

# Optimally Sample-Efficient Phase Retrieval with Deep Generative Models

Paul Hand\*, Oscar Leong†, and Vladislav Voroninski‡

\*Department of Mathematics & Computer Science, Northeastern University

p.hand@northeastern.edu

†Department of Computational & Applied Mathematics, Rice University

oscar.f.leong@rice.edu

‡Helm.ai

vlad@helm.ai

**Abstract**—We consider the phase retrieval problem, which asks to recover a structured  $n$ -dimensional signal from  $m$  quadratic measurements. In many imaging contexts, it is beneficial to enforce a sparsity prior on the signal to reduce the number of measurements necessary for recovery. However, the best known methodologies for sparse phase retrieval have a sub-optimal quadratic dependency on the sparsity level of the signal at hand. In this work, we instead model signals as living in the range of a deep generative neural network  $G : \mathbb{R}^k \rightarrow \mathbb{R}^n$ . We show that under the model of a  $d$ -layer feed forward neural network with Gaussian weights,  $m = O(kd^2 \log n)$  generic measurements suffice for the  $\ell_2$  empirical risk minimization problem to have favorable geometry. In particular, we exhibit a descent direction for all points outside of two arbitrarily small neighborhoods of the true  $k$ -dimensional latent code and a negative reflection of it. Our proof is based on showing the sufficiency of two deterministic conditions on the generator and measurement matrices, which are satisfied with high probability under random Gaussian ensembles. We corroborate these results with numerical experiments showing that enforcing a generative prior via empirical risk minimization outperforms sparse phase retrieval methods.

## I. INTRODUCTION

We study the phase retrieval problem which asks to recover a structured signal  $y_0 \in \mathbb{R}^n$  from  $m$  non-linear measurements  $b = |Ay_0|$  where  $A \in \mathbb{R}^{m \times n}$  is known,  $|\cdot|$  acts entrywise, and  $m$  is to be minimized. The phase retrieval problem arises in a number of different contexts, with its main area of importance being X-ray crystallography [7]. In practice, the missing phase or sign information of the linear measurements is lost during the measurement process.

Researchers have considered exploiting structural priors on the signal with the aim of lowering the sample complexity necessary for recovery. In many imaging contexts, a natural prior to enforce is that the signal is sparse since images are compressible, or nearly sparse, with respect to a wavelet basis. However, sparsity-based methods for phase retrieval have been met with severe computational bottlenecks as the best known methods to reconstruct an  $s$ -sparse  $n$ -dimensional signal require  $O(s^2 \log n)$  generic measurements. Non-convex optimization methods, for example, require careful initialization to succeed, the best of which require  $O(s^2 \log n)$  measurements [3]. Moreover, it has been shown that the semidefinite

programming approach, PhaseLift, provably cannot surpass the  $O(s^2 \log n)$  bottleneck via direct  $\ell_1$  penalization [10]. Hence in order to overcome such a complexity barrier, improved initialization schemes must be made or new forms of regularization must be studied. Here we take the latter approach through the utilization of generative models borrowed from deep-learning.

Generative models have become profound tools to approximate distributions of various complex signal classes. Generative Adversarial Networks (GANs), for example, have demonstrated the impressive ability to create photorealistic, yet synthetic human faces [8] and have pushed forward the state-of-the-art in image generation. Moreover, they are able to express such complex classes in a *low-dimensional* way and provide a direct parametrization of the natural signal manifold. Based on this critical attribute, they exhibit the potential to regularize ill-posed inverse problems for a variety of natural signal classes, a property we aim to explore in this work in the context of phase retrieval.

Previous works have also considered using generative models as priors in a closely related imaging inverse problem, compressed sensing. [2] showed that, empirically, 5-10X fewer measurements were necessary for recovery compared to standard sparsity-based techniques. Furthermore, [6] showed that the optimization landscape of the  $\ell_2$  empirical risk minimization problem over the latent code space of the generator exhibits favorable geometry for gradient methods under the model of a random feed forward network. Based on these positive results for enforcing generative priors in compressed sensing along with the potential theoretical and practical limitations of sparse phase retrieval, we consider enforcing a generative prior as a form of regularization. The present manuscript is a shortened version of that which appears in [5] with more extensive experimentation which will be discussed in Section IV.

## II. FAVORABLE GEOMETRY OF EMPIRICAL RISK MINIMIZATION UNDER A GENERATIVE PRIOR

In this work, we suppose the signal of interest is the output of a  $d$ -layer feed forward neural network with Rectifying Linear Unit (ReLU) activation functions and no bias terms.

Specifically, we assume  $y_0 = G(x_0)$  where  $G : \mathbb{R}^k \rightarrow \mathbb{R}^{n_d}$  is defined by

$$G(x) := \text{relu}(W_d \dots \text{relu}(W_2 \text{relu}(W_1 x)) \dots)$$

where each  $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$  denote the weights in the  $i$ -th layer of our network with  $k := n_0 < n_1 < \dots < n_d =: n$ . Given measurements of the form  $|Ay_0|$ , we aim to recover the latent code  $x_0$  by solving the following  $\ell_2$  empirical risk minimization problem:

$$\min_{x \in \mathbb{R}^k} f(x) := \frac{1}{2} \left\| |AG(x)| - |Ay_0| \right\|^2 \quad (\text{II.1})$$

where  $\|\cdot\|$  is the  $\ell_2$  norm. Due to the non-linearities introduced by the measurements and generative model, this is a non-convex optimization problem. Theoretically, spurious local minima may exist, prohibiting gradient descent with random initialization from succeeding. However, our main result asserts that with high probability, the objective function has favorable geometry for gradient methods with information theoretically optimal sample complexity with respect to the latent code dimension. To formally state our result, we define some notation. Note that while the objective function is not differentiable, its one-sided directional derivative exists everywhere, i.e. for any  $x \in \mathbb{R}^k$  and direction  $v$ ,  $D_v f(x) := \lim_{t \rightarrow 0^+} \frac{f(x+tv) - f(x)}{t}$  exists.

**Theorem II.1.** Fix  $\epsilon > 0$  such that  $K_1 d^8 \epsilon^{1/4} \leq 1$  and let  $d \geq 2$ . Suppose  $G$  is such that  $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$  has i.i.d.  $\mathcal{N}(0, 1/n_i)$  entries for  $i = 1, \dots, d$ . Suppose that  $A \in \mathbb{R}^{m \times n_d}$  has i.i.d.  $\mathcal{N}(0, 1/m)$  entries independent from  $\{W_i\}$ . Then if  $m \geq C_\epsilon dk \log(n_1 n_2 \dots n_d)$  and  $n_i \geq C_\epsilon n_{i-1} \log n_{i-1}$  for  $i = 1, \dots, d$ , then with probability at least  $1 - \sum_{i=1}^d \gamma n_i e^{-c_\epsilon n_{i-1}} - \gamma m^{4k+1} e^{-c_\epsilon m}$ , the following holds: for all non-zero  $x, x_0 \in \mathbb{R}^k$ , there exists  $v_{x, x_0} \in \mathbb{R}^k$  such that the one-sided directional derivatives of  $f$  satisfy

$$\begin{aligned} D_{-v_{x, x_0}} f(x) &< 0, \quad \forall x \notin \mathcal{B}(x_0, K_2 d^3 \epsilon^{1/4} \|x_0\|) \\ &\cup \mathcal{B}(-\rho_d x_0, K_2 d^{1/4} \epsilon^{1/4} \|x_0\|) \cup \{0\}, \\ D_x f(0) &< 0, \quad \forall x \neq 0, \end{aligned}$$

where  $\rho_d > 0$  converges to 1 as  $d \rightarrow \infty$  and  $K_1$  and  $K_2$  are universal constants. Here  $C_\epsilon$  depends polynomially on  $\epsilon^{-1}$ ,  $c_\epsilon$  depends on  $\epsilon$ , and  $\gamma$  is a universal constant.

The theorem essentially states that, with a sufficient number of Gaussian measurements and a sufficiently expansive ReLU network with Gaussian weights, the constructed descent direction  $v_{x, x_0}$  does not vanish outside of two small neighborhoods of the minimizer and a negative multiple of it. This eliminates the existence of spurious local minima away from these two neighborhoods. We first show the sufficiency of two deterministic conditions on the measurement matrix and weights of our generative model for this result to hold. We then prove that measurement matrices and weights with i.i.d. Gaussian entries satisfy these conditions with high probability.

The first condition concerns the spatial arrangement and number of neurons in our network. We first define the following: for a matrix  $W \in \mathbb{R}^{n \times k}$  and  $x \in \mathbb{R}^k$ , let  $W_{+,x} :=$

$\text{diag}(Wx > 0)W$ . This quantity can be viewed as determining which neurons are active in the neural network when the input to it is  $x$ .

**Definition II.2.** We say that  $W$  satisfies the Weight Distribution Condition (WDC) with constant  $\epsilon > 0$  if for all  $x, y \in \mathbb{R}^k$ ,

$$\|W_{+,x}^\top W_{+,y} - Q_{x,y}\| \leq \epsilon$$

where

$$Q_{x,y} := \frac{\pi - \theta_{x,y}}{2\pi} I_k + \frac{\sin \theta_{x,y}}{2\pi} M_{\hat{x} \leftrightarrow \hat{y}}.$$

Here  $\theta_{x,y} = \angle(x, y)$ ,  $\hat{x} = x/\|x\|$ ,  $\hat{y} = y/\|y\|$ ,  $I_k$  is the  $k \times k$  identity matrix, and  $M_{\hat{x} \leftrightarrow \hat{y}}$  is the matrix that sends  $\hat{x} \mapsto \hat{y}$ ,  $\hat{y} \mapsto \hat{x}$ , and  $z \mapsto 0$  for any  $z \in \text{span}(\{\hat{x}, \hat{y}\})^\perp$ .

Note that if  $W_{ij} \sim \mathcal{N}(0, 1/n)$ , then  $\mathbb{E}[W_{+,x}^\top W_{+,y}] = Q_{x,y}$  for any  $x, y \in \mathbb{R}^k$ . In [6], it was shown that Gaussian  $W$  satisfies the WDC with high probability:

**Lemma II.3 (WDC).** Fix  $0 < \epsilon < 1$  and suppose  $W \in \mathbb{R}^{n \times k}$  has i.i.d.  $\mathcal{N}(0, 1/n)$  entries. Then if  $n \geq C_\epsilon k \log k$ , then with probability at least  $1 - 8n \exp(-\gamma_\epsilon k)$ ,  $W$  satisfies the WDC with constant  $\epsilon$ . Here  $C_\epsilon$  and  $\gamma_\epsilon^{-1}$  depend polynomially on  $\epsilon^{-1}$ .

The assumption that the weights of our generative neural network follow a Gaussian distribution is supported by recent results showcasing real neural networks, post-training, exhibiting statistics similar to Gaussians [1]. In any case, we note that the conditions we impose on the generator and measurement matrix are deterministic, meaning they could potentially be satisfied with other distributions.

The next condition quantifies whether the measurement matrix behaves like a Gaussian when acting on the range of our generative model. For  $A \in \mathbb{R}^{m \times n}$  and  $z \in \mathbb{R}^n$ , define  $A_z := \text{diag}(\text{sgn}(Az))A$  where  $\text{sgn} : \mathbb{R} \rightarrow \mathbb{R}$  acts entrywise and is defined by  $\text{sgn}(a) = a/|a|$  for  $a \neq 0$  and 0 otherwise. Note that for any  $z \in \mathbb{R}^n$ ,  $|Az| = A_z z$ .

**Definition II.4.** We say that  $A$  satisfies the Range Restricted Concentration Property (RRCP) with constant  $\epsilon > 0$  if for any  $x, y \in \mathbb{R}^k$ , the matrices  $A_{G(x)}$  and  $A_{G(y)}$  satisfy the following for all  $x_1, x_2, x_3, x_4 \in \mathbb{R}^n$ :

$$\begin{aligned} &| \langle (A_{G(x)}^\top A_{G(y)} - \Phi_{G(x), G(y)})(G(x_1) - G(x_2)), G(x_3) - G(x_4) \rangle | \\ &\leq 31\epsilon \|G(x_1) - G(x_2)\| \|G(x_3) - G(x_4)\| \end{aligned}$$

where

$$\Phi_{z,w} := \frac{\pi - 2\theta_{z,w}}{\pi} I_n + \frac{2 \sin \theta_{z,w}}{\pi} M_{\hat{z} \leftrightarrow \hat{w}}.$$

Note that if  $A_{ij} \sim \mathcal{N}(0, 1/m)$ , then for any  $z, w \in \mathbb{R}^n$ ,  $\mathbb{E}[A_z^\top A_w] = \Phi_{z,w}$ . In our work, we establish that Gaussian  $A$  satisfies the RRCP with high probability:

**Lemma II.5.** Fix  $0 < \epsilon < 1$  and suppose  $A \in \mathbb{R}^{m \times n}$  has i.i.d.  $\mathcal{N}(0, 1/m)$  entries. If  $m \geq C_\epsilon kd \log(n_1 n_2 \dots n_d)$ , then with probability  $1 - \gamma m^{4k+1} \exp(-c_\epsilon m)$ ,  $A$  satisfies the RRCP with constant  $\epsilon$ . Here  $\gamma$  is a universal constant and  $C_\epsilon$  and  $c_\epsilon^{-1}$  depend polynomially on  $\epsilon^{-1}$ .

Proving this particular property for a Gaussian measurement matrix  $A$  requires establishing a generalized Restricted Isometry Property for  $A_z^\top A_w$  uniformly over all elements  $z, w \in \text{range}(G)$ . Due to the ReLU non-linearities, the generative model  $G$  is piecewise linear. Thus its range lies in the union of finitely many subspaces. Hence one must control  $A_{G(x)}^\top A_{G(y)}$  over these finitely many subspaces for all  $x, y \in \mathbb{R}^k$ . This quantity that can be interpreted as a spatially dependent sensing matrix.

Our main deterministic result shows that these two conditions on the generator and measurement matrix are sufficient to establish favorable geometry of the empirical risk minimization problem:

**Theorem II.6.** Fix  $\epsilon > 0$  such that  $K_1 d^8 \epsilon^{1/4} \leq 1$  and let  $d \geq 2$ . Suppose that  $G$  is such that  $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$  satisfies the WDC with constant  $\epsilon$  for all  $i = 1, \dots, d$ . Suppose  $A \in \mathbb{R}^{m \times n_d}$  satisfies the RRCP with constant  $\epsilon$ . Then the same conclusion as Theorem II.1 holds.

### III. A VARIATION OF GRADIENT DESCENT

Based on our geometric result concerning the energy landscape of the empirical risk minimization problem, gradient descent would converge either to a neighborhood of the true solution or a negative multiple of it. In order to avoid the latter neighborhood, we propose a variation of gradient descent to converge to the true solution. To aid in our exposition of the algorithm, consider Figure 1 which showcases the objective function *in expectation* which can be interpreted as the case when the number of measurements  $m \rightarrow \infty$  and the network layers become arbitrarily expansive.

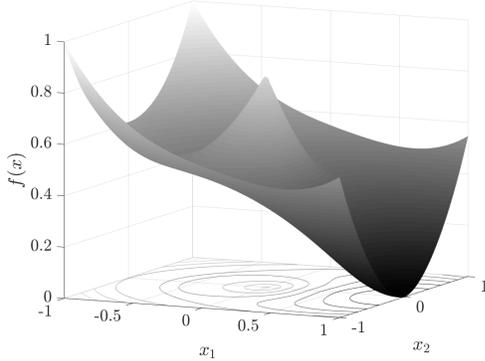


Fig. 1. Surface plot of (II.1) with  $x_0 = [1 \ 0]^\top \in \mathbb{R}^2$  and  $m \rightarrow \infty$ .

We see that there are three critical points: the global optimum  $x_0$ , a negative multiple of the true solution, and the origin which is a local maximum. Moreover, the function value near the negative multiple is higher than near the true solution. While gradient descent may, in principle, be attracted to the negative multiple, we can exploit this property of the objective function by checking the objective function value of our current iterate and its negation, and then choosing the result with lower objective function value as our main iterate.

This would allow us to escape the neighborhood of the critical point and move towards the true minimizer.

To formally describe our algorithm, we define some further notation. Let  $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$  denote the neurons in the  $i$ -th layer of our network. Then, for any  $x \in \mathbb{R}^k$  and  $i \in [d]$ , define

$$W_{i,+x} := \text{diag}(W_{i-1,+x} \dots W_{2,+x} W_{1,+x} x > 0) W_i.$$

Similarly to the definition presented in the WDC, this quantity determines which neurons are active in the  $i$ -th layer of our network when the input to the generative model is  $x$ . Finally, let  $\Pi_{i=d}^1 W_i = W_d \dots W_2 W_1$ . To initialize the algorithm, choose a random initial iterate  $x_1 \neq 0$ . Then, for each  $i = 1, 2, \dots$ , compute the descent direction

$$v_{x_i, x_0} := (\Pi_{i=d}^1 W_{i,+x_i})^\top A_{G(x_i)}^\top (|AG(x_i)| - |Ay_0|).$$

This is the gradient of our objective function when it is differentiable. It is important to note that, even at points of non-differentiability, this descent direction is well-defined due to the definitions of  $A_{G(x)}$  and each  $W_{i,+x}$ . Once computed, we can then take a step in the direction of the negative gradient  $-v_{x_i, x_0}$  as in standard gradient descent. However, prior to taking this step, we check the objective function value for our current iterate and its negation. If the objective function value is lower for the negation, we switch to that point, compute the descent direction with respect to it, and update the iterate. The intuition is that, if gradient descent begins to converge toward the negative multiple, this negation step will choose a point of lower objective function value, which will be closer to a neighborhood of the minimizer. Algorithm 1 formally outlines this process.

---

#### Algorithm 1 Deep Phase Retrieval (DPR) Gradient Method

---

**Require:** Weights  $W_i$ , measurement matrix  $A$ , observations  $|Ay_0|$ , and step size  $\alpha > 0$

- 1: Choose an arbitrary initial point  $x_1 \in \mathbb{R}^k \setminus \{0\}$
  - 2: **for**  $i = 1, 2, \dots$  **do**
  - 3:   **if**  $f(-x_i) < f(x_i)$  **then**
  - 4:      $x_i \leftarrow -x_i$ ;
  - 5:   **end if**
  - 6:    $v_{x_i, x_0} = (\Pi_{i=d}^1 W_{i,+x_i})^\top A_{G(x_i)}^\top (|AG(x_i)| - |Ay_0|)$ ;
  - 7:    $x_{i+1} = x_i - \alpha v_{x_i, x_0}$ ;
  - 8: **end for**
- 

### IV. IMAGE EXPERIMENTS ON MNIST

We now consider the task of recovering a natural image  $y_0 \in \mathbb{R}^n$  from phaseless linear measurements  $|Ay_0|$  where  $A \in \mathbb{R}^{m \times n}$  has i.i.d.  $\mathcal{N}(0, 1)$  entries. The images are taken from the test set of the MNIST Handwritten digit database. We will be comparing our results to three sparse phase retrieval methods: SPARTA [11], CoPRAM [4], and Thresholded Wirtinger Flow (TWF) [3]. We will refer to our algorithm as DPR throughout.

As our generative model, we used the decoder network of a pretrained Variational Autoencoder (VAE) from [2]. The decoder network is of size 20 – 500 – 500 – 784. To recover the image  $y_0$  from measurements  $|Ay_0|$ , we find the closest

image in the range of our generative model  $G$  that best fits our measurements in an  $\ell_2$  sense. In these experiments, we use a variation of Algorithm 1 to solve Equation (II.1). We found that, in practice, the negation step would only occur at the beginning of the algorithm. Hence instead of checking the objective function value of the negation at each step, we choose the best solution  $G(\hat{x})$  from two final outputs of gradient descent: one initialized with a random point  $x_1$  and the other with its negation  $-x_1$ . As our gradient descent scheme, we used the Adam optimizer [9].

For the sparse phase retrieval methods, we performed sparse recovery by transforming the images using the Daubechies-4 Wavelet Transform. More specifically, we considered the vector of wavelet coefficients  $v_0 = \Psi y_0$  which is compressible. Since the wavelet transform  $\Psi$  is orthogonal and the measurement matrix  $A$  is Gaussian,  $A\Psi^\top$  and  $A$  are equal in distribution due to the rotational invariance of  $A$ . Hence, instead of recovering  $v_0$  from measurements  $|Ay_0| = |A\Psi^\top v_0|$ , we recover  $v_0$  from  $|Av_0|$ . We then take an inverse wavelet transform to retrieve the approximate image. To appropriately take the wavelet transform, we pad the images with zeros uniformly around the border so that they are of size  $32 \times 32$ . If a sparse recovery algorithm required a sparsity parameter, we ran the algorithm with a range of sparsity parameter values, choosing the best reconstruction in terms of lowest reconstruction error. The resulting images generated by our algorithm were also uniformly padded with zeros around the border to obtain  $32 \times 32$  images.

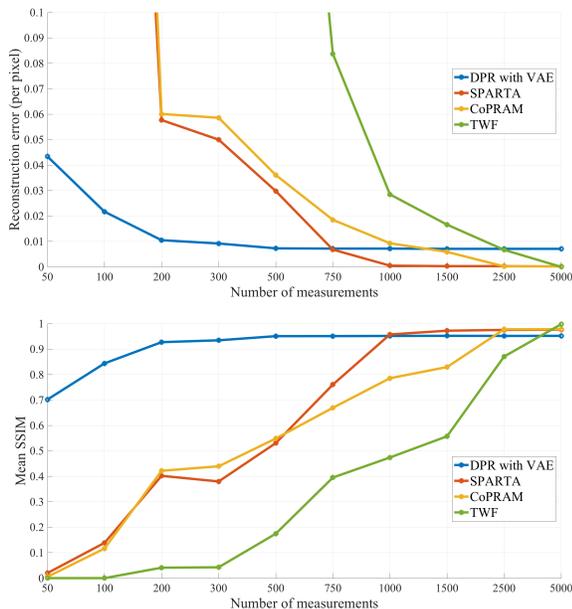


Fig. 2. Each algorithm’s average reconstruction error (top) and mean SSIM (bottom) over 10 images from the MNIST test set for different numbers of measurements.

We attempted to reconstruct 10 images from the MNIST test set. We allowed 5 random restarts for each algorithm and recorded the result with the least  $\ell_2$  reconstruction error per

pixel. In addition, we calculated the Structural Similarity Index Measure (SSIM) [12] for each reconstruction and computed the average for various numbers of measurements. The results in Figure 2 demonstrate the success of our algorithm with very few measurements. For 200 measurements, we can achieve accurate recovery with a mean SSIM value of over 0.9 while other algorithms require 1000 measurements or more. In terms of reconstruction error, our algorithm exhibits recovery with 200 measurements comparable to the alternatives requiring 750 measurements or more, which is where they begin to succeed. We finally present a comparison of qualitative results for low numbers of measurements in Figure 3.



Fig. 3. Each algorithm’s reconstructed images with 100 measurements (left) and 300 measurements (right). If an image is blank, then the reconstruction error between the blank image and the original image was lower than that of the algorithm’s reconstructed image and the original image. We note that even for as few as 100 measurements, nearly all of DPR’s reconstructions are semantically correct.

#### ACKNOWLEDGEMENTS

OL acknowledges support by the NSF Graduate Research Fellowship under Grant No. DGE-1450681. PH acknowledges funding by the grant NSF DMS-1464525.

#### REFERENCES

- [1] S. ARORA, Y. LIANG, & T. MA *Why are deep nets reversible: A simple theory with implications for training.* (2015) CoRR, abs/1511.05653
- [2] A. BORA, A. G. DIMAKIS, A. JALAL, & E. PRICE *Compressed sensing using generative models.* (2017) arXiv preprint arXiv:1703.03208
- [3] T. CAI, X. LI, & Z. MA *Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow.* (2016) The Annals of Statistics, 44(5):2221–2251
- [4] C. HEDGE & G. JAGATAP *Sample-efficient algorithms for recovering structured signals from magnitude-only measurements.* (2017) Neural Information Processing Systems
- [5] P. HAND, O. LEONG, & V. VORONINSKI *Phase Retrieval Under a Generative Prior.* (2018) Neural Information Processing Systems
- [6] P. HAND & V. VORONINSKI *Global guarantees for enforcing deep generative priors by empirical risk.* (2018) arXiv preprint arXiv:1705.07576
- [7] R. HARRISON *Phase problem in crystallography.* (1993) J. Opt. Soc. Am. A, 10(5):1046–1055
- [8] T. KARRAS, S. LAINE, & T. AILA *A style-based generator architecture for generative adversarial networks.* (2018) arXiv preprint arXiv:1812.04948
- [9] D. KINGMA & J. BA *Adam: A method for stochastic optimization.* (2014) arXiv preprint, arXiv:1412.6980
- [10] X. LI & V. VORONINSKI *Sparse signal recovery from quadratic measurements via convex programming.* (2013) SIAM Journal on Mathematical Analysis, 45(5):3019–3033
- [11] G. WANG, L. ZHANG, G. GIANNAKIS, M. AKÇAKAYA, & J. CHEN *Sparse phase retrieval via truncated amplitude flow.* (2018) IEEE Transactions on Signal Processing, 66:479–491
- [12] Z. WANG, A. BOVIK, H. SHEIKH, & E. SIMONCELLI *Image quality assessment: from error visibility to structural similarity.* (2004) IEEE Transactions on Image Processing, 13(4):600–612