Algorithmic Guarantees for Inverse Imaging with Untrained Neural Priors

Gauri Jagatap and Chinmay Hegde
Department of Electrical and Computer Engineering
IOWA STATE UNIVERSITY

Inverse Imaging

Given a d-dimensional image signal x^* and a sensing operator $f(\cdot) : \mathbb{R}^d \to \mathbb{R}^n$, measurements y take the form: $y = f(x^*)$

Task: Recover x^* from measurements y.

- Optimize: $\hat{x} = \arg\min_x L(x) = \arg\min_x \|y - f(x)\|_2^2$
- f can in general be ill-posed; exact recovery $\hat{x} = x^*$ is not guaranteed.

Denoising Super-resolution In-painting Compressive sensing

Untrained Neural Priors

Prior S	Data?	Guarantees?
Sparsity (structure, total variation)	No	Yes
Deep generative priors	Yes	Yes, Limited
Deep image prior (+this paper)	No	No $\to$ Yes

- Deep image prior (D. Ulyanov et. al., CVPR, '18).
  - The structure of the neural network impose a good prior in imaging.
  - Use a neural network to represent one image, instead of thousands.

Our contributions

- Deep image prior for compressive imaging.
- Algorithmic guarantees for reconstruction.

Deep Decoder

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

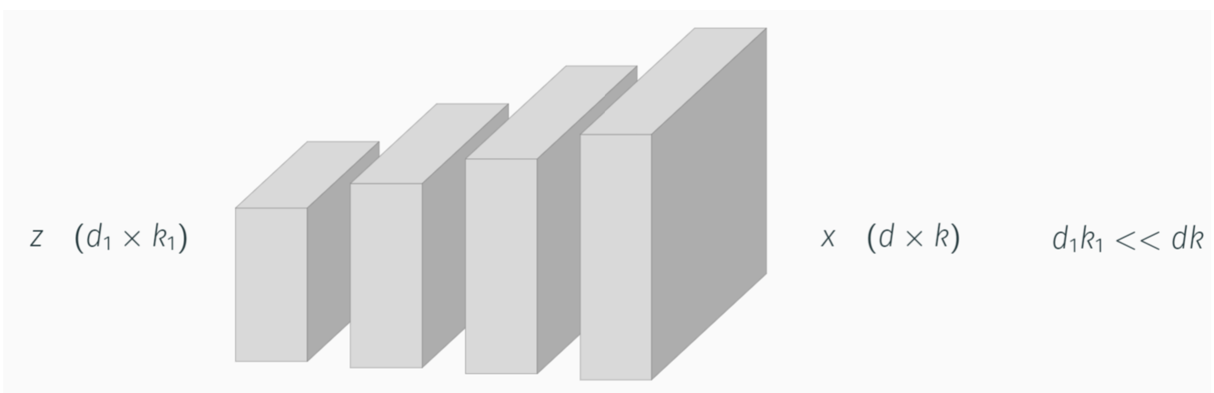

$x = G(\mathbf{w}, z) = U_{L-1}\sigma(Z_{L-1}W_{L-1})W_L = Z_L W_L$, (Heckel et. al., ICLR '19)

$\sigma(\cdot)$ represents ReLU, $Z_i^{d_i\times k_i} = U_{i-1}\sigma(Z_{i-1}W_{i-1})$, for $i = 2 \ldots L$, U is bi-linear upsampling, $z = \text{vec}(Z_1) \in \mathbb{R}^{d_1\times k_1}$, $d_L = d$ and $W_L \in \mathbb{R}^{k_L\times k}$.

Compressive Imaging (CS and CPR)

Compressive imaging, with operator $f : \mathbb{R}^d \to \mathbb{R}^n$, such that $y = f(x^*)$ and f takes the forms as below:

- Linear compressive sensing (CS): $y = Ax^*$
- Compressive phase retrieval (CPR): $y = |Ax^*|$

and entries of A are from $\mathcal{N}(0, 1/n)$ with $n < d$.

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

Solve: $\min_{x,\mathbf{w}} \|y - f(x)\|_2^2$ s.t. $x = G(\mathbf{w}, z) \in \mathcal{S}$.

Reconstruction Algorithm

Algorithm 1 Net-PGD for compressive imaging.

1: **Input:** A, Z_1, η, T, (CPR only) x^0 s.t. $\|x^0 - x^*\|_2 \le \delta_i \|x^*\|_2$.
2: **for** $t = 1, \cdots, T$ **do**
3: $p^t \leftarrow \text{sign}(Ax^t)$ (CPR) or $p^t \leftarrow \mathbf{1}$ (CS) {phase estimation}
4: $v^t \leftarrow x^t - \eta A^\top(Ax^t - y \circ p^t)$ {gradient step}
5: $\mathbf{w}^t \leftarrow \arg\min_{\mathbf{w}} \|v^t - G(\mathbf{w}; z)\|_2^2$ {projection to $\mathcal{S}$}
6: $x^{t+1} \leftarrow G(\mathbf{w}^t; z)$
7: **end for**
8: **Output** $\hat{x} \leftarrow x^T$.

Theoretical Guarantees (I)

Lemma: Set-RIP for Gaussian matrices

If $x \in \mathbb{R}^{d\times 1}$ has a decoder prior $\mathcal{S}$, then $A \in \mathbb{R}^{n\times d}$ with elements from $\mathcal{N}(0, 1/n)$, satisfies $(\mathcal{S}, 1-\alpha, 1+\alpha)$-RIP, with probability $1 - e^{-c\alpha^2 n}$, if $n = O(\frac{k_1}{\alpha^2}\sum_{l=2}^{L} k_l \log d)$, for small constant c and $0 < \alpha < 1$.

$$(1-\alpha)\|x\|_2^2 \le \|Ax\|_2^2 \le (1+\alpha)\|x\|_2^2.$$

- For a fixed linearized subspace, x has a representation: $x = UZw$, where U absorbs all upsampling operations, Z is latent code which is fixed and known and w is the direct product of all weight matrices with $w \in \mathbb{R}^{k_1}$.
- Oblivious subspace embedding (OSE) of x:
$$(1-\alpha)\|x\|_2^2 \le \|Ax\|_2^2 \le (1+\alpha)\|x\|_2^2,$$
where A is a Gaussian matrix, and holds for all possible $w \in \mathbb{R}^{k_1}$, with high probability , if $n = O(k_1/\alpha^2)$.
- Union of all possible linearized subspaces to capture the range of a deep untrained network.

Theoretical Guarantees (II)

Convergence of Net-PGD for Compressive Imaging

Suppose $A^{n\times d}$ satisfies $(\mathcal{S}, 1-\alpha, 1+\alpha)$-RIP with high probability, η is small enough, (for CPR, the weights are initialized such that $\|x^0 - x^*\|_2 \le \delta_i \|x^*\|_2$ and the number of measurements is $n = O\left(k_1 \sum_{l=2}^{L} k_l \log d\right)$), Net-PGD produces $\hat{x}$, s.t. $\|\hat{x} - x^*\|_2 \le \epsilon$.

- $(\mathcal{S}, 1-\alpha, 1+\alpha)$-RIP for $x^*, x^t, x^{t+1} \in \mathcal{S}$,
- gradient update rule,
- exact projection criterion $\|x^{t+1} - v^t\|_2 \le \|x^* - v^t\|_2$,
- bound $\varepsilon_p^t := A^\top Ax^* \circ (1 - \text{sign}(Ax^*) \circ \text{sign}(Ax^t))$ (requires delta-close initialization), phase estimation error,

to establish *linear convergence of Net-PGD* $\|x^{t+1} - x^*\|_2 \le \nu \|x^t - x^*\|_2$, with $\nu < 1$.

Results

Net-GD: Solve $\min_{\mathbf{w}} \|y - f(G(\mathbf{w}; z))\|_2^2$, nMSE: $\|\hat{x} - x^*\|_2^2 / \|x^*\|_2^2$

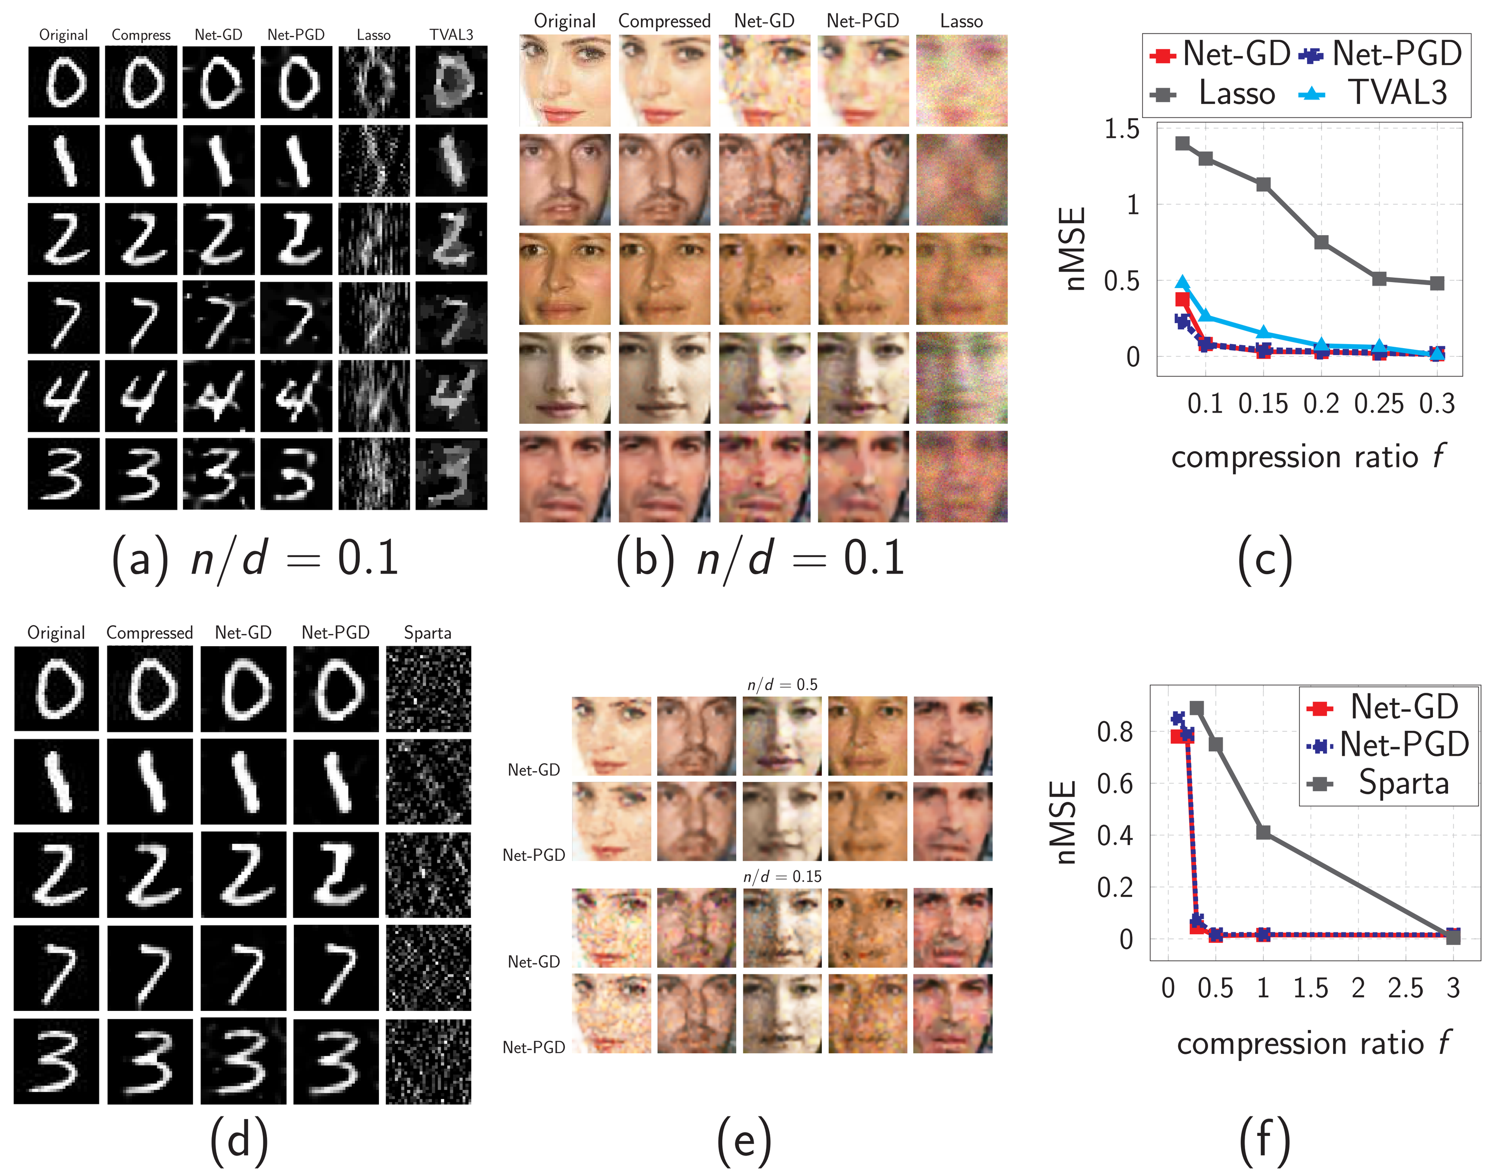

Figure 1: CS on (a) MNIST images (b) CelebA images (c) digit '0' of MNIST; CPR on (d) MNIST images (b) CelebA images (c) fixed celeb image. [Code:https://github.com/GauriJagatap/invimaging-deeppriors]

Acknowledgements

This work was supported in part by NSF grants CAREER CCF-2005804, CCF-1815101, and a faculty fellowship from the Black and Veatch Foundation.