

# Inverse Imaging using Deep Untrained Neural Networks

NeurIPS 2019

---

Gauri Jagatap and Chinmay Hegde

Iowa State University

# Imaging models

---

# Imaging models

Given a  $n$ -dimensional image signal  $x$  and a sensing operator  $A: \mathbb{R}^n \rightarrow \mathbb{R}^m$ , measurements  $y$  take the form:

$$y = Ax$$

**Task:** Recover  $x$  from measurements  $y$ .

- Posed as an optimization problem:

$$\hat{x} = \arg \min_x \|Ax - y\|_2^2 = \arg \min_x \|k - Ax\|_2^2$$

- $n$ -dimensional image  $\rightarrow$  requires  $m = \tilde{E}(n)$  measurements in conventional sensing systems for stable estimation (i.e.  $\hat{x} = x$ ).
- can in general be ill-posed  $\rightarrow$  exact recovery  $\hat{x} = x$  is not guaranteed.

## Examples of inverse problems: observation $y$

$$y = x + n \quad x = \mathcal{I} \quad y = D$$

- Introduce a regularization that makes the problem more tractable.
- Degrees of freedom of natural images is typically lower than  $n$ .
- Constrain the search space to this lower-dimensional set  $S$ .

$$\hat{x} = \arg \min_{x \in S} \|y - x\|_2 = \arg \min_{x \in S} \|k(x)\|_2^2$$

# Leveraging concise representations for regularization

- Natural images have lower dimensional structure → this can be enforced as a prior for inverse imaging problems.

Prior $S$	Data?	Guarantees?
Sparsity (w or w/o structure, total variation)	No	Yes
Deep generative priors	Yes	Yes, Limited

**Table 1:** Low-dimensional priors

# Leveraging concise representations for regularization

- Natural images have lower dimensional structure → this can be enforced as a prior for inverse imaging problems.

Prior $S$	Data?	Guarantees?
Sparsity (w or w/o structure, total variation)	No	Yes
Deep generative priors	Yes	Yes, Limited

**Table 1:** Low-dimensional priors

! Are there other lower-dimensional representations that are more efficient?

# Deep image prior

Prior $S$	Data?	Guarantees?
Sparsity (w or w/o structure, total variation)	No	Yes
Deep generative priors	Yes	Yes, Limited
Deep image prior (+this paper)	No	No ! Yes

**Table 2:** Low-dimensional priors

Deep image prior<sup>1</sup>: Using untrained neural networks as priors.

---

<sup>1</sup>D. Ulyanov et. al., Deep image prior, IEEE CVPR, 2018.

<sup>2</sup>G. Jagatap and C. Hegde, "Algorithmic Guarantees for Inverse Imaging with Untrained Network Priors," NeurIPS (2019).

# Deep image prior

Prior $S$	Data?	Guarantees?
Sparsity (w or w/o structure, total variation)	No	Yes
Deep generative priors	Yes	Yes, Limited
Deep image prior (+this paper)	No	No ! Yes

**Table 2:** Low-dimensional priors

Deep image prior<sup>1</sup>: Using untrained neural networks as priors.

Our contributions<sup>2</sup>:

- New applications of deep image prior for inverse imaging.
  - Linear compressive sensing.
  - Compressive phase retrieval.
- Algorithmic guarantees for reconstruction.

---

<sup>1</sup>D. Ulyanov et. al., Deep image prior, IEEE CVPR, 2018.

<sup>2</sup>G. Jagatap and C. Hegde, "Algorithmic Guarantees for Inverse Imaging with Untrained Network Priors," NeurIPS (2019).



# Deep Neural Networks for Inverse Imaging

---

# Trained deep neural networks for inverse imaging

- Deep neural networks have been used successfully for learning image representations.
- Autoencoders, generative adversarial networks, trained on thousands of images learn latent representations which are lower-dimensional.
- Exploit global statistics across dataset images.

# Trained deep neural networks for inverse imaging

- Deep neural networks have been used successfully for learning image representations.
- Autoencoders, generative adversarial networks, trained on thousands of images learn latent representations which are lower-dimensional.
- Exploit global statistics across dataset images.

! Does the structure of the neural network impose a good prior for image-related problems?

# Trained deep neural networks for inverse imaging

- Deep neural networks have been used successfully for learning image representations.
- Autoencoders, generative adversarial networks, trained on thousands of images learn latent representations which are lower-dimensional.
- Exploit global statistics across dataset images.


! Does the structure of the neural network impose a good prior for image-related problems?

! Can a neural network be used to represent one image, instead of thousands of images?

# Deep Image Prior : Untrained Neural Priors

## Untrained networks as priors

A given image  $I \in \mathbb{R}^{h \times w}$  is said to obey an untrained neural network prior if it belongs to a set  $S$  defined as:  $S := \{f_j = r(\mathbf{w}; \mathbf{z})\}$  where  $\mathbf{z}$  is a (randomly chosen, fixed, dimensionally smaller than  $I$ ) latent code vector and  $r(\mathbf{w}; \cdot)$  has the form as below.



$$r(\mathbf{w}; \mathbf{z}) = \text{vec}^{-1}(\sigma(\mathbf{z})) \otimes \text{vec}^{-1}(\sigma(\mathbf{z})) \otimes \dots \otimes \text{vec}^{-1}(\sigma(\mathbf{z})) \otimes \text{vec}^{-1}(\sigma(\mathbf{z})); (\text{Heckel et. al. 2019})$$

$\sigma(\cdot)$  represents ReLU,  $\otimes$  = bi-linear upsampling,  $\text{vec}^{-1}(\cdot) \in \mathbb{R}^{h \times w}$ ,  $\mathbf{z} = \text{vec}(\mathbf{z}) \in \mathbb{R}^{1 \times 1}$ ,  $a = 1, 2, \dots, a$  and  $a \in \mathbb{R}^a$ .

# Applications of Untrained Neural Priors

---

# Applications of Untrained Neural Priors in Inverse Imaging

Denoising

Super-resolution <sup>3</sup>

Inpainting <sup>4</sup>

---

<sup>3</sup>D. Ulyanov et. al., Deep image prior, IEEE CVPR, 2018.

<sup>4</sup>R. Heckel et. al., Deep Decoder: Concise Image Representations from Untrained Non-convolutional Networks, ICLR, 2019

# Application to Compressive Imaging (Our contribution)

We consider two models for compressive imaging, with operator  $\Phi$ , such that  $y = \Phi(x)$ , and  $x$  takes the forms as below:

- Linear compressive sensing:  $x = \sum_j a_j \phi_j$
- Compressive phase retrieval:  $y_j = |x_j| \phi_j$

where  $a_j \in \mathbb{R}$ ,  $\phi_j \in \mathbb{R}^n$ , and entries of  $\phi_j$  are from  $N(0;1)$  with  $n > m$ .



# Application to Compressive Imaging (Our contribution)

We consider two models for compressive imaging, with operator  $\mathbf{A}$ , such that  $\mathbf{y} = \mathbf{A}\mathbf{x}$ , and  $\mathbf{x}$  takes the forms as below:

- Linear compressive sensing:  $\mathbf{x} = \mathbf{w}$
- Compressive phase retrieval:  $\mathbf{x} = \mathbf{w} \odot \mathbf{e}^j$

where  $\mathbf{w} \in \mathbb{R}^n$ ,  $\mathbf{e}^j \in \mathbb{R}^n$ , and entries of  $\mathbf{e}^j$  are from  $N(0; 1)$  with  $n \gg 1$ .

! both of these problems are ill-posed in this form and require prior information (or regularization) to yield unique solutions.

Pose as the following optimization problem:

$$\min_{\mathbf{w}} \|\mathbf{w}\|_2 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{r}(\mathbf{w}; \mathbf{e}^j) \in S$$

where weights  $\mathbf{w}$  need to be estimated and  $S$  is the range of signals that can be represented as  $\mathbf{y} = \mathbf{r}(\mathbf{w}; \mathbf{e}^j)$ .

Projected gradient descent for  
compressive imaging with  
untrained neural priors

---

# PGD for compressive sensing with untrained neural priors

$$\text{Solve: } \min_{z \in \mathcal{S}} a(z) = \min_{z \in \mathcal{S}} \|k(z)\|_1 \quad k^2$$

---

**Algorithm 1** Net-PGD for linear compressive sensing.

---

```
1: Input:  $y$ ;  $A$ ;  $\mathcal{S}$ ;  $\beta = \log \frac{1}{\epsilon}$ ;  $w^0 = r(w^0; y)$ 
2: for  $i = 1$ ;  $i \leq \beta$ ; do
3:    $w^i = \text{argmin}_{w} \|k(w; y)\|_1$  {gradient step for least squares}
4:    $w^i = \text{argmin}_{w} \|k(w; y)\|_1$  {projection to  $\mathcal{S}$ }
5:    $w^{i+1} = r(w^i; y)$ 
6: end for
7: Output  $w^\beta$ .
```

---

# PGD for compressive phase retrieval with untrained priors

$$\text{Solve: } \min_{z \in S} \|z - y\|_2^2 = \min_{z \in S} \sum_j |z_j - y_j|^2$$

---

**Algorithm 2** Net-PGD for compressive phase retrieval.

---

```

1: Input:  $y$ ;  $\mathbf{A} = \text{vec}(\mathbf{A}_1)$ ;  $\mathbf{A}_1 = \log \mathbf{A}^{-1}$ ;  $\mathbf{0} \leq z \leq 1$  s.t.  $\|z\|_0 \leq k$ 
2: for  $i = 1$ ;  $i \leq k$ ; do
3:    $\hat{z} = \text{sign}(\mathbf{A}^{-1} \mathbf{A} z)$  {phase estimation}
4:    $\tilde{z} = \hat{z} + \alpha (\mathbf{A}^{-1} \mathbf{A} z - \hat{z})$  {gradient step for phase retrieval}
5:    $\mathbf{w} = \arg \min_{\mathbf{w}} \|\mathbf{A} \tilde{z} - \mathbf{w}\|_2$  {projection to  $S$ }
6:    $z = \mathbf{w}$ 
7: end for
8: Output  $\hat{z}$ 

```

---

# Theoretical guarantees

---

# Unique recovery : Set-restricted Isometry Property

To establish unique recovery of  $\mathbf{x}$  from  $\mathbf{y}$ , we need the measurement matrix  $\mathbf{A}$  to satisfy a set-restricted isometry property as follows:

## Lemma: Set-RIP for Gaussian matrices

If an image  $\mathbf{x} \in \mathbb{R}^k$  has a decoder prior (captured in set  $S$ ), where the decoder consists of weights  $\mathbf{w}$  and piece-wise linear activation (ReLU), a random Gaussian matrix  $\mathbf{A} \in \mathbb{R}^{m \times k}$  with elements from  $N(0; 1/\sqrt{m})$ , satisfies  $(S; 1 - \epsilon; 1 + \epsilon)$ -RIP, with probability  $1 - \epsilon^2$ , as long as  $m = \tilde{E} \frac{P}{\epsilon^2} \log \frac{1}{\epsilon}$ , for small constant  $\tilde{E}$  and  $0 < \epsilon < 1$ .

$$(1 - \epsilon)k \leq \|\mathbf{A}\mathbf{x}\|_2 \leq (1 + \epsilon)k$$

## Proof sketch

- For a fixed linearized subspace, the image has a representation of the form

$$= \dots ;$$

where  $\dots$  absorbs all upsampling operations,  $\dots$  is latent code which is fixed and known and  $\dots$  is the direct product of all weight matrices with  $\dots \in \mathbb{R}^1$ .

- An oblivious subspace embedding (OSE) of  $\dots$  takes the form

$$(1 - \dots)k^2 \leq \dots \leq (1 + \dots)k^2;$$

where  $\dots$  is a Gaussian matrix, and holds for all  $\dots$ -dimensional vectors  $\dots$ , with high probability as long as  $\dots = \tilde{E}(\dots^2)$ .

- Counting argument for the number of such linearized networks followed by union of subspaces argument to capture the range of a deep untrained network.

# Theoretical guarantees : Convergence of Net-PGD

## Convergence of Net-PGD for Linear Compressive Sensing

Suppose the sampling matrix  $S$  satisfies  $(S; 1 - \epsilon; 1 + \epsilon)$ -RIP with high probability then, Algorithm 1 produces  $\hat{x}$  such that  $\|\hat{x} - x\|_2 \leq \epsilon \|x\|_2$  and requires  $\frac{1}{\epsilon^2} \log \frac{1}{\epsilon}$  iterations.

### Proof approach:

- $(S; 1 - \epsilon; 1 + \epsilon)$ -RIP for  $x \in S$
- gradient update rule
- exact projection criterion  $\|x^{k+1} - x^k\|_2 \leq \epsilon \|x^k - x\|_2$

to establish the contraction  $\|x^{k+1} - x^k\|_2 \leq \epsilon \|x^k - x\|_2$ , with  $\epsilon < 1$  to guarantee recovery.



# Theoretical guarantees

## Convergence of Net-PGD for Compressive Phase Retrieval

Suppose the sampling matrix satisfies  $(S; 1 - \epsilon; 1 + \epsilon)$ -RIP with high probability, Algorithm 2 produces  $\hat{x}$ , such that  $\| \hat{x} - x \|_2 \leq \epsilon \|x\|_2$ , as long as the weights are initialized such that  $\|x^0 - x\|_2 \leq \epsilon \|x\|_2$  and the number of measurements is  $m = \frac{1}{\epsilon^2} \log \frac{1}{\epsilon}$ .

- $(S; 1 - \epsilon; 1 + \epsilon)$ -RIP for  $S$ ;  $\epsilon > 0$ ;  $\epsilon^{-1} \geq S$
- gradient update rule
- exact projection criterion  $\|x^{+1} - x\|_2 \leq \epsilon \|x - x^0\|_2$
- bound on the error  $\|x^0 - x\|_2$ ,  
 $\epsilon := \epsilon \left( 1 - \text{sign}(\epsilon) \right) \text{sign}(\epsilon)$  (requires good initialization)

to establish the contraction  $\|x^{+1} - x\|_2 \leq \epsilon \|x - x^0\|_2$ , with  $\epsilon < 1$  to guarantee  $\epsilon \geq 0$  or  $D$  for compressive phase retrieval.

# Experiments

---

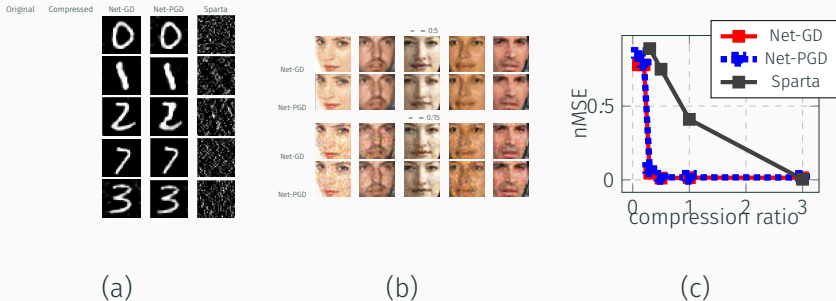
# Linear compressive sensing



Figure 1: (CS) Reconstructed images from linear measurements (at compression rate = 0.1) with (a) = 78 measurements for examples from MNIST, (b) = 1228 measurements for examples from CelebA, and (c) nMSE at different compression rates = for MNIST.

Net-GD: Solve  $\min_{\mathbf{w}} k \quad (r(\mathbf{w}; ))k_2^2$

# Compressive phase retrieval



**Figure 2:** (CPR) Reconstructed images from magnitude-only measurements (a) at compression rate of  $r = 0.3$  for MNIST, (b) for CelebA with Net-GD and Net-PGD, (c) nMSE at different compression rates  $r = 0.3$  for MNIST.

Net-GD: Solve  $\min_{\mathbf{w}} k(r(\mathbf{w}; \cdot))k_2^2$

# Conclusion and future directions

- Our contributions:
  - Novel applications of untrained neural priors to two problems: compressive sensing and phase retrieval with superior empirical performance.
  - Algorithmic guarantees for convergence of PGD for both applications.
- Future directions:
  - Explore other applications such as signal demixing, modulo imaging.
  - Theoretical guarantees for projection oracle.
  - Testing invertible architectures like Glow, instead of decoder structure as prior.

Thank you!

# Deep network configuration

- Fit our example images such that  $r(\mathbf{w}; \cdot)$  (referred as “compressed” image).
- For MNIST images, the architecture was fixed to a 2 layer configuration  $n_1 = 15; n_2 = 15; n_3 = 10$ .
- For CelebA images, a 3 layer configuration with  $n_1 = 120; n_2 = 15; n_3 = 15; n_4 = 10$  was sufficient to represent most images.
- Both architectures use bilinear upsampling operators each with upsampling factor of 2,  $u^2; u = f_{1;2;3}g$ .
- The outputs after each ReLU operation are normalized, by calling for batch normalization subroutine in Pytorch.
- Finally a sigmoid activation is added to the output of the deep network, which smoothens the output; however this is not mandatory for the deep network configuration to work.