Alternating Phase Projected Gradient Descent with Generative Prior for Solving Compressive Phase Retrieval

Rakib Hyder [1] Viraj Shah [2] Chinmay Hegde [2] M. Salman Asif [1]

[1]University of California Riverside [2]Iowa State University

Introduction

- **Phase retrieval** is broadly defined as a problem that deals with the recovery of a real- or complex-valued signal from its amplitude observations. It naturally arises in optical imaging where sensors measure intensity.
- **General problem formulation:** Suppose we are given noisy, amplitude measurements as

$$y = |Ax^*| + e,$$

  - $x^* \in \mathbb{R}^n$ is the unknown signal or image
  - $A \in \mathbb{R}^{m\times n}$ is the measurement operator matrix
  - $y \in \mathbb{R}^m$ is the measurement vector
  - $e \in \mathbb{R}^m$ is the measurement noise.
- **Aim:** To recover unknown signal x^* given y and A.

Phase Retrieval as an Inverse Problem

- In general, phase retrieval is an under-determined problem with infinitely many possible solutions.
- A standard approach for solving such a problem is to restrict the solution space to a set $\mathcal{M} \subset \mathbb{R}^n$ that captures some known structure x^* is expected to obey.
- The phase retrieval problem can then be formulated as the following constrained optimization problem

$$\min_x \ \text{loss}(y; |Ax|) \quad \text{s.t.} \quad x \in \mathcal{M},$$

where loss$(\cdot)$ denotes some appropriate loss function between the given and estimated observations and $\mathcal{M}$ denotes the constraint set.

Generative Prior for Phase Retrieval

- **Traditional constraints** for signals and images include known support for nonzero entries, positivity, and sparsity in some basis.
- **Generative prior** can learn the structure of "natural" images from large training data using an autoencoder or a generative adversarial network (GAN).
- **Learn a generative model** $G(z)$ that maps a latent vector $z \in \mathbb{R}^k$ to a natural image $x \in \mathbb{R}^n$. See an example of a generator in Figure 1. The learned generative model, $G(z)$, is expected to approximate the high-dimensional probability distribution of the image set $\mathcal{M}$. That is,

$$\mathcal{M} = \{x \in \mathbb{R}^n | x = G(z) \text{ for some } z \in \mathbb{R}^k\}.$$

- **Solve an inverse problem with generative priors** as
  - Gradient descent (GD) [3,5]: $\min_z \ \|y - |AG(z)|\|_2^2$
  - Alternating phase gradient descent (APGD): $\min_z \ \|p \odot y - AG(z)\|_2^2$
  - Alternating phase projected gradient descent (APPGD): $\min_{x,z} \ \|p \odot y - Ax\|_2^2 \quad \text{s.t.} \quad x = G(z).$
- Our proposed APPGD improves upon [3,5] by combining the gradient descent and projected gradient descent methods for generative priors [2] with AltMin-based non-convex optimization techniques used in sparse phase retrieval [4,6].

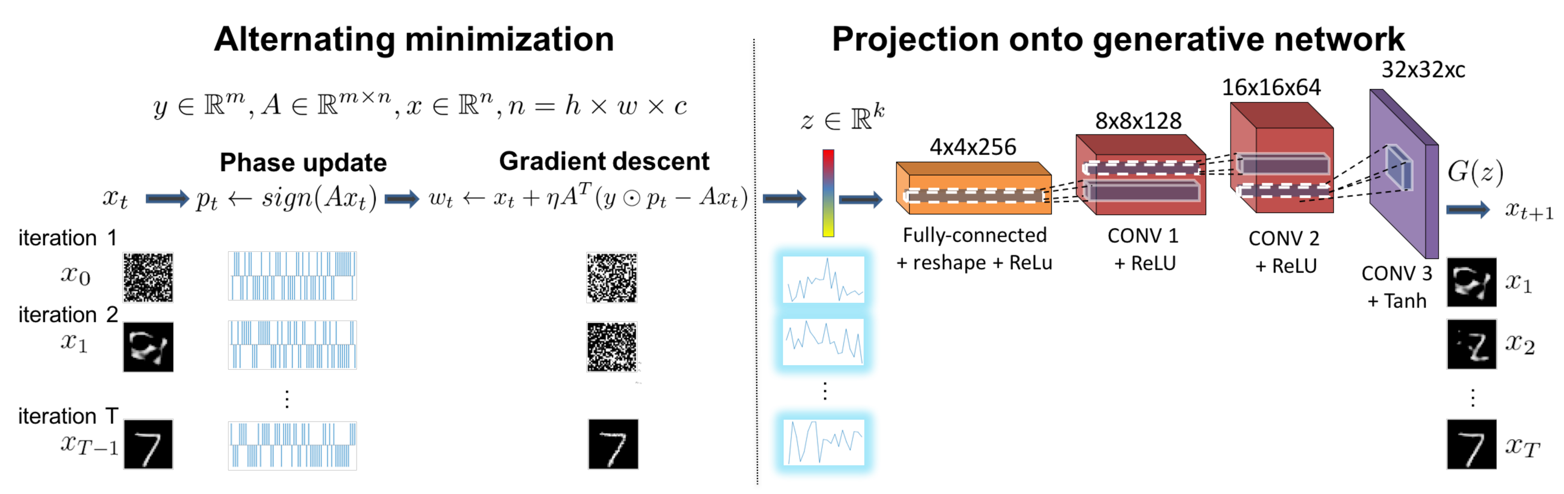

Figure 1. *Illustration of Alternating Phase Projected Gradient Descent (APPGD) algorithm.*

Alternating Phase Projected Gradient Descent

- Our algorithm performs three main tasks at every iteration:
  - Phase update
  - Gradient descent update
  - Projection onto the generator
- We initialize z as a random vector and repeat the three steps until convergence.

Algorithm 1 APPGD

```
Inputs: y, A, G, T, Output: x̂
Choose an initial point x_0 ∈ R^n
for t = 1,...T do
    p_{t-1} ← sign(A x_{t-1})
    w_{t-1} ← x_{t-1} + η A^T (y ⊙ p_{t-1} − A x_{t-1})
    x_t ← P_G(w_{t-1}) = G(arg min_z ||w_{t-1} − G(z)||)
end for
x̂ ← x_T
```

Convergence Result

Theorem: Suppose we have an initialization $x_0 \in$ satisfying dist $(x_0, x^*) \le \delta_0$, 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 1 satisfy

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

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

Experiments

- **Datasets:** MNIST digits and celebA images.
- **Generator:** DCGAN architecture as shown in Figure 1.
- **Latent code dimension:** $k = 32$ for MNIST and $k = 256$ for celebA.
- **Optimizer:** Stochastic gradient descent (SGD) as in [1].
- **Measurement matrix** A is created by selecting entries independently from Gaussian distribution with zero mean and variance $1/m$.

References

[1] P. Bojanowski, A. Joulin, D. Lopez-Paz, and A. Szlam, "Optimizing the latent space of generative networks," ICML 2018.
[2] A. Bora, A. Jalal, E. Price, and A. Dimakis, "Compressed sensing using generative models," ICML 2017.
[3] P. Hand, O. Leong, and V. Voroninski, "Phase retrieval under a generative prior," NeurIPS 2018.
[4] G. Jagatap and C .Hegde, "Fast, sample-efficient algorithms for structured phase retrieval," NeurIPS 2017.
[5] F. Shamshad and A. Ahmed, "Robust compressive phase retrieval via deep generative priors," arXiv:1808.05854, 2018.
[6] H. Zhang and Y. Liang, "Reshaped wirtinger flow for solving quadratic system of equations," NeurIPS 2016.

Results: MNIST

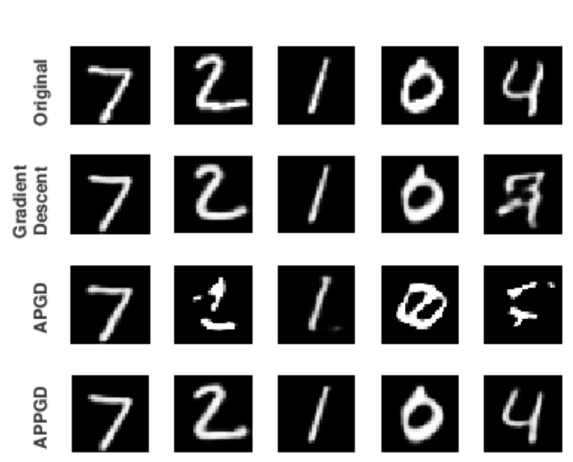

Figure 2. *Reconstruction results on MNIST for three approaches of compressive phase retrieval with $m = 60$ measurements.*

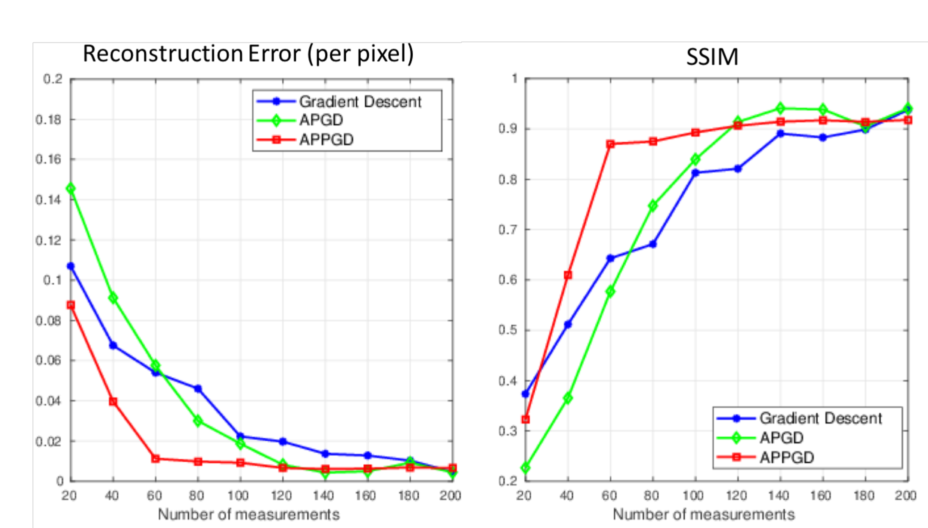

Figure 3. *Performance comparison for three approaches on MNIST test set (10 images).*

Further Information

Full paper available at https://arxiv.org/abs/1903.02707
For additional question please email sasif@ece.ucr.edu

Results: celebA

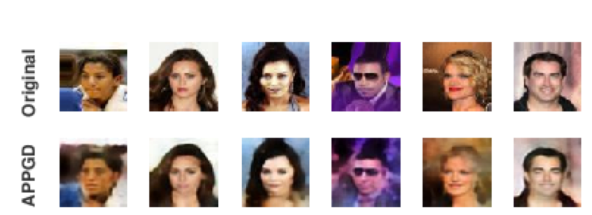

Figure 4. *Reconstruction results on celeb A for APPGD with $m = 1000$ measurements.*

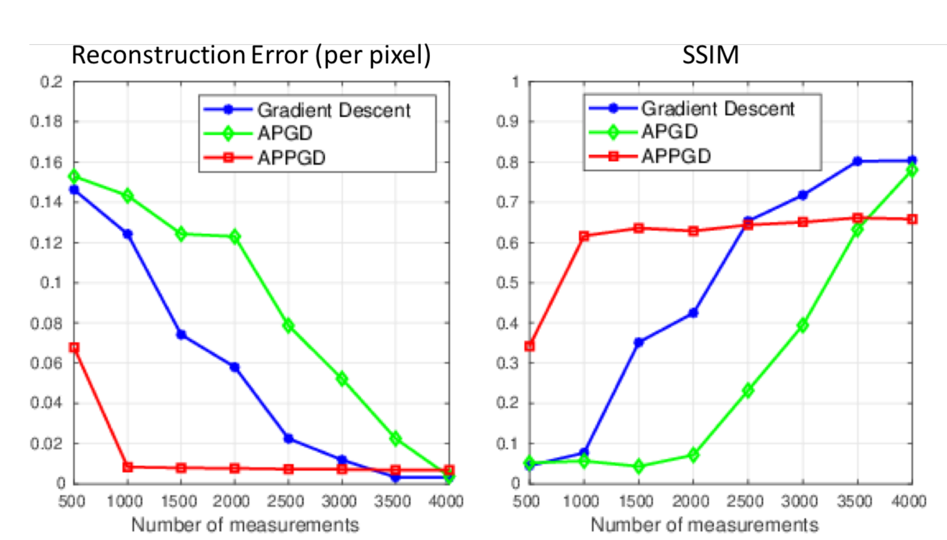

Figure 5. *Performance comparison for three approaches on celeb A test set (10 images).*

Acknowledgments: RH and MA were supported by the DARPA REVEAL Program and an equipment donation from NVIDIA Corporation. VS and CH were supported in part by grants from NSF CCF-1815101, a Faculty Fellowship from the Black and Veatch Foundation, and an equipment donation from NVIDIA Corporation.