_t\| \|v_{t+1} - v_t\| \\ & \leq \frac{\mu}{2} \left( \|u_{t+1} - u_t\|^2 + \|v_{t+1} - v_t\|^2 \right) \\ & \leq \frac{\mu}{2} \left( \|u_{t+1} + v_{t+1} - u_t - v_t\|^2 \right) \\ & \quad + \mu |\langle u_{t+1} - u_t, v_{t+1} - v_t \rangle|. \end{aligned}$$

This gives us the following inequalities:

$$\begin{aligned} & |\langle u_{t+1} - u_t, v_{t+1} \rangle| \\ & \leq \frac{\mu}{2(1-\mu)} \|x_{t+1} - x_t\|^2 \\ & = \frac{\mu}{2(1-\mu)} \left( \|x_{t+1} - x^*\|^2 + \|x_t - x^*\|^2 + \right. \\ & \quad \left. 2|\langle x_{t+1} - x^*, x_t - x^* \rangle| \right) \\ & \leq \frac{\mu}{1-\mu} \left( \|x_{t+1} - x^*\|^2 + \|x_t - x^*\|^2 \right). \end{aligned}$$




Figure 2: (a) Comparison of our algorithm (Alg. 1) with CSGM [20] and Lasso on MNIST; (b) Reconstruction results with  $m = 100$  measurements; (c) Reconstruction results on celebA dataset with  $m = 1000$  measurements.

Similarly,

$$\begin{aligned}
|\langle u^* - u_t, v^* - v_t \rangle| &\leq \mu \|u^* - u_t\| \|v^* - v_t\| \\
&\leq \frac{\mu}{2} (\|u^* - u_t\|^2 + \|v^* - v_t\|^2) \\
&= \frac{\mu}{2} (\|u^* + v^* - u_t - v_t\|^2) \\
&\quad + \mu |\langle u^* - u_t, v^* - v_t \rangle|,
\end{aligned}$$

which gives:

$$|\langle u^* - u_t, v^* - v_t \rangle| \leq \frac{\mu}{2(1-\mu)} \|x^* - x_t\|^2.$$

Combining, we get:

$$\mathbb{T}_2 \leq \frac{\beta\mu}{2(1-\mu)} (3\|x^* - x_t\|^2 + \|x^* - x_{t+1}\|^2).$$

Moreover, by invoking RSC and Cauchy-Schwartz (similar to the proof of Theorem 2), we have:

$$\begin{aligned}
\|x^* - x_t\|^2 &\leq \frac{1}{\alpha} (F(x_t) - F(x^*)) + O(\varepsilon), \\
\|x^* - x_{t+1}\|^2 &\leq \frac{1}{\alpha} (F(x_{t+1}) - F(x^*)) + O(\varepsilon).
\end{aligned}$$

Therefore we obtain the upper bound on  $\mathbb{T}_2$ :

$$\begin{aligned}
\mathbb{T}_2 &\leq \frac{3\beta\mu}{2\alpha(1-\mu)} (F(x_t) - F(x^*)) \\
&\quad + \frac{\beta\mu}{2\alpha(1-\mu)} (F(x_{t+1}) - F(x^*)) + C'\varepsilon.
\end{aligned}$$

Plugging in the upper bounds on  $\mathbb{T}_1$  and  $\mathbb{T}_2$  and re-arranging terms, we get:

$$\begin{aligned}
&\left(1 - \frac{\beta\mu}{2\alpha(1-\mu)}\right) (F(x_{t+1}) - F(x^*)) \\
&\leq \left(2 - \frac{\beta}{\alpha} + \frac{3\beta\mu}{2\alpha(1-\mu)}\right) (F(x_t) - F(x^*)) + C'\varepsilon,
\end{aligned}$$

which leads to the desired result.  $\square$

## 4 Experimental Results

Our primary focus in this paper is theoretical. However, we supplement our theory with representative numerical experiments that show the promise of our proposed algorithms. In this section, we describe our experimental setup and report algorithm performance.






Figure 3: Comparison of our algorithm (Alg. 2) with Lasso for non-linear forward models (a) with  $\mathcal{A}(x^*) = Ax^* + \sin(Ax^*)$ ; (b) with  $\mathcal{A}(x^*) = \text{sigmoid}(Ax^*)$ . Reconstruction results are on celebA dataset with  $m = 1000$  measurements.

#### 4.1 Experiments with compressed sensing

We use two different GAN architectures and two different datasets in our experiments to show that our approach [Alg. 1] can work with variety of GAN architectures and datasets. We provide comparisons with the CSGM algorithm proposed in [20] as well as the Lasso.

In our experiments, we choose the entries of the matrix  $A$  independently from a Gaussian distribution with zero mean and variance  $1/m$ . In these experiments, we ignore the presence of noise; however, our experiments can be replicated in the presence of additive Gaussian noise. For comparison with CSGM [20], we use a gradient descent optimizer keeping the total number of update steps ( $T \times T_{in}$ ) fixed for both algorithms to produce fair comparisons.


In the first experiment, we use a very simple GAN model trained on the MNIST dataset, which is collection of 60,000 handwritten digit images, each of size  $28 \times 28$  [72]. In our GAN, both the generator and the discriminator are fully-connected neural networks with only one hidden layer. The generator consists of 20 input neurons, 200 hidden-layer neurons and 784 output neurons, while the discriminator consists of 784 input neurons, 128 hidden layer neurons and 1 output neuron. The size of the latent space is set to  $k = 20$ , i.e., the input to our generator is a standard normal vector  $z \in R^{20}$ . We train the GAN using the method described in [12]. We use the Adam optimizer [73] with learning rate 0.001 and mini-batch size 128 for the training.

We test the MNIST GAN with 10 test images taken from the span of the generator to avoid model mismatch issues, and provide both quantitative and qualitative results. For PGD-GAN, because of the zero initialization, a high learning rate is required to get a meaningful output before passing it to the projection step. Therefore, we choose  $\eta = 0.5$ . The parameter  $\eta_{in}$  is set to 0.01 with  $T = 15$  and  $T_{in} = 200$ . Thus, the total number of update steps is fixed to 3000. Similarly, the algorithm of [20] is tested with 3000 updates and  $\eta = 0.01$ . For reporting purposes, we use the reconstruction error  $= \|\hat{x} - x^*\|^2$ . In Fig. 2(a), we show the reconstruction error comparisons for increasing values of number of measurements. It can be seen that our algorithm performs better than the other two methods. Also, as the input images are chosen from the span of the generator itself, it is possible to get close to zero error with only 100 measurements. Fig. 2(b) depicts reconstruction results for selected MNIST images.


The second set of our experiments are performed on a Deep Convolutional GAN (DCGAN) trained on the celebA dataset, which contains more than 200,000 face images of celebrities [74]. We use a pre-trained DCGAN model, which was made available by [20]. The dimension of the latent space for the DCGAN model is  $k = 100$ . We report the results on a held out test dataset, unseen by the GAN at the time of training. The total number of updates is set to 1000, with  $T = 10$  and  $T_{in} = 100$ . Learning rates for PGD-GAN are set as  $\eta = 0.5$  and  $\eta_{in} = 0.1$ . The algorithm of [20] is run with  $\eta = 0.1$  and 1000 update steps. Image reconstruction results from  $m = 1000$  measurements with our algorithm are displayed in Fig. 2(c). We observe that our algorithm produces better reconstructions compared to the other baselines.

#### 4.2 Experiments with nonlinear inverse problems


We extend our experiments to nonlinear models to depict the performance of our algorithm as described in Sec. 3.2. We present image reconstructions from the measurements obtained using two non-linear forward models: a sinusoidal model with  $\mathcal{A}(x^*) = Ax^* + \sin(Ax^*)$ , and a sigmoidal model with  $\mathcal{A}(x^*) = \text{sigmoid}(Ax^*)$ . Similar to the linear case, these experiments are performed using a DCGAN trained on celebA. It is evident that our algorithm produces superior reconstructions as depicted in Fig. 3.



(a) Reconstruction results on MNIST with  $m = 60$  measurements.




(b) Reconstruction error




(c) Mean SSIM


Figure 4: Comparison of Phase-PGD (ours) and DPR [75] on MNIST test set.



(a) Reconstruction results on celebA dataset with  $m = 1000$  measurements.



(b) Reconstruction error



(c) Mean SSIM

Figure 5: Comparison of Phase-PGD (ours) and DPR [75] on celebA test set.

### 4.3 Experiments with phase retrieval

In this section, we describe our experimental setup and report the performance comparisons of our proposed Phase-PGD and deep phase retrieval (DPR) method proposed in [75]. The DPR method estimates the signal as

$$\hat{x} = G \left( \arg \min_z \|y - |AG(z)|\|_2^2 \right), \quad (18)$$

which requires solving an optimization problem directly over the latent space  $z$  of the generative model  $G(\cdot)$ .

In our experiments, we choose the entries of the matrix  $A$  independently from the  $\mathcal{N}(0, \frac{1}{m})$  distribution. Although we ignore the presence of noise, it is possible to replicate our experiments with additive Gaussian noise.

We use two different generative models for the MNIST and CelebA datasets. The generative model for CelebA follows the DCGAN framework [76] except that we do not use any batchnorm layer since the gradient for this layer is dependent on batch size and the distribution of the batch. We train our generators by jointly optimizing generator parameters and the latent code  $z$  using SGD by following the procedure in [77]. We use the squared-loss function,  $\|x - \hat{x}\|^2$  to train the generators. We choose  $z$  from the standard normal distribution on  $\mathbb{R}^k$  and normalize it to unit norm. We project  $z$  back to the unit norm after each gradient update.

In our first set of experiments, we use a generator trained over the MNIST training dataset resized to  $32 \times 32$  pixel. We test two approaches on 10 images from the test set of MNIST dataset and provide both quantitative and qualitative results. For Phase-PGD, we choose gradient descent step size  $\eta = 0.9$  and  $T = 50$ . For fair comparison, we use 2500 iterations for DPR. We show the reconstruction error comparison in Fig. 4b and SSIM comparison in Fig. 4c for increasing number of measurements. Since the input images are not chosen from the span of the generator, the reconstruction error does not necessarily reduce to zero. Nevertheless, we observe in Fig. 4b that Phase-PGD error gets close to zero with nearly 60 measurements, which is significantly smaller than those of DPR method. Fig. 4a depicts reconstruction results for some of the selected MNIST images for the two approaches.

For our second set of experiments, we train a generator for the CelebA dataset. For training, we resize 202,599 color images of celebrity faces in celebA dataset to  $64 \times 64 \times 3$  and kept  $\frac{1}{32}$  of the images for testing. We do not use the aligned and cropped version, which includes only the faces in the images.

We experiment on a subset of 10 images from the test dataset and report reconstruction results. We set the total number of updates to 1500 with  $T = 50$  for Phase-PGD. Image reconstruction results from  $m = 1000$  measurements using Phase-PGD algorithm are presented in Fig. 5a. A comparison between our method and DPR in terms of reconstruction error and SSIM is shown in Fig. 5b and Fig. 5c, respectively. We observe that Phase-PGD can achieve good reconstruction with significantly fewer measurements compared to DPR method.

## 5 DISCUSSION

Our contributions in this paper are primarily theoretical. We also explored the practical benefits of our approach in the context of inverse problems such as compressive sensing and phase retrieval. The algorithms proposed in this paper are generic and can be potentially used to solve other nonlinear inverse problems as well.

We make several assumptions to enable our analysis. Some of them (for example, restricted strong convexity/smoothness; incoherence) are standard analysis tools and are common in the high-dimensional statistics and compressive sensing literature. However, in order to be applicable, they need to be verified for specific problems. A broader characterization of problems that do satisfy these assumptions will be of great interest.

## References

- [1] V. Shah and C. Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Apr. 2018.
- [2] C. Hegde. Algorithmic aspects of inverse problems using generative models. In *Proc. Allerton Conf. on Comm., Contr., and Comp.*, Oct. 2018.
- [3] Rakib Hyder, Viraj Shah, Chinmay Hegde, and M Salman Asif. Alternating phase projected gradient descent with generative priors for solving compressive phase retrieval. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7705–7709. IEEE, 2019.
- [4] E. Candès et al. Compressive sampling. In *Proc. of the intl. congress of math.*, volume 3, pages 1433–1452. Madrid, Spain, 2006.
- [5] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56(4):1982–2001, Apr. 2010.
- [6] C. Hegde, P. Indyk, and L. Schmidt. Fast algorithms for structured sparsity. *Bulletin of the EATCS*, 1(117):197–228, Oct. 2015.
- [7] M. Duarte, C. Hegde, V. Cevher, and R. Baraniuk. Recovery of compressible signals from unions of subspaces. In *Proc. IEEE Conf. Inform. Science and Systems (CISS)*, March 2009.
- [8] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. Image Processing*, 15(12):3736–3745, 2006.
- [9] M. Aharon, M. Elad, and A. Bruckstein. *rmk*-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Processing*, 54(11):4311–4322, 2006.
- [10] T. Chan, J. Shen, and H. Zhou. Total variation wavelet inpainting. *Jour. of Math. imaging and Vision*, 25(1):107–125, 2006.
- [11] S. Ravishanker and Y. Bresler. Learning sparsifying transforms. *IEEE Trans. Signal Processing*, 61(5):1072–1086, 2013.
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 2672–2680, 2014.
- [13] J. Zhu, P. Krähenbühl, E. Shechtman, and A. Efros. Generative visual manipulation on the natural image manifold. In *Proc. European Conf. Comp. Vision (ECCV)*, 2016.
- [14] A. Brock, T. Lim, J. Ritchie, and N. Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016.
- [15] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 2172–2180, 2016.

- [16] J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.
- [17] D. Berthelot, T. Schumm, and L. Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017.
- [18] R. Yeh, C. Chen, T. Lim, M. Hasegawa-Johnson, and M. Do. Semantic image inpainting with perceptual and contextual losses. *arXiv preprint arXiv:1607.07539*, 2016.
- [19] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, pages 105–114, 2017.
- [20] A. Bora, A. Jalal, E. Price, and A. Dimakis. Compressed sensing using generative models. *Proc. Int. Conf. Machine Learning*, 2017.
- [21] E. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Trans. Inform. Theory*, 61(4):1985–2007, 2015.
- [22] L. Demanet and P. Hand. Stable optimizationless recovery from phaseless linear measurements. *Journal of Fourier Analysis and Applications*, 20(1):199–221, 2014.
- [23] C. Hegde and R. Baraniuk. SPIN: Iterative signal recovery on incoherent manifolds. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, July 2012.
- [24] C. Hegde and R. Baraniuk. Signal recovery on incoherent manifolds. *IEEE Trans. Inform. Theory*, 58(12):7204–7214, Dec. 2012.
- [25] M. Dhar, A. Grover, and S. Ermon. Modeling sparse deviations for compressed sensing using generative models. In *Proc. Int. Conf. Machine Learning*, 2018.
- [26] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian designs. *J. Machine Learning Research*, 11(Aug):2241–2259, 2010.
- [27] P. Jain and P. Kar. Non-convex optimization for machine learning. *Foundations and Trends in Machine Learning*, 10(3-4):142–336, 2017.
- [28] M. McCoy and J. Tropp. Sharp recovery bounds for convex demixing, with applications. *Foundations of Comp. Math.*, 14(3):503–567, 2014.
- [29] M. Soltani and C. Hegde. Fast algorithms for demixing signals from nonlinear observations. *IEEE Trans. Sig. Proc.*, 65(16):4209–4222, Aug. 2017.
- [30] D. Donoho. De-noising by soft-thresholding. *IEEE Trans. Inform. Theory*, 41(3):613–627, 1995.
- [31] Z. Xu and J. Sun. Image inpainting by patch propagation using patch sparsity. *IEEE Trans. Image Processing*, 19(5):1153–1165, 2010.
- [32] W. Dong, L. Zhang, G. Shi, and X. Wu. Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization. *IEEE Trans. Image Processing*, 20(7):1838–1857, 2011.
- [33] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [34] K. Kulkarni, S. Lohit, P. Turaga, R. Kerviche, and A. Ashok. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, pages 449–458, 2016.
- [35] A. Mousavi, A. Patel, and R. Baraniuk. A deep learning approach to structured signal recovery. In *Proc. Allerton Conf. Communication, Control, and Computing*, pages 1336–1343, 2015.
- [36] A. Mousavi and R. Baraniuk. Learning to invert: Signal recovery via deep convolutional networks. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, 2017.
- [37] L. Xu, J. Ren, C. Liu, and J. Jia. Deep convolutional neural network for image deconvolution. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 1790–1798, 2014.
- [38] C. Dong, C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Machine Intell.*, 38(2):295–307, 2016.
- [39] J. Kim, J. Kwon Lee, and K. Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, pages 1646–1654, 2016.
- [40] R. Yeh, C. Chen, T.-Y. Lim, A. Schwing, M. Hasegawa-Johnson, and M. Do. Semantic image inpainting with deep generative models. In *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, volume 2, page 4, 2017.

- [41] Christopher A Metzler, Felix Heide, Prasana Rangarajan, Muralidhar Madabhushi Balaji, Aparna Viswanath, Ashok Veeraraghavan, and Richard G Baraniuk. Deep-inverse correlography: towards real-time high-resolution non-line-of-sight imaging. *Optica*, 7(1):63–71, 2020.
- [42] J. Rick Chang, C. Li, B. Poczos, B. Vijaya Kumar, and A. Sankaranarayanan. One network to solve them all—solving linear inverse problems using deep projection models. In *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, pages 5888–5897, 2017.
- [43] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [44] J.-Y. Zhu, T. Park, P. Isola, and A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. IEEE Conf. Comp. Vision and Pattern Recog. (CVPR)*, 2017.
- [45] A. Brock, J. Donahue, and K. Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [46] D. Van Veen, A. Jalal, E. Price, S. Vishwanath, and A. Dimakis. Compressed sensing with deep image prior and learned regularization. *arXiv preprint arXiv:1806.06438*, 2018.
- [47] Paul Hand and Vladislav Voroninski. Global guarantees for enforcing deep generative priors by empirical risk. *IEEE Transactions on Information Theory*, 66(1):401–418, 2019.
- [48] G Jagatap and C Hegde. Algorithmic guarantees for inverse imaging with untrained network priors. *Advances in neural information processing systems*, 2019.
- [49] A. Raj, Y. Li, and Y. Bresler. Gan-based projector for faster recovery in compressed sensing with convergence guarantees. *arXiv preprint arXiv:1902.09698*, 2019.
- [50] R.W. Gerchberg and A Saxton W. O. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–250, 11 1971.
- [51] James R Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.
- [52] Emmanuel J Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [53] Jason Holloway, M. Salman Asif, Manoj Kumar Sharma, Nathan Matsuda, Roarke Horstmeyer, Oliver Cossairt, and Ashok Veeraraghavan. Toward long-distance subdiffraction imaging using coherent camera arrays. *IEEE Transactions on Computational Imaging*, 2(3):251–265, 2016.
- [54] Lei Tian, Xiao Li, Kannan Ramchandran, and Laura Waller. Multiplexed coded illumination for fourier ptychography with an led array microscope. *Biomedical optics express*, 5(7):2376–2389, 2014.
- [55] Jianwei Miao, Pambos Charalambous, Janos Kirz, and David Sayre. Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*, 400(6742):342–344, 1999.
- [56] S. Bahmani and J. Romberg. Efficient compressive phase retrieval with constrained sensing vectors. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 523–531, 2015.
- [57] Tom Goldstein and Christoph Studer. Phasemax: Convex phase retrieval via basis pursuit. *IEEE Transactions on Information Theory*, 64(4):2675–2689, 2018.
- [58] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 2796–2804, 2013.
- [59] G. Wang, G. Giannakis, Y. Saad, and J. Chen. Solving most systems of random quadratic equations. In *Proc. Adv. in Neural Processing Systems (NIPS)*, 2017.
- [60] Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Applied and Computational Harmonic Analysis*, 39(2):277–299, 2015.
- [61] G. Jagatap and C. Hegde. Fast, sample efficient algorithms for structured phase retrieval. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 4924–4934, 2017.
- [62] G. Jagatap and C. Hegde. Sample-efficient algorithms for recovering structured signals from magnitude-only measurements. *IEEE Trans. Inform. Theory*, 65(7):4435–4456, July 2019.
- [63] George Barbastathis, Aydogan Ozcan, and Guohai Situ. On the use of deep learning for computational imaging. *Optica*, 6(8):921–943, 2019.
- [64] Ayan Sinha, Justin Lee, Shuai Li, and George Barbastathis. Lensless computational imaging through deep learning. *Optica*, 4(9):1117–1125, 2017.

- [65] P. Shah and V. Chandrasekaran. Iterative projections for signal identification on manifolds: Global recovery guarantees. In *Proc. Allerton Conf. Communication, Control, and Computing*, pages 760–767, 2011.
- [66] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Springer, 2013.
- [67] Qi Lei, Ajil Jalal, Inderjit S Dhillon, and Alexandros G Dimakis. Inverting deep generative models, one layer at a time. *Advances in Neural Information Processing Systems*, 32:13910–13919, 2019.
- [68] S. Negahban, P. Ravikumar, M. Wainwright, B. Yu, et al. A unified framework for high-dimensional analysis of  $m$ -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [69] M. Soltani and C. Hegde. Fast algorithms for demixing sparse signals from nonlinear observations. *IEEE Transactions on Signal Processing*, 65(16):4209–4222, 2017.
- [70] Gauri Jagatap and Chinmay Hegde. Fast, sample-efficient algorithms for structured phase retrieval. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4924–4934, 2017.
- [71] Gauri Jagatap and Chinmay Hegde. Sample-efficient algorithms for recovering structured signals from magnitude-only measurements. *IEEE Transactions on Information Theory*, 65(7):4434–4456, 2019.
- [72] Y. LeCun, L. on Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. of the IEEE*, 86(11):2278–2324, 1998.
- [73] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [74] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proc. of the IEEE Intl. Conf. on Comp. Vision*, pages 3730–3738, 2015.
- [75] Paul Hand, Oscar Leong, and Vladislav Voroninski. Phase retrieval under a generative prior. In *Proc. Adv. in Neural Processing Systems (NIPS)*, pages 9154–9164, 2018.
- [76] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *Proc. Int. Conf. Learning Representations (ICLR)*, 2016.
- [77] Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the latent space of generative networks. In *International Conference on Machine Learning*, pages 600–609. PMLR, 2018.