

---

# Low Shot Learning with Untrained Neural Networks for Imaging Inverse Problems

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1       Employing deep neural networks as natural image priors to solve inverse problems  
2       either requires large amounts of data to sufficiently train expressive generative  
3       models or can succeed with no data via untrained neural networks. However, very  
4       few works have considered how to interpolate between these no- to high-data  
5       regimes. In particular, how can one use the availability of a small amount of data to  
6       one’s advantage in solving these inverse problems and can a system’s performance  
7       increase as the amount of data increases as well? In this work, we consider solving  
8       linear inverse problems when given a small number of examples of images that are  
9       drawn from the same distribution as the image of interest. Comparing to untrained  
10       neural networks that use no data, we show how one can pre-train a neural network  
11       with a few given examples to improve reconstruction results in compressed sensing.  
12       Our approach leads to improved reconstruction as the amount of available data  
13       increases and is on par with fully trained generative models, while requiring less  
14       than 1% of the data needed to train a generative model. The Latnix author helped  
15       develop the approach and conducted the experiments.

## 16   1 Introduction

17   We study the problem of recovering an image  $\mathbf{x}_0 \in \mathbb{R}^n$  from  $m$  linear measurements of the form  
18    $\mathbf{y}_0 = \mathbf{A}\mathbf{x}_0 \in \mathbb{R}^m$  where  $\mathbf{A} \in \mathbb{R}^{m \times n}$  is a known measurement operator and  $m \leq n$ . The problem’s  
19   difficulty is a result of its ill-posedness due to the underdetermined nature of the system. To resolve  
20   this ambiguity, many approaches enforce that the image must obey a natural image model or prior.

21   One of the most successful image models from deep learning for inverse problems are generative  
22   models such as Generative Adversarial Networks (GANs) [5]. When enforced as a natural image  
23   prior, these models have achieved state-of-the-art results in problems such as compressed sensing  
24   [4, 9, 12, 17, 13], phase retrieval [8, 18, 14], and blind deconvolution [2]. However, these models  
25   require large amounts of data to train and suffer from a non-trivial representation error. On the  
26   opposite end of the data spectrum, recent works have shown that randomly initialized neural  
27   networks can act as natural image priors without any learning such as in [19, 11]. These works  
28   highlighted how convolutional networks exhibit an architectural bias towards natural images, lending  
29   themselves to be successful in a variety of problems [11, 10, 20, 15].

30   Based on these two approaches, we would like to investigate how can one interpolate between these  
31   data regimes. We would like to develop an algorithm that 1) performs just as well as untrained neural  
32   networks with no data and 2) improves performance as the amount of provided data increases.

## 2 Low Shot Learning For Imaging Inverse Problems

We consider the problem of recovering an image  $\mathbf{x}_0 \in \mathbb{R}^n$  from linear measurements of the form  $\mathbf{y}_0 = \mathbf{A}\mathbf{x}_0 \in \mathbb{R}^m$  where  $\mathbf{A} \in \mathbb{R}^{m \times n}$  and  $m \leq n$ . We also assume that  $\mathbf{x}_0$  is drawn from a particular data distribution  $\mathcal{D}$  and that we are given a low number of examples drawn from the same distribution (low shots), i.e., examples  $\mathbf{x}_i \sim \mathcal{D}$  where  $i \in [S]$ . We propose using the range of a deep neural network as a natural image model. In particular, we model the  $\mathbf{x}_0$  as the output of  $\mathcal{G}(\mathbf{z}; \boldsymbol{\theta})$ , where  $\mathbf{z} \in \mathbb{R}^k$  is a latent code and  $\boldsymbol{\theta} \in \mathbb{R}^P$  are the parameters of the neural network.

**Pre-training:** Prior to solving the inverse problem, we propose to first pre-train the network using the low shots that are given by jointly learning its weights and latent space similarly to [3]. Given low shots  $\{\mathbf{x}_i\}_{i=1}^S$ , we aim to find latent codes  $\{\mathbf{z}_i\}_{i=1}^S$  and parameters  $\boldsymbol{\theta}$  that solve

$$\hat{\boldsymbol{\theta}}, \hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_S := \operatorname{argmin}_{\boldsymbol{\theta}, \mathbf{z}_1, \dots, \mathbf{z}_S} \frac{1}{S} \sum_{i=1}^S \mathcal{L}(\mathcal{G}(\mathbf{z}_i; \boldsymbol{\theta}), \mathbf{x}_i).$$

where  $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$  is a loss function. In our experiments, we investigate the use of an  $\ell_2$  loss and an estimate of the kernel Maximum Mean Discrepancy (MMD) [6].

**Solving the inverse problem:** We solve the inverse problem via the image-adaptivity approach of [13]. Using the weights found via pre-training, we begin solving the inverse problem by first optimizing over the latent code space to minimize

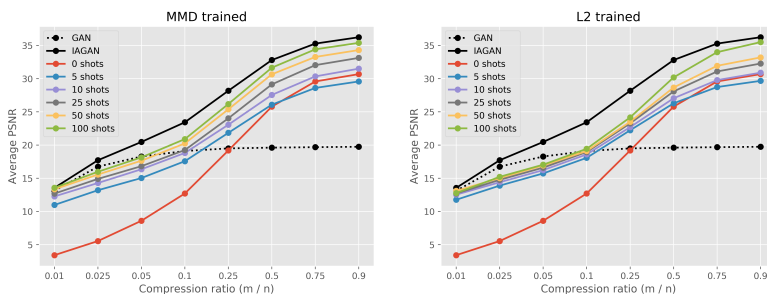
$$f(\mathbf{z}, \boldsymbol{\theta}) := \|\mathbf{A}\mathcal{G}(\mathbf{z}; \boldsymbol{\theta}) - \mathbf{y}_0\|_2^2.$$

The intuition is that we want to first use the semantics regarding the data distribution learned via pre-training the network’s parameters. Once a solution  $\hat{\mathbf{z}}$  is found, we then refine our solution by minimizing  $f(\mathbf{z}, \boldsymbol{\theta})$  with respect to both  $\boldsymbol{\theta}$  and  $\mathbf{z}$  with  $\hat{\boldsymbol{\theta}}$  and  $\hat{\mathbf{z}}$  as our initial iterates.

## 3 Experimental Results

We now consider solving inverse problems with our approach and compare to three different baselines: an untrained neural network, optimizing the latent space of a trained Wasserstein GAN (WGAN) [1, 7], and the image-adaptivity approach of [13] (IAGAN). Each method uses the same DCGAN architecture with a latent code dimension of 128. In each problem, the image of interest is from a hold-out test set from the CelebA dataset [16]. The WGAN was trained on a corpus of over 200,000  $64 \times 64$  RGB celebrity images and our low-shot models were trained on small subsets of this.

**Compressed Sensing:** We want to recover an image  $\mathbf{x}_0 \in \mathbb{R}^n$  from measurements of the form  $\mathbf{y}_0 = \mathbf{A}\mathbf{x}_0 \in \mathbb{R}^m$  where  $\mathbf{A} \in \mathbb{R}^{m \times n}$  has i.i.d.  $\mathcal{N}(0, 1)$  entries with  $m \ll n$ . We refer to amount of undersampling  $\frac{m}{n}$  as the *compression ratio*. We trained our models using the two different loss functions proposed in the previous section for various numbers of shots  $S \in [5, 10, 15, 25, 50, 100]$ .



The figure above compares the average PSNR for each method at various compression ratios over 50 different test images and different loss functions. We note that as the number of shots increases, our method continues to improve and we see comparable performance between our method using an MMD loss and optimizing over the latent code space of a fully trained GAN. While we expect IAGAN to be superior due to being trained with over 200,000 images, the MMD trained model’s performance with 100 images is comparable.

## 68 References

- 69 [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial  
70 networks. *International Conference on Machine Learning (ICML)*, 2017.
- 71 [2] Muhammad Asim, Fahad Shamshad, and Ali Ahmed. Blind image deconvolution using deep  
72 generative priors. *arXiv preprint*, arXiv:1802.04073, 2018.
- 73 [3] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent  
74 space of generative networks. *International Conference on Machine Learning (ICML)*, 2018.
- 75 [4] Ashish Bora, Alexandros G. Dimakis, Ajil Jalal, and Eric Price. Compressed sensing using  
76 generative models. *International Conference on Machine Learning (ICML)*, 2017.
- 77 [5] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil  
78 Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *arXiv preprint*,  
79 arXiv:1406.2661, 2014.
- 80 [6] Arthur Gretton, Karsten M. Borgwardt, Bernhard Schölkopf Malte J. Rasch, and Alexander  
81 Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.
- 82 [7] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville.  
83 Improved training of wasserstein gans. *Neural Information Processing Systems (NIPS)*, 2017.
- 84 [8] Paul Hand, Oscar Leong, and Vladislav Voroninski. Phase retrieval under a generative prior.  
85 *Neural Information Processing Systems (NeurIPS)*, 2018.
- 86 [9] Paul Hand and Vladislav Voroninski. Global guarantees for enforcing generative priors by  
87 empirical risk. *Conference on Learning Theory (COLT)*, 2018.
- 88 [10] Reinhard Heckel. Regularizing linear inverse problems with convolutional neural networks.  
89 *arXiv preprint*, arXiv:1907.0310, 2019.
- 90 [11] Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained  
91 non-convolutional networks. *International Conference on Learning Representations (ICLR)*,  
92 2019.
- 93 [12] Wen Huang, Paul Hand, Reinhard Heckel, and Vladislav Voroninski. A provably conver-  
94 gent scheme for compressive sensing under random generative priors. *arXiv preprint*,  
95 arXiv:1812.04176, 2018.
- 96 [13] Shady Abu Hussein, Tom Tirer, and Raja Giryes. Image-adaptive gan based reconstruction.  
97 *arXiv preprint*, arXiv:1906.05284, 2019.
- 98 [14] Rakib Hyder, Viraj Shah, Chinmay Hegde, and M. Salman Asif. Alternating phase projected  
99 gradient descent with generative priors for solving compressive phase retrieval. *arXiv preprint*,  
100 arXiv:1903.02707, 2019.
- 101 [15] Gauri Jagatap and Chinmay Hegde. Algorithmic guarantees for inverse imaging with untrained  
102 network priors. *arXiv preprint*, arXiv:1906.08763, 2019.
- 103 [16] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the  
104 wild. *Proceedings of the IEEE International Conference on Computer Vision*, pages 3730–3738,  
105 2015.
- 106 [17] Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm  
107 with provable guarantees. *arXiv preprint*, arXiv:1802.08406, 2018.
- 108 [18] Fahad Shamshad and Ali Ahmed. Robust compressive phase retrieval via deep generative priors.  
109 *arXiv preprint*, arXiv:1808.05854:0, 2018.
- 110 [19] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *arXiv preprint*,  
111 arXiv:1711.10925, 2017.
- 112 [20] Dave Van Veen, Ajil Jalal, Mahdi Soltanolkotabi, Eric Price, Sriram Vishwanath, and Alexan-  
113 dros G. Dimakis. Compressed sensing with deep image prior and learned regularization. *arXiv*  
114 *preprint*, arXiv:1806.06438, 2019.