Inverse Problems Under a Learned Generative Prior

Paul Hand

Collaborators: Vladislav Voroninski, Wen Huang, Oscar Leong, Reinhard Heckel

Examples of inverse problem

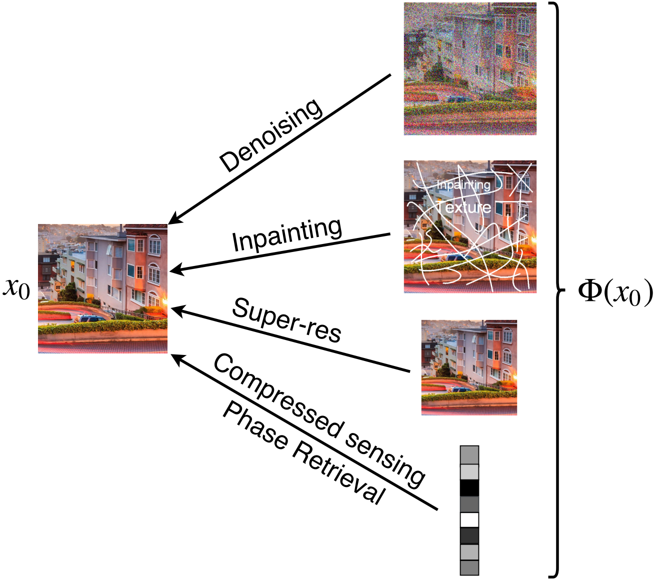

A common prior: sparsity

Original image

Sparse approximation

Sparsity can be optimized via a convex relaxation

$$\min_{x \in \mathbb{R}^n} \|x\|_0 \quad \text{s.t. } \Phi(x) = \Phi(x_0) \quad \xrightarrow{\text{Relaxation}} \quad \min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{s.t. } \Phi(x) = \Phi(x_0)$$

Recovery guarantee for sparse signals

Fix k-sparse vector $x_0 \in \mathbb{R}^n$.
Let $A \in \mathbb{R}^{m \times n}$ be a random gaussian matrix with $m = \Omega(k \log n)$.

$$\begin{aligned} \min \quad & \|x\|_1 \\ \text{s.t.} \quad & Ax = Ax_0 \end{aligned} \tag{L1}$$

Theorem (Candes, Romberg, Tao. 2004. Donoho, 2004.)

The global minimizer of (L1) is x_0 with high probability.

Generative models learn to impressively sample from complex signal classes

Faces

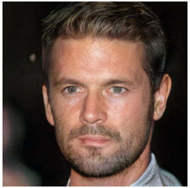

Bedrooms

MRI scans

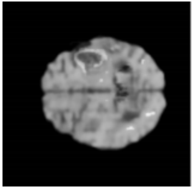

Cells

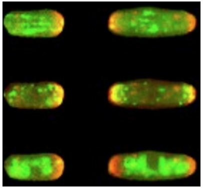

Fingerprints

Main takeaways

1. Generative models provide SOTA performance in imaging inverse problems
2. Generative models provide lower dimensional priors that can be directly and efficiently exploited
3. Suggests new framework for experimental scientists in imaging

Main takeaways

1. Generative models provide SOTA performance in imaging inverse problems

2. Generative models provide lower dimensional priors that can be directly and efficiently exploited

3. Suggests new framework for experimental scientists in imaging

How are generative models used in inverse problems?

1. Train generative model to output signal class:

$$G : \mathbb{R}^k \rightarrow$$

2. Directly optimize over range of generative model via empirical risk:

$$\min_{z \in \mathbb{R}^k} \left\| \Phi(G(z)) - \Phi(x_0) \right\|^2$$

Generative models provide SOTA performance

- Denoising
- Inpainting
- Super resolution
- Compressed sensing
  - 5-10X less measurements than Lasso (Bora et al.)
  - 2 orders of magnitude speedup in MRI imaging (Mardani et al.)

Main takeaways

1. Generative models provide SOTA performance in imaging inverse problems

2. Generative models provide lower dimensional priors that can be directly and efficiently exploited
   - Deep Compressive Sensing
   - Deep Phase Retrieval

3. Suggests new framework for experimental scientists in imaging

Deep Compressive Sensing

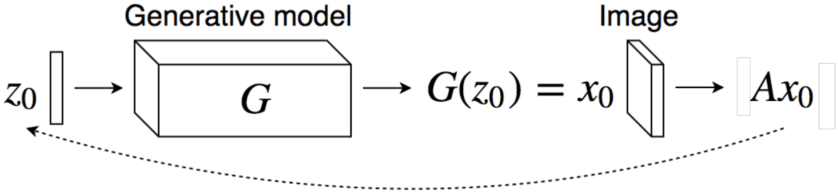

$$\min_{z \in \mathbb{R}^k} \left\| AG(z) - Ax_0 \right\|^2$$

Bora, Jalal, Price, Dimakis

Initial theory for generative priors analyzed global minimizers, which may be hard to find

Theorem (Bora, Jalal, Price, Dimakis)

Let $G : \mathbb{R}^k \to \mathbb{R}^n$ be a d-layer NN with ReLU activations. Let $C \in \mathbb{R}^{m \times n}$ have iid. Gaussian entries with $m = O(kd \log n)$. Fix $y \in \mathbb{R}^n$ and $b = Cy$. ***If*** *$\hat{z}$* ***minimizes*** *$\|b - CG(z)\|_2$* ***to within*** *ϵ, then, with probability $1 - e^{-\Omega(m)}$,*

$$||G(\hat{z}) - y\|_2 \leq 6 \min_x \|G(z) - y\|_2 + 2\epsilon.$$

Random generative priors allow rigorous recovery guarantees

$$
\begin{aligned}
\text{Let:} \quad & \mathcal{G} : \mathbb{R}^k \to \mathbb{R}^n \\
& \mathcal{G}(z) = \text{relu}(W_d \ldots \text{relu}(W_2 \text{relu}(W_1 z)) \ldots) \\
\text{Given:} \quad & W_i \in \mathbb{R}^{n_i \times n_{i-1}}, A \in \mathbb{R}^{m \times n}, y := A\mathcal{G}(z_0) \in \mathbb{R}^m \\
\text{Find:} \quad & x_0
\end{aligned}
$$

- **Expansivity**: Let $n_i > c n_{i-1} \log n_{i-1}$
- **Gaussianicity**: Let W_i and A have iid Gaussian entries.
- **Biasless**: No bias terms in $\mathcal{G}$.

Compressive sensing with random generative prior has favorable geometry for optimization

Theorem (Hand and Voroninski)

Let $d \geq 2$. If

1. *$\mathcal{G}$ is gaussian and sufficiently expansive,*
2. $m = \Omega(kd \log(\Pi_{i=1}^{d} n_i)$,

then w.h.p. there exists v_z such that $D_{-v_z} f < 0$ for all z outside a neighborhood of z_0 and $-\rho_d z_0$. Additionally, 0 is a local max.

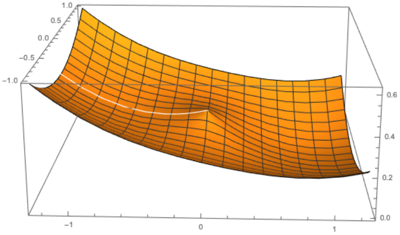

Proof Outline

- Explicit formula for $v_z = \nabla \text{objective}(z)$
- Explicit formula for $g_z = \mathbb{E}\, \nabla \text{objective}(z)$
- Show $v_z \approx g_z$ uniformly in z
- Show $g_z \neq 0$ away from $z_0, 0, -\rho_d z_0$

Deterministic Condition for Recovery

A matrix $W \in \mathbb{R}^{n \times k}$ satisfies the **Weight Distribution Condition** with constant ϵ if for all $x, y \neq 0 \in \mathbb{R}^k$,

$$\left\| \sum_{i=1}^{n} 1_{w_i \cdot x > 0} 1_{w_i \cdot y > 0} \cdot w_i w_i^T - Q_{x,y} \right\| \leq \epsilon, \text{ with } Q_{x,y} = \frac{\pi - \theta}{2\pi} I + \frac{\sin \theta}{2\pi} M_{\hat{x} \leftrightarrow \hat{y}}.$$

Here, w_i^T is the ith row of W; $M_{\hat{x} \leftrightarrow \hat{y}} \in \mathbb{R}^{k \times k}$ is the matrix such that $\hat{x} \mapsto \hat{y}$, $\hat{y} \mapsto \hat{x}$, and $\hat{z} \mapsto 0$ for all $z \in (\{x, y\})^{\perp}$; $\theta = \angle(x, y)$.

Compressive sensing with random generative prior has a provably convergent subgradient descent algorithm

Theorem (Huang, Heckel, Hand, Voroninski)

Let $d \geq 2$. If

1. *$\mathcal{G}$ is gaussian and sufficiently expansive,*
2. $m = \Omega(kd\log(\Pi_{i=1}^{d} n_i))$,
3. *measurements have sufficiently small noise*

then w.h.p. a subgradient descent with a twist converges to within the noise level of z_0.

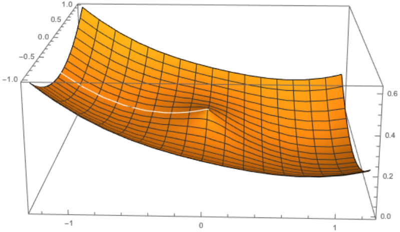

Guarantees for compressive sensing under generative priors have been extended to convolutional architectures

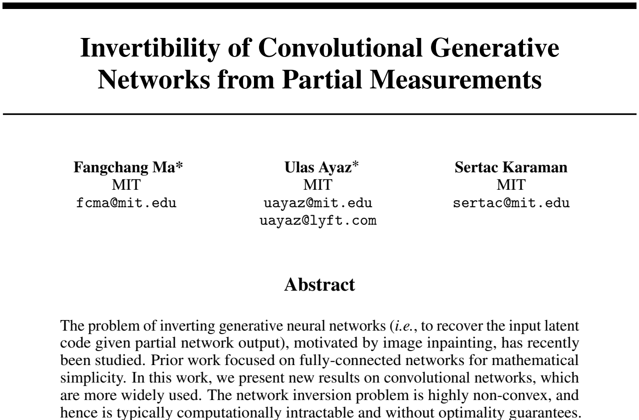

Why can generative models outperform sparsity models?

Generative models can admit lower dimensional representations than sparsity

Why can generative models outperform sparsity models?

Generative models can admit lower dimensional representations than sparsity

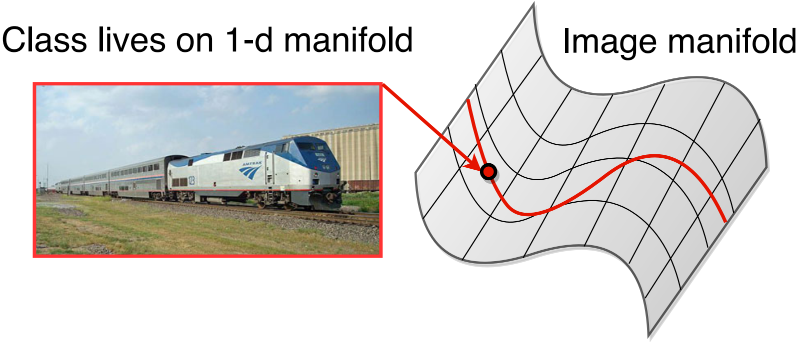

Why can generative models outperform sparsity models?

Generative priors allow direct optimization

$$\min_{x \in \mathbb{R}^n} \|x\|_0 \quad \text{s.t. } \Phi(x) = \Phi(x_0) \quad \xrightarrow{\text{Relaxation}} \quad \min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{s.t. } \Phi(x) = \Phi(x_0)$$

$$\min_{z \in \mathbb{R}^k} \left\| \Phi(G(z)) - \Phi(x_0) \right\|^2$$

Sparsity appears to fail in Compressive Phase Retrieval

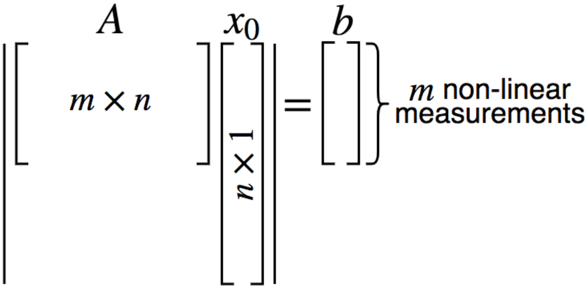

Open problem: there is no known efficient algorithm to recover s-sparse x_0 from $O(s)$ generic phaseless measurements

Our formulation: Deep Phase Retrieval

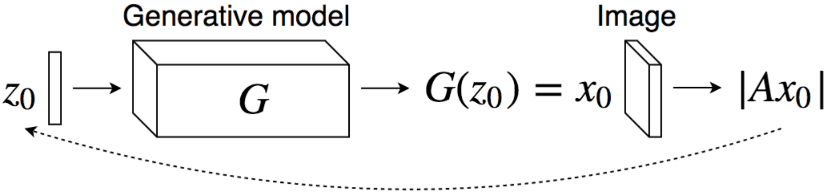

$$\min_{z \in \mathbb{R}^k} \left\| |AG(z)| - |Ax_0| \right\|^2$$

Generative priors can be efficient exploited for compressive phase retrieval

Assumptions

1. *# measurements* $= O(k)$
2. *network layers are sufficiently expansive*
3. *A and weights of G have i.i.d. Gaussian entries*

Theorem (Hand, Leong, Voroninski)

The objective function has a strict descent direction outside of two small neighborhoods of the minimizer and a negative multiple thereof with high probability.

Comparison on MNIST

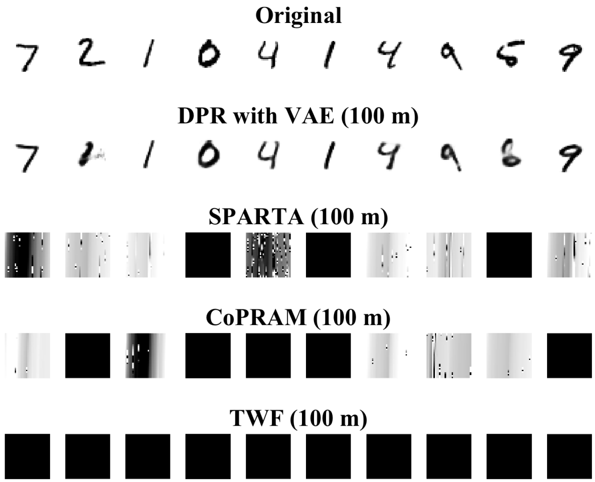

Main takeaways

1. Generative models provide SOTA performance in imaging inverse problems
2. Generative models provide lower dimensional priors that can be directly and efficiently exploited
3. Suggests new framework for experimental scientists in imaging

New workflow for scientists

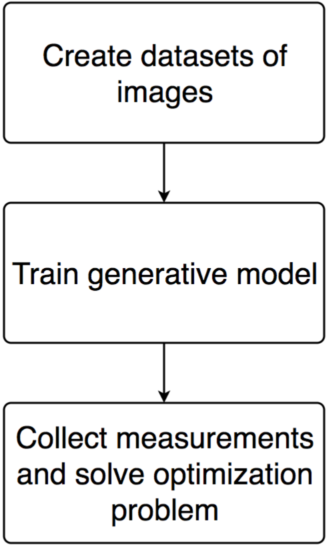

Concrete steps have already been taken

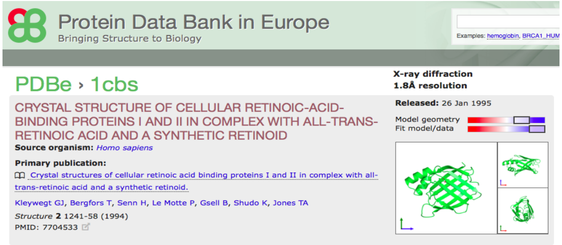

NYU, Facebook Offer Dataset for MRI Artificial Intelligence Project

NYU and Facebook have released a large set of open source MRI data as they work to use artificial intelligence to speed up the diagnostic process.

Further Theory Needed

1. More realistic models of neural nets
2. Other inverse problems
3. Make proofs simpler
4. Understanding of adversarial examples
5. Supplement sparsity approaches in other contexts

Main takeaways

1. Generative models provide SOTA performance in imaging inverse problems

2. Generative models provide lower dimensional priors that can be directly and efficiently exploited

3. Suggests new framework for experimental scientists in imaging