Signal Recovery with Generative Priors

Paul Hand

Northeastern University

Collaborators: V. Voroninski, O. Leong, W. Huang, R. Heckel, A. Ahmed, M. Asim, J. Cocola, M. Daniels

Funding: National Science Foundation

Examples of inverse problems

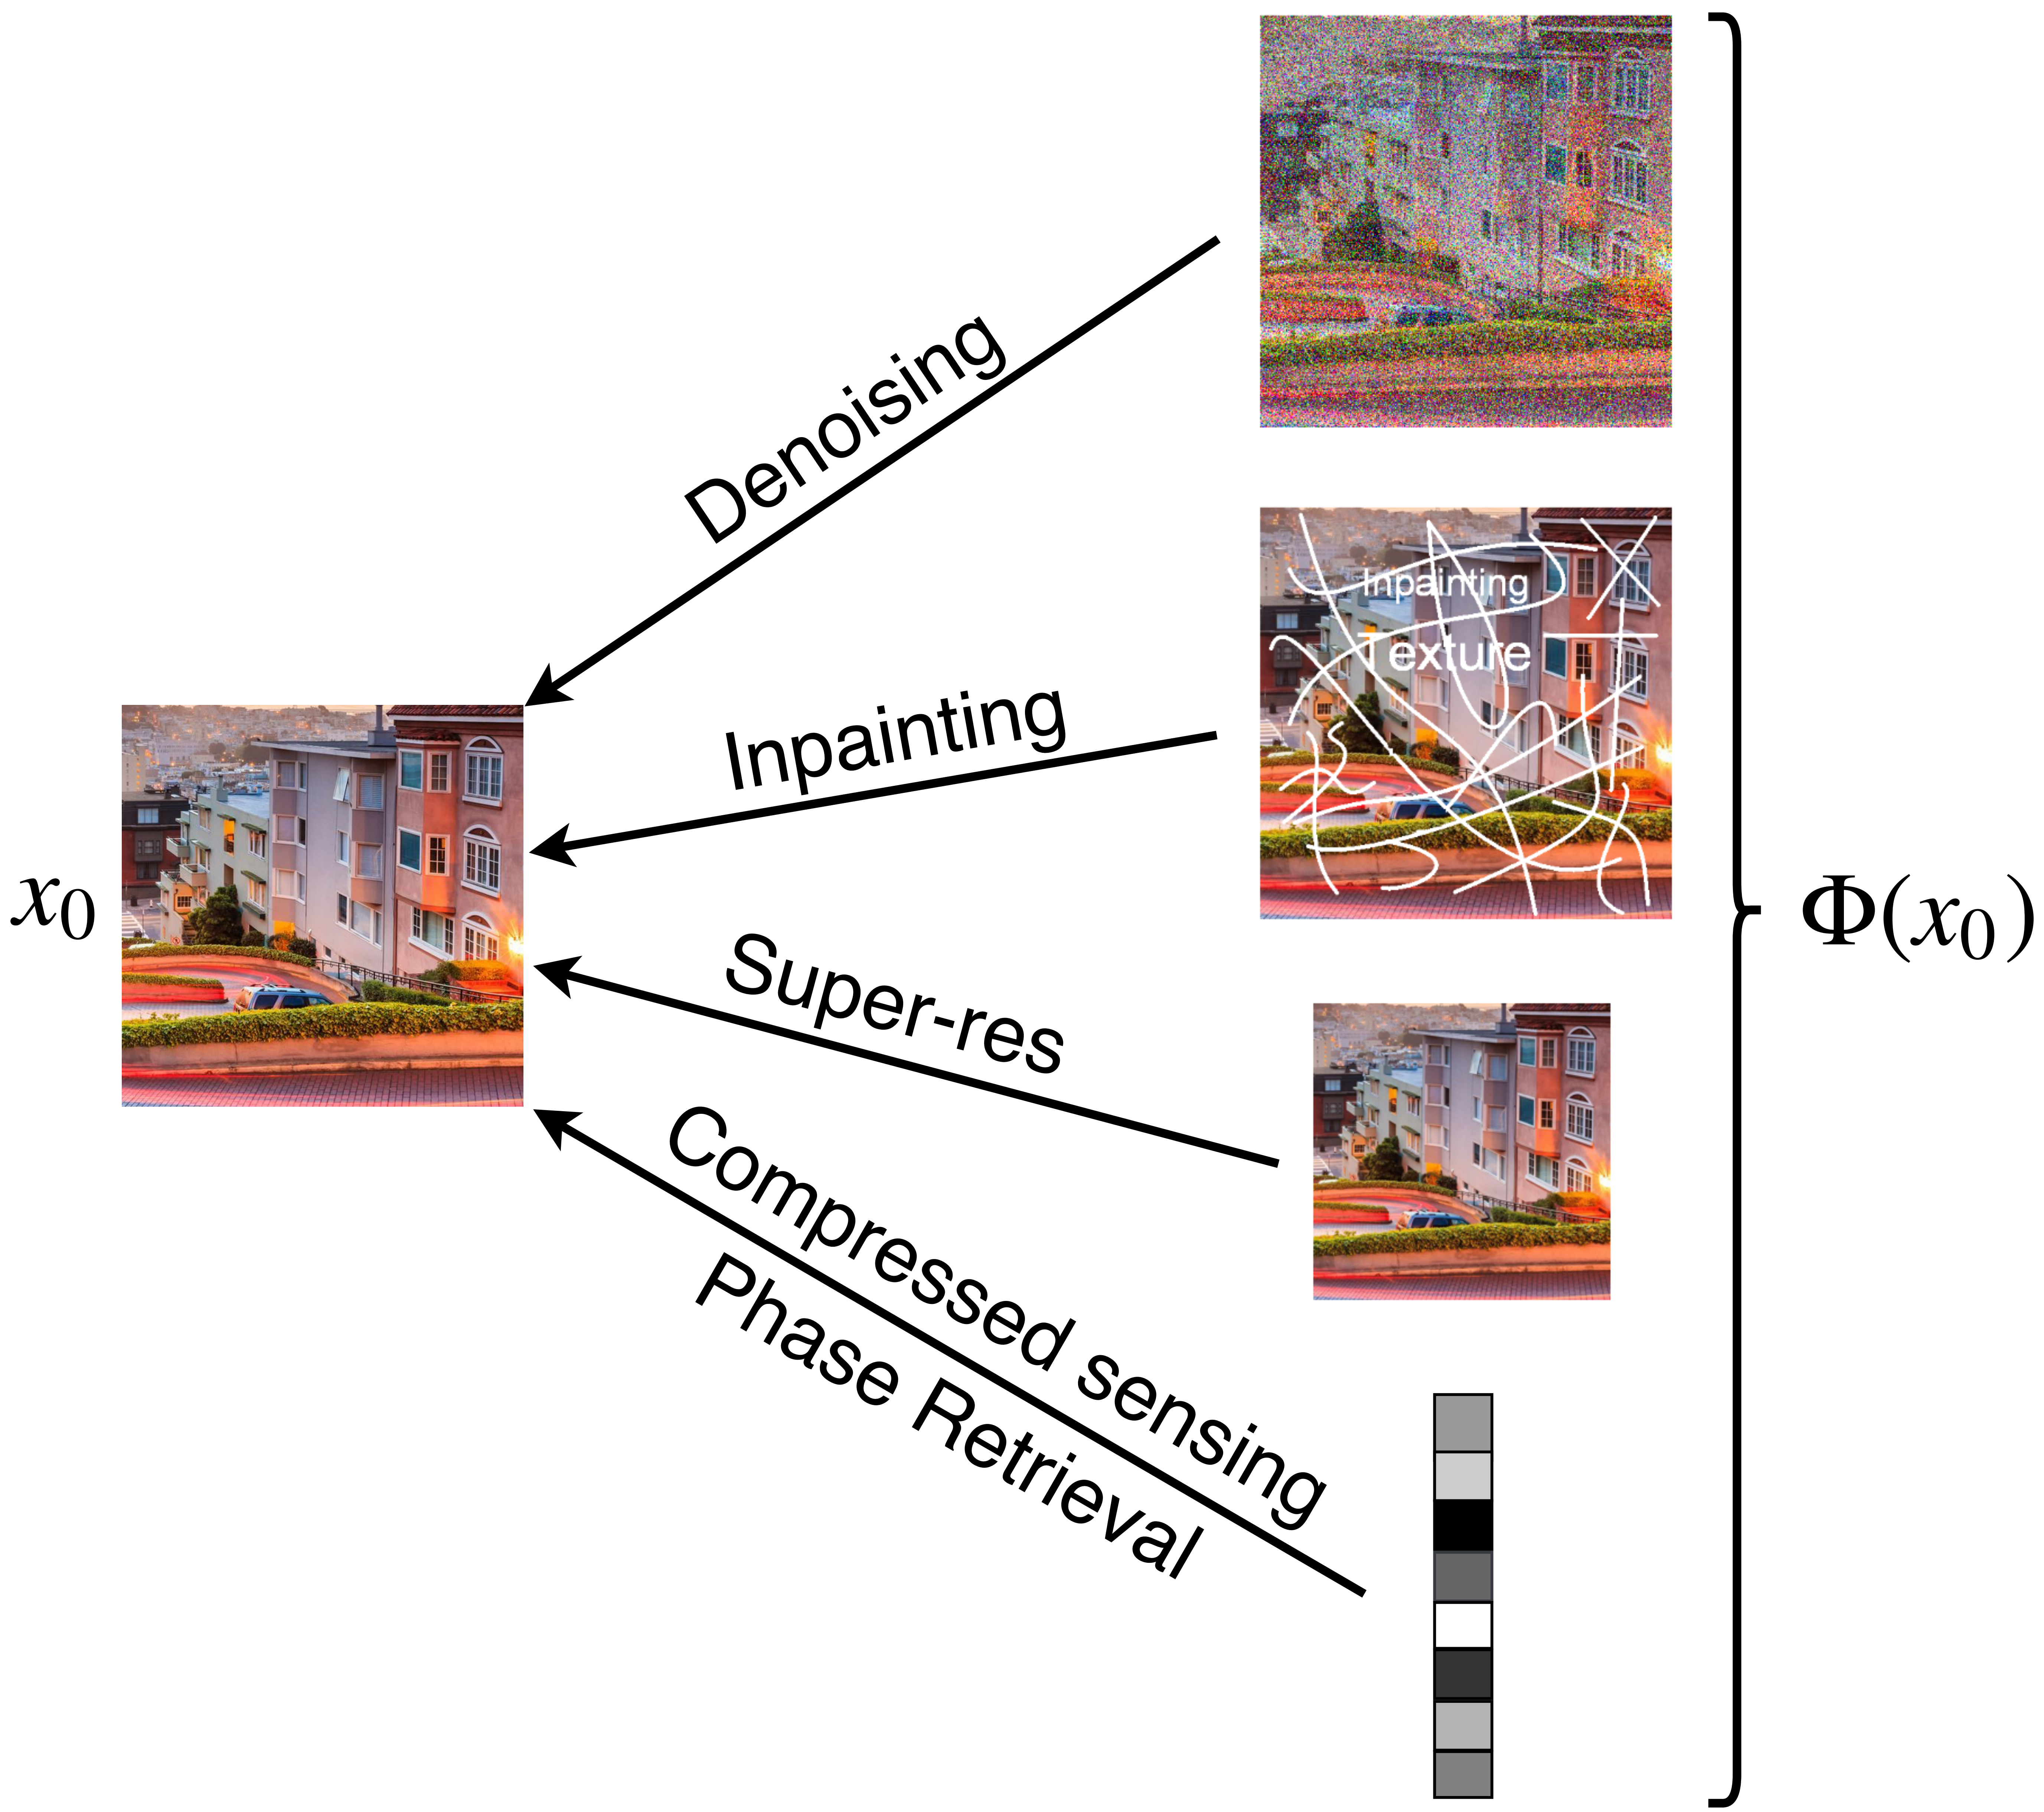

Sparsity can be sometimes be optimized via a convex relaxation

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

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.

Sparsity appears to fail in Compressive Phase Retrieval

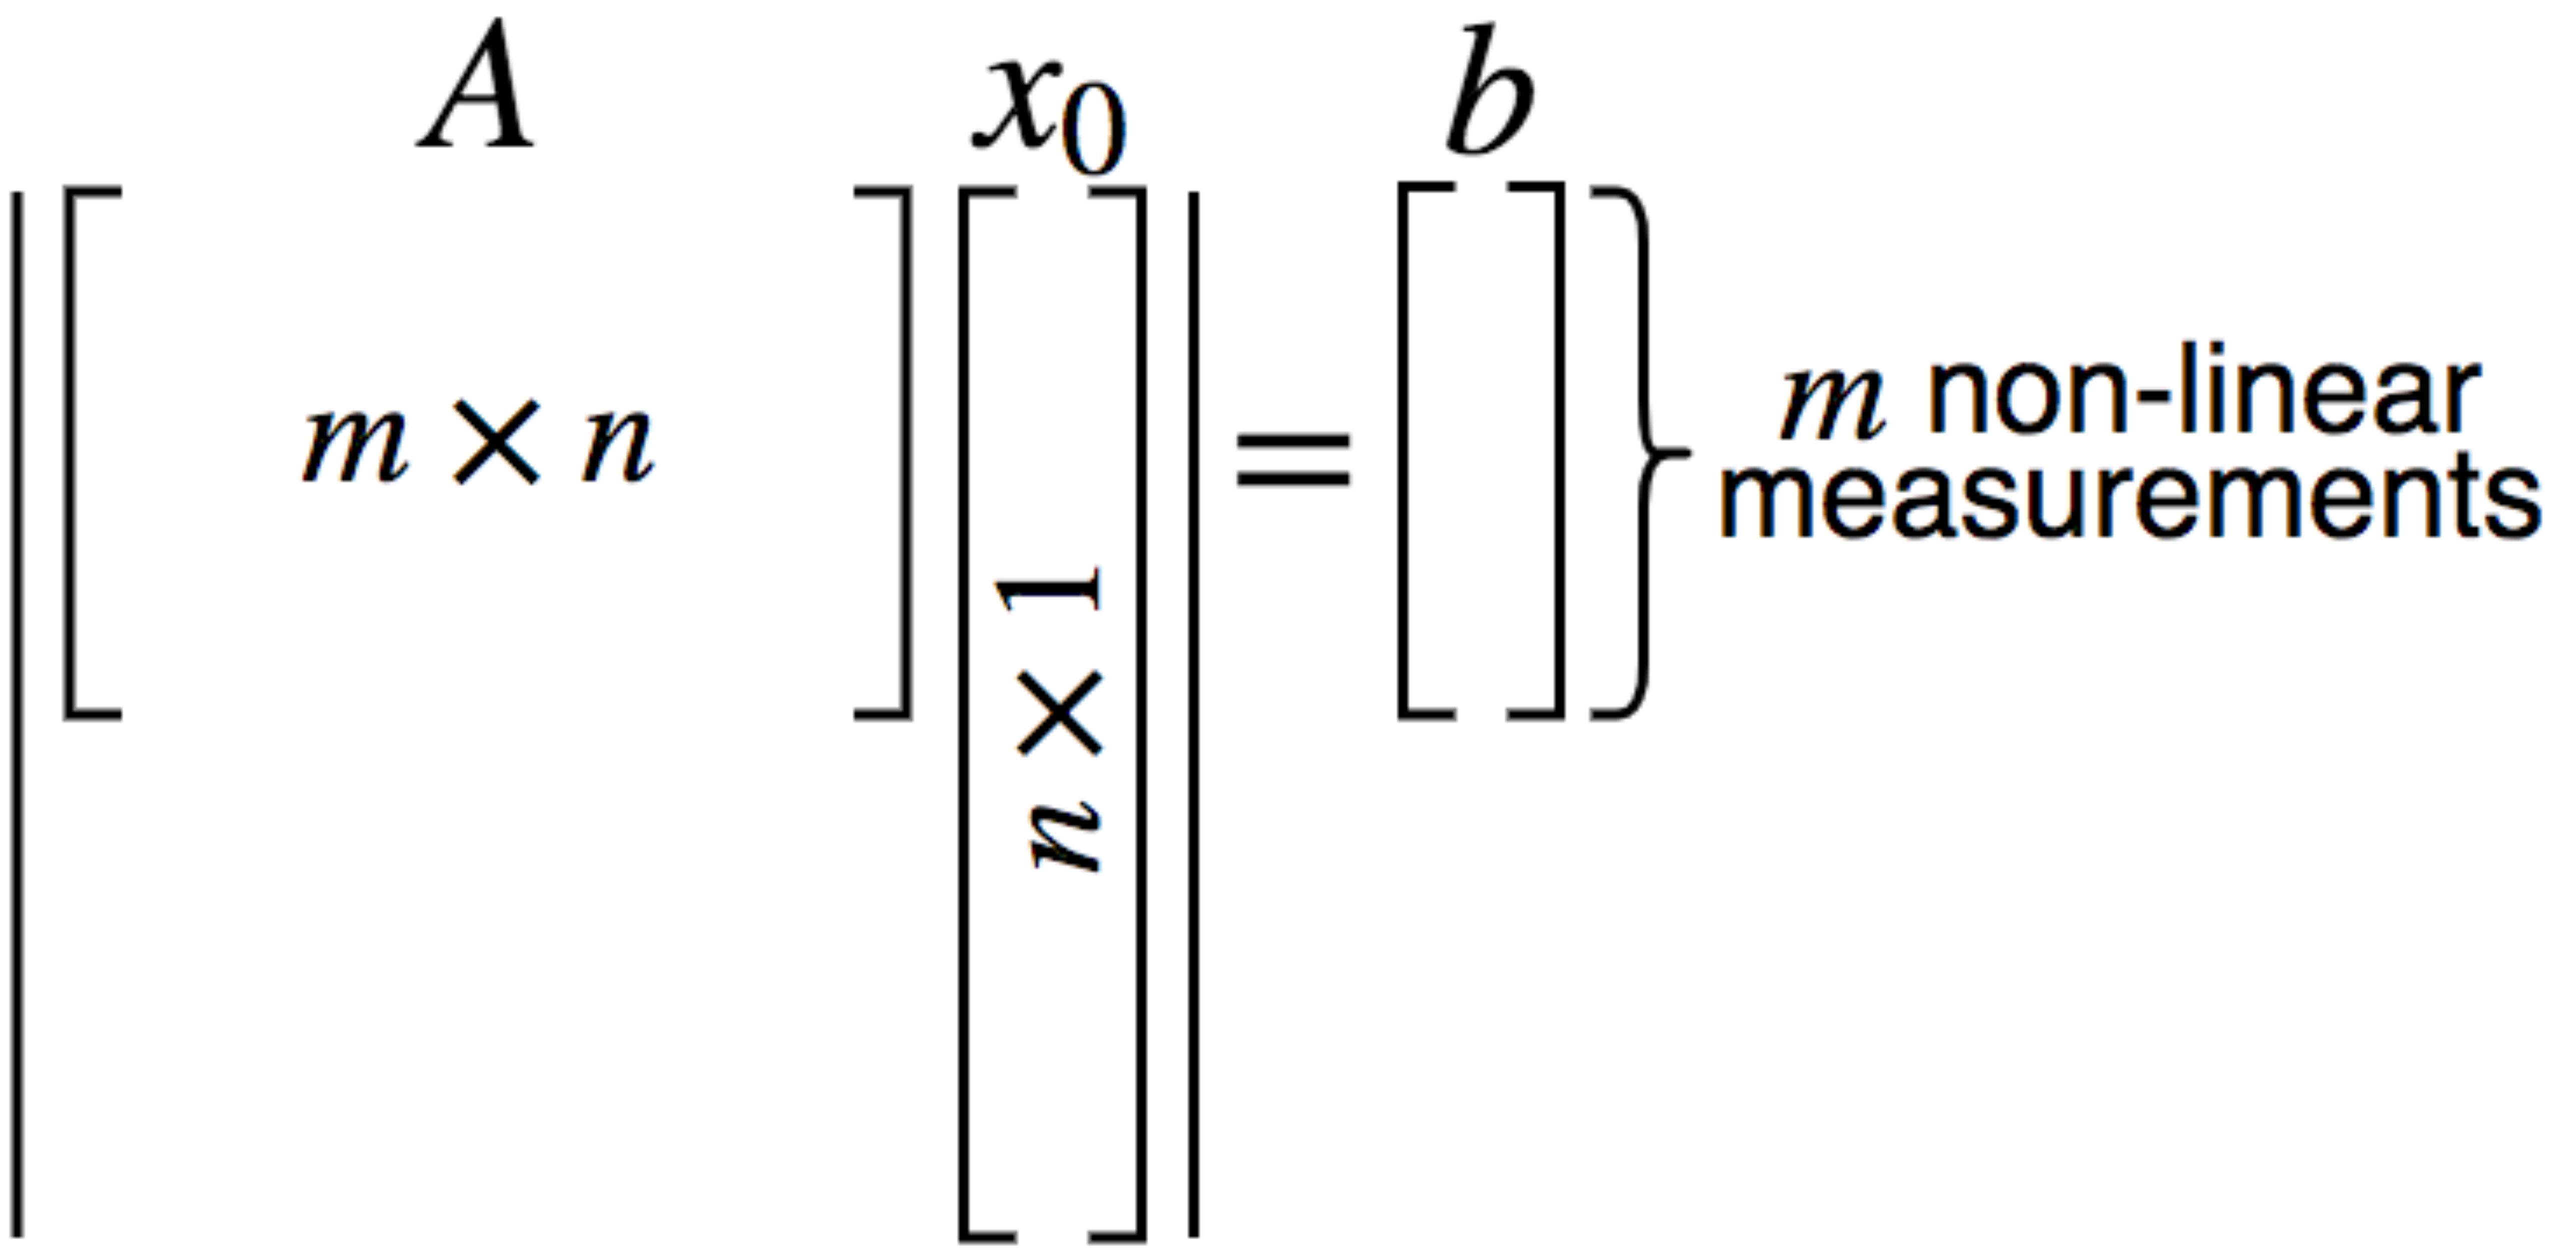

Open problem: there is no known efficient algorithm to recover s-sparse x_0 from $O(s)$ generic measurements

With generic measurements, Sparse PhaseLift gives suboptimal sample complexity

$$\begin{aligned} \min \quad & \lambda \operatorname{tr}(X) + \|X\|_1 \\ \text{s.t.} \quad & a_i^* X a_i = a_i^* (x_0 x_0^*) a_i, \quad i = 1 \ldots m \\ & X \succeq 0 \end{aligned}$$

Theorem (Li and Voroninski, 2012)

If $x_0 \in \mathbb{R}^n$ is s-sparse with constant-magnitude coefficients, and $a_i \sim \mathcal{N}(0, I)$, then with high probability: If $\exists \lambda \in \mathbb{R}$ such that $x_0 x_0^*$ minimizes Sparse PhaseLift, $m = \Omega(s^2 / \log^2 n)$.

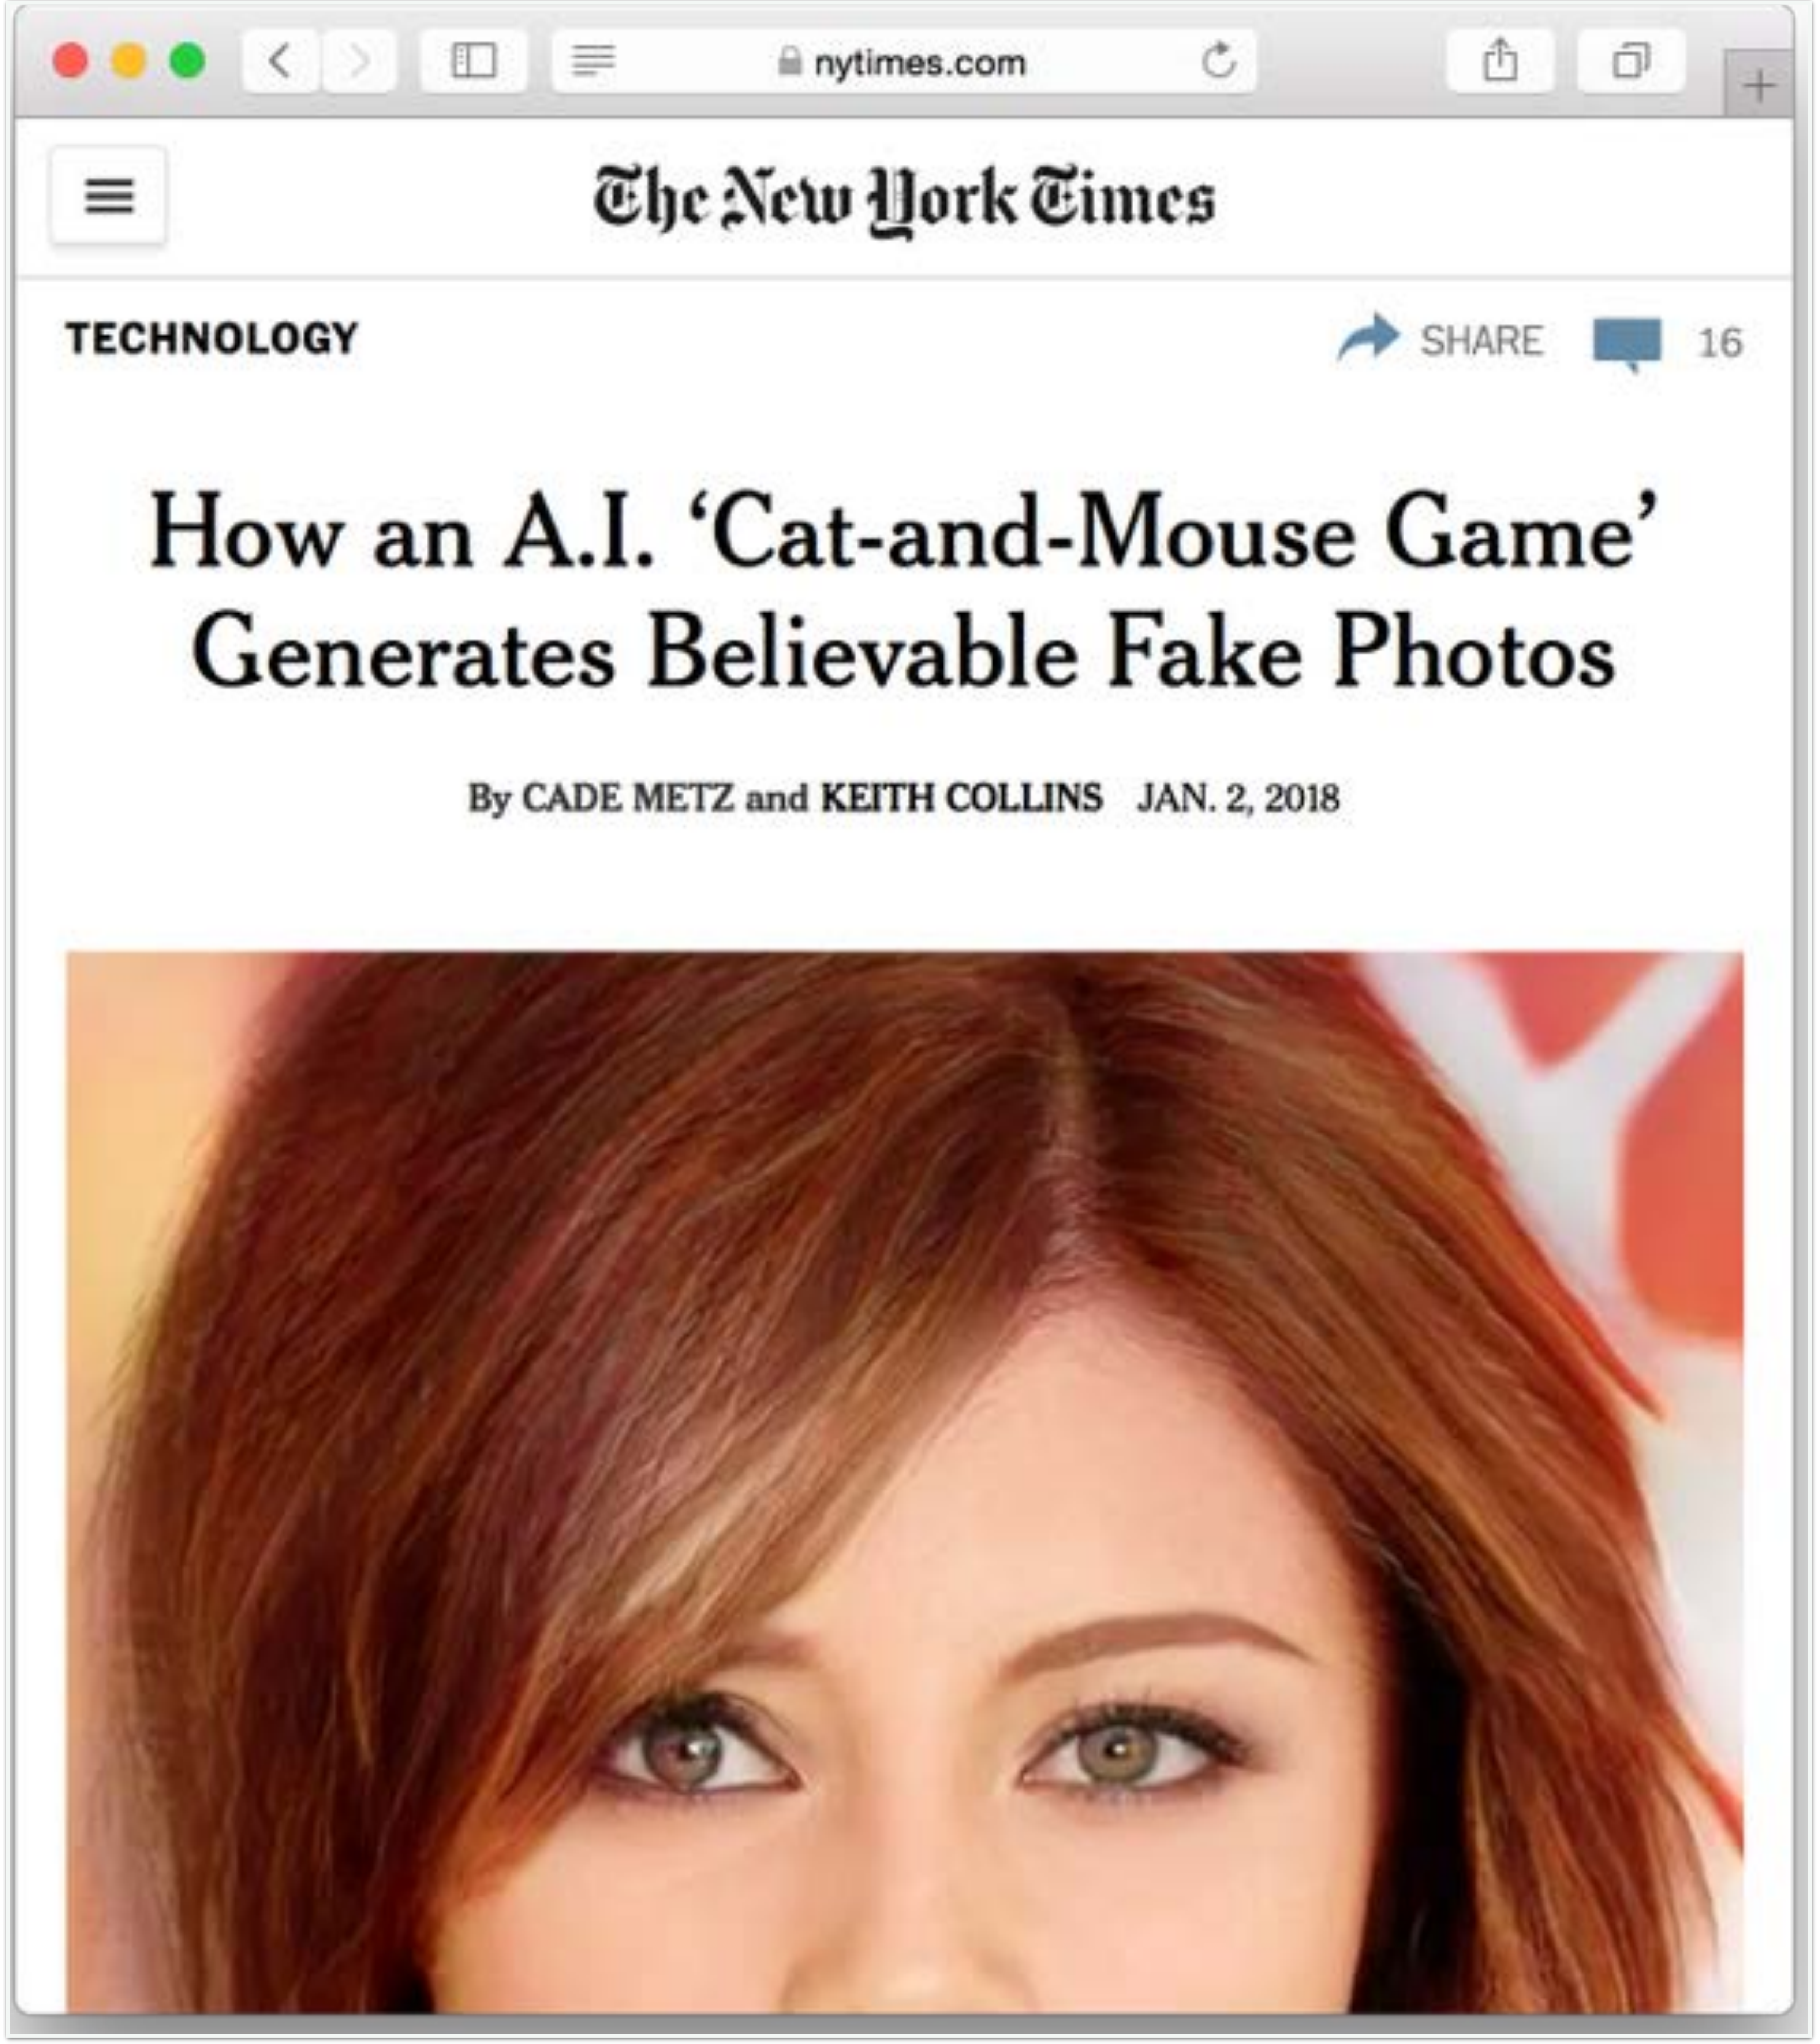
nytimes.com

The New York Times

TECHNOLOGY

SHARE 16

How an A.I. 'Cat-and-Mouse Game' Generates Believable Fake Photos

By CADE METZ and KEITH COLLINS JAN. 2, 2018

Deep Generative Models

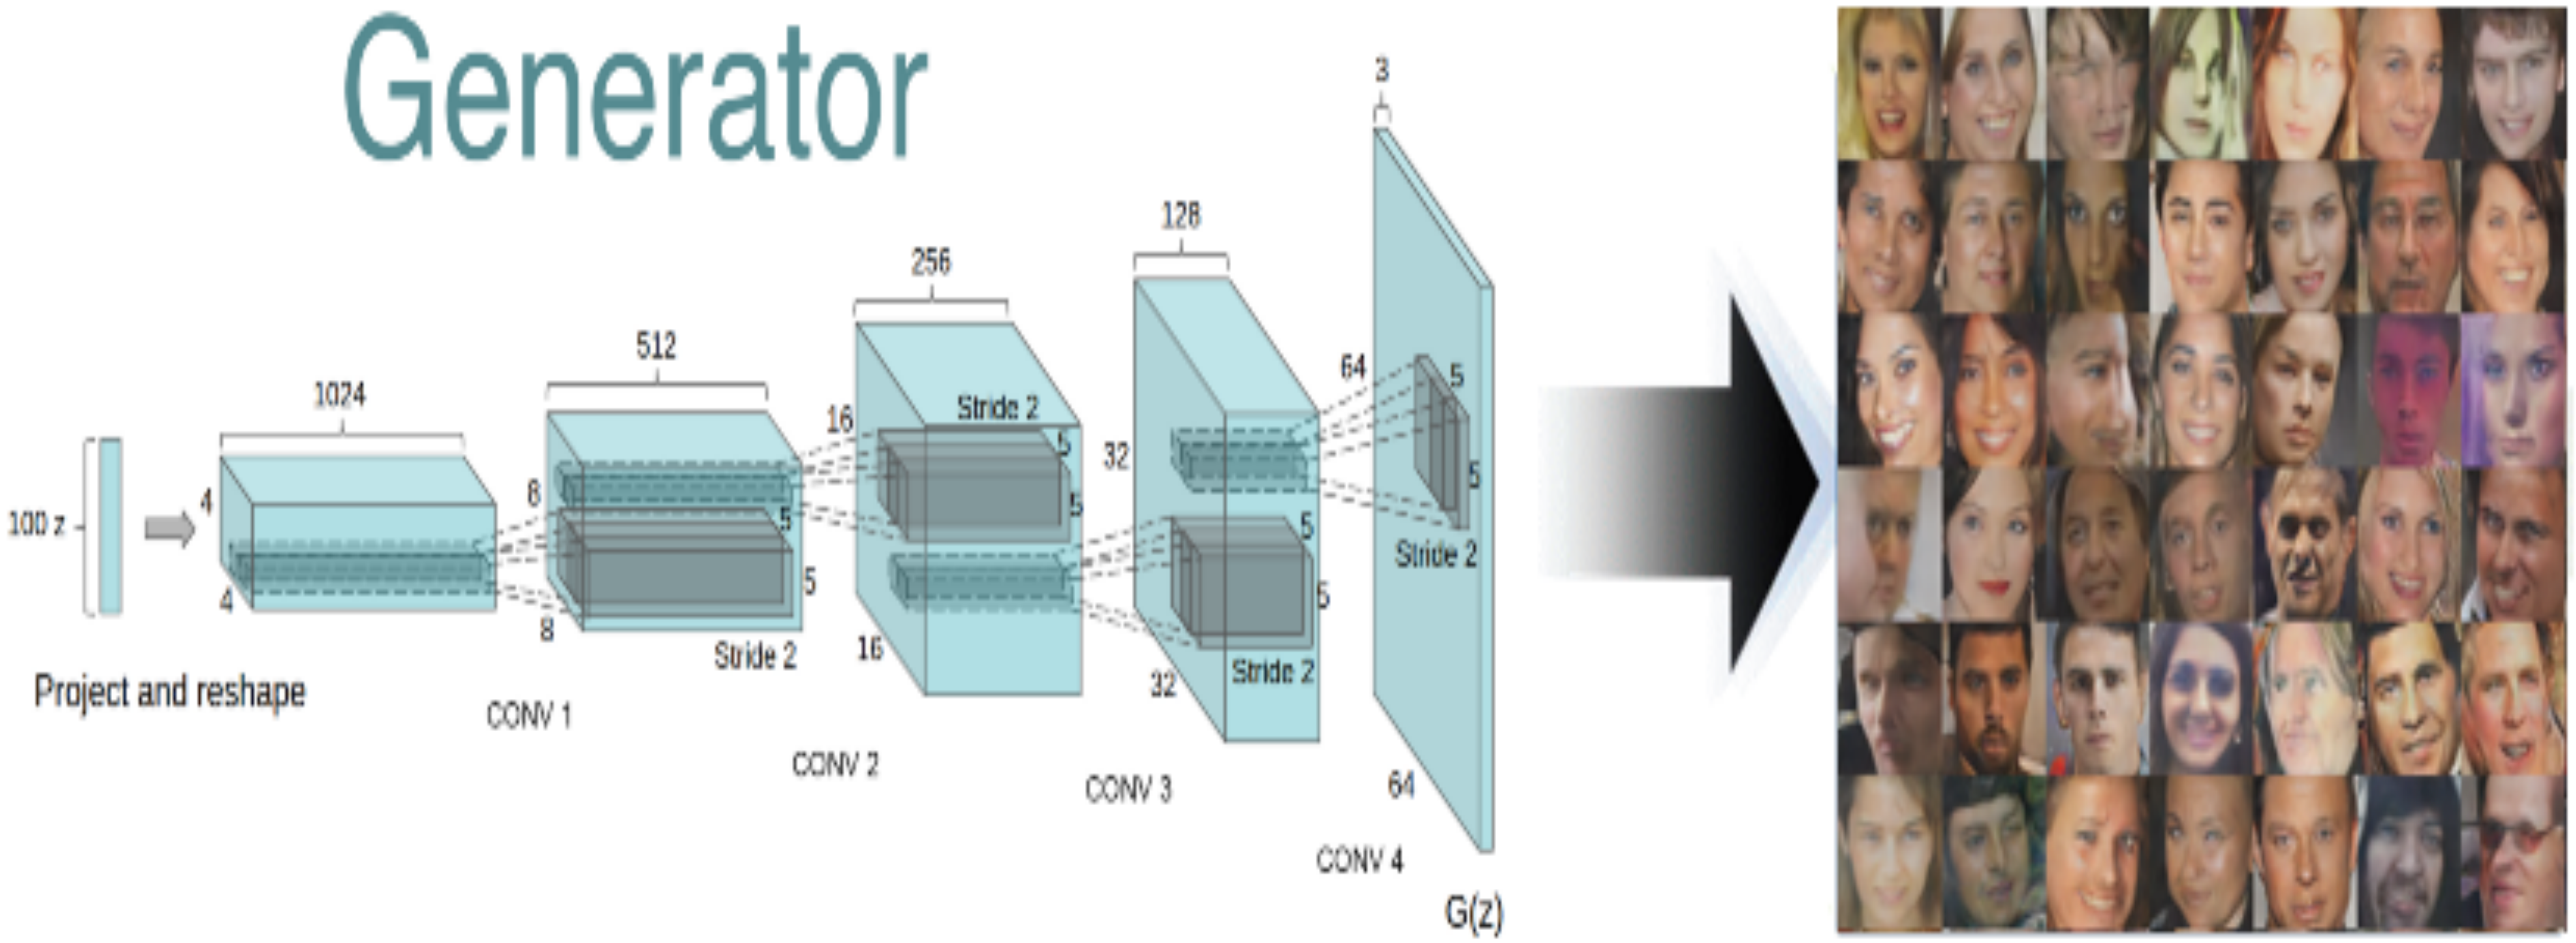

- Given samples in $\mathbb{R}^n$ from distribution π_0, learn π_0.
- Model π_0 as $G(Z)$, where $Z \sim \mathcal{N}(0, I_{k\times k})$ and $G : \mathbb{R}^k \to \mathbb{R}^n$.

Visualization of a generative model and its latent space

Inference

$x \sim \hat{p}_X$

$z = f(x)$

Generation

$z \sim p_Z$

$x = f^{-1}(z)$

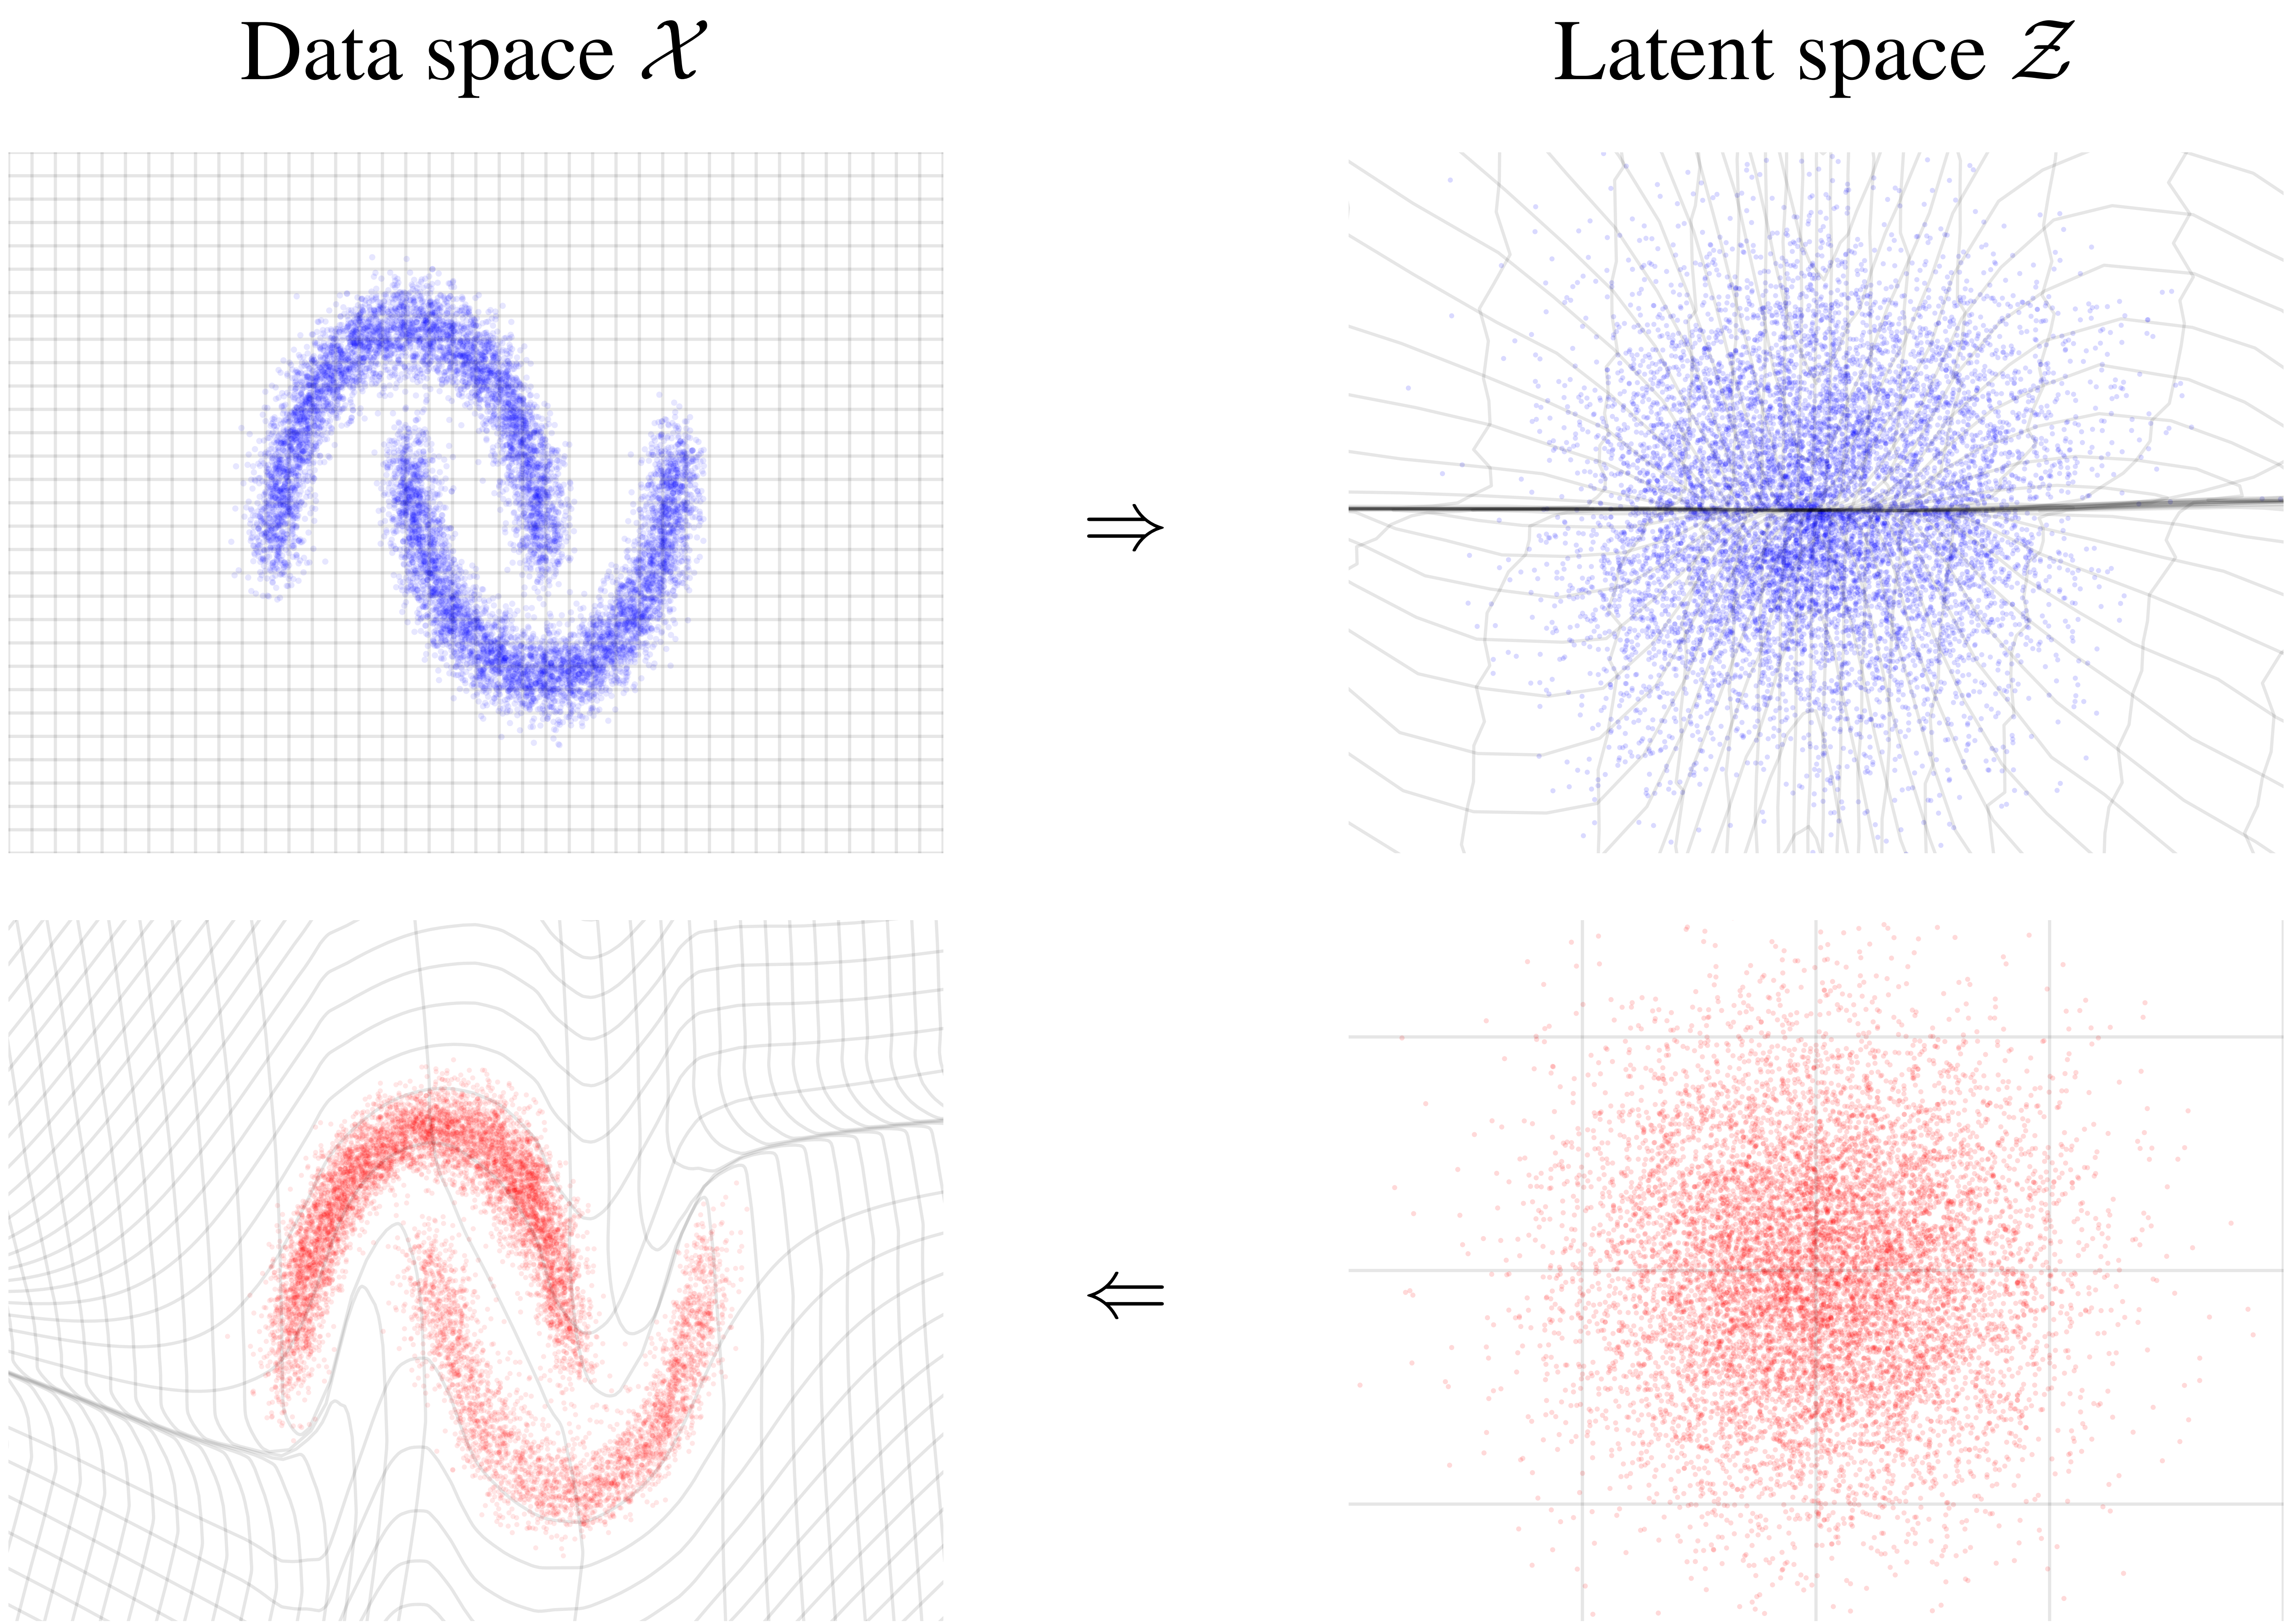

Dinh et al., 2017

How are generative models used in inverse problems?

1. Train generative model to output signal class:

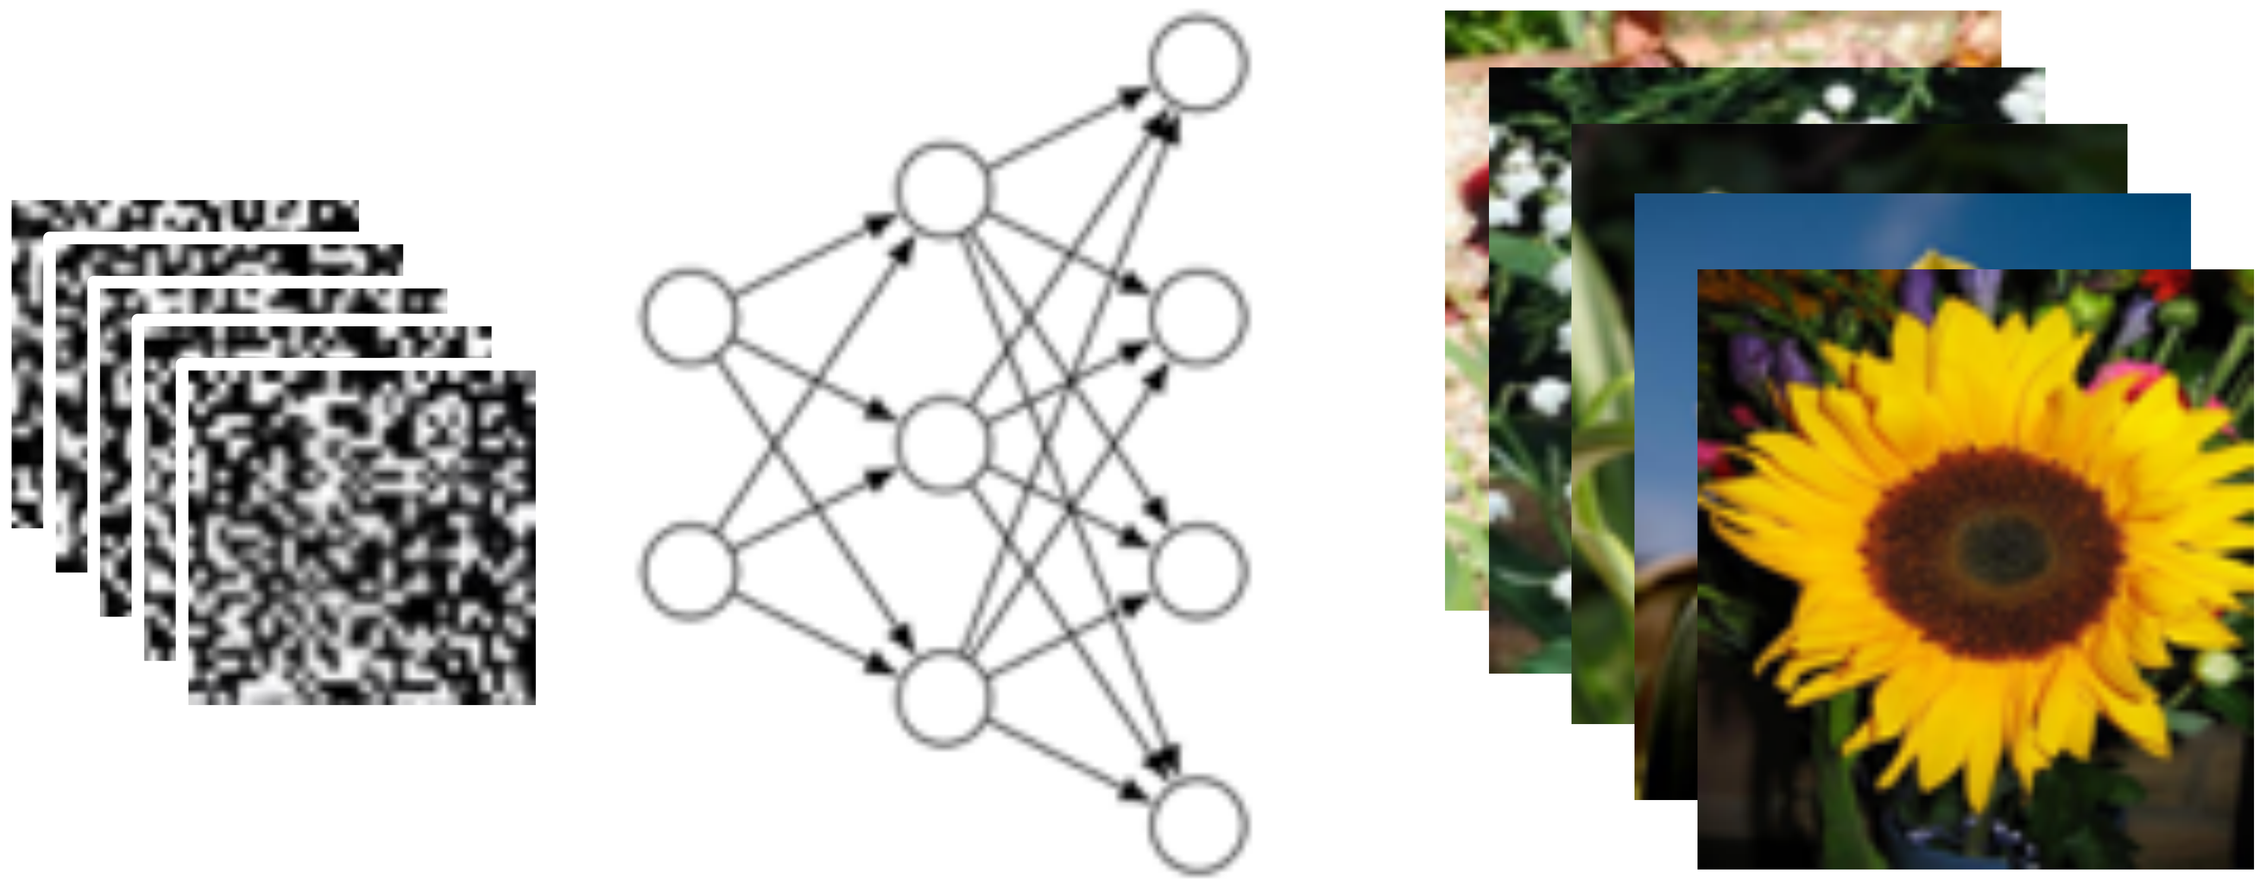

2. Directly optimize over range of generative model via empirical risk:

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

Compressed Sensing with Generative Models

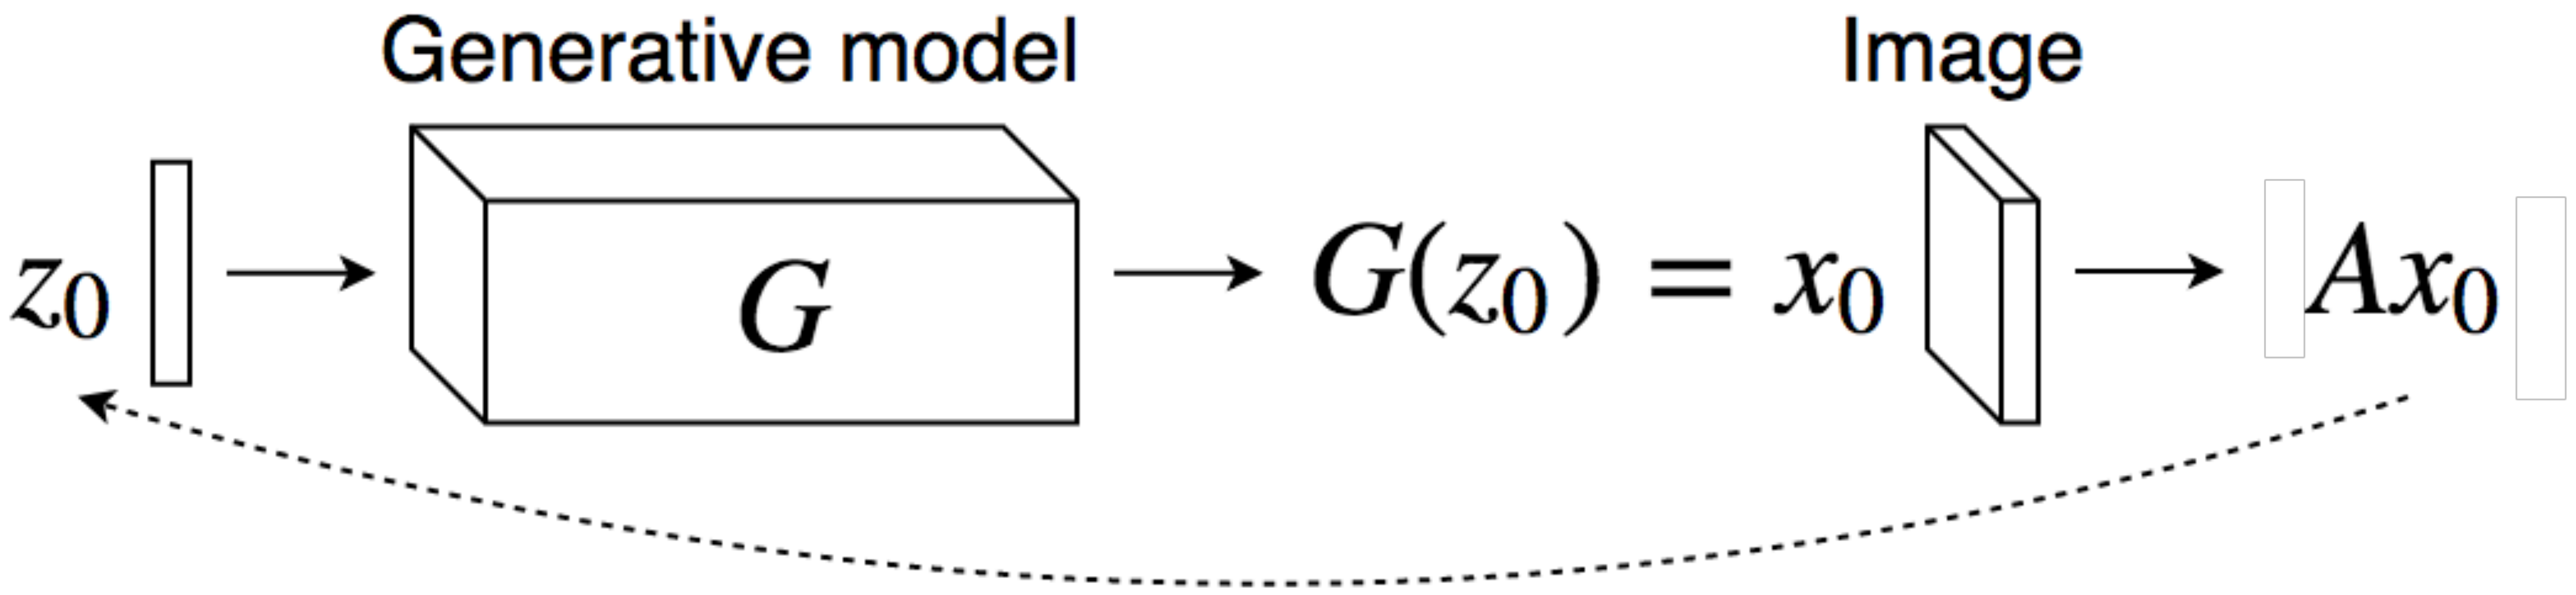

$$\min_{z \in \mathbb{R}^k} \left\| AG(z) - Ax_0 \right\|^2$$

Bora, Jalal, Price, Dimakis

Our formulation: Phase Retrieval with Generative Models

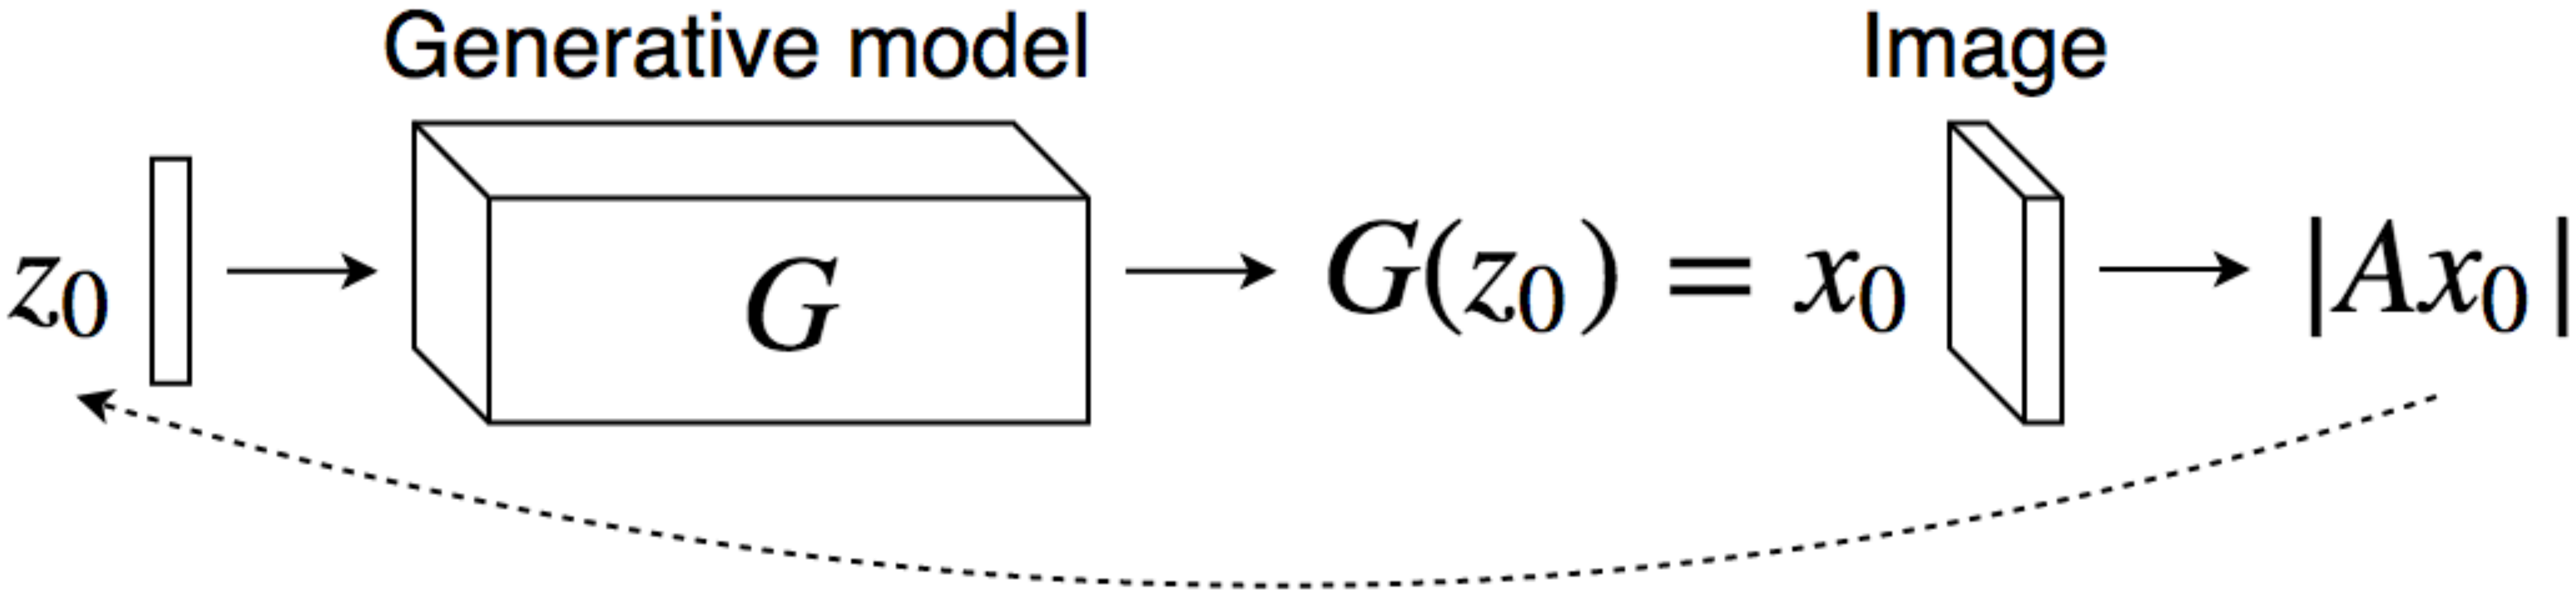

$$\min_{z \in \mathbb{R}^k} \Big\| |AG(z)| - |Ax_0| \Big\|^2$$

Geometric picture of signal recovery with a low-dimensional generative prior

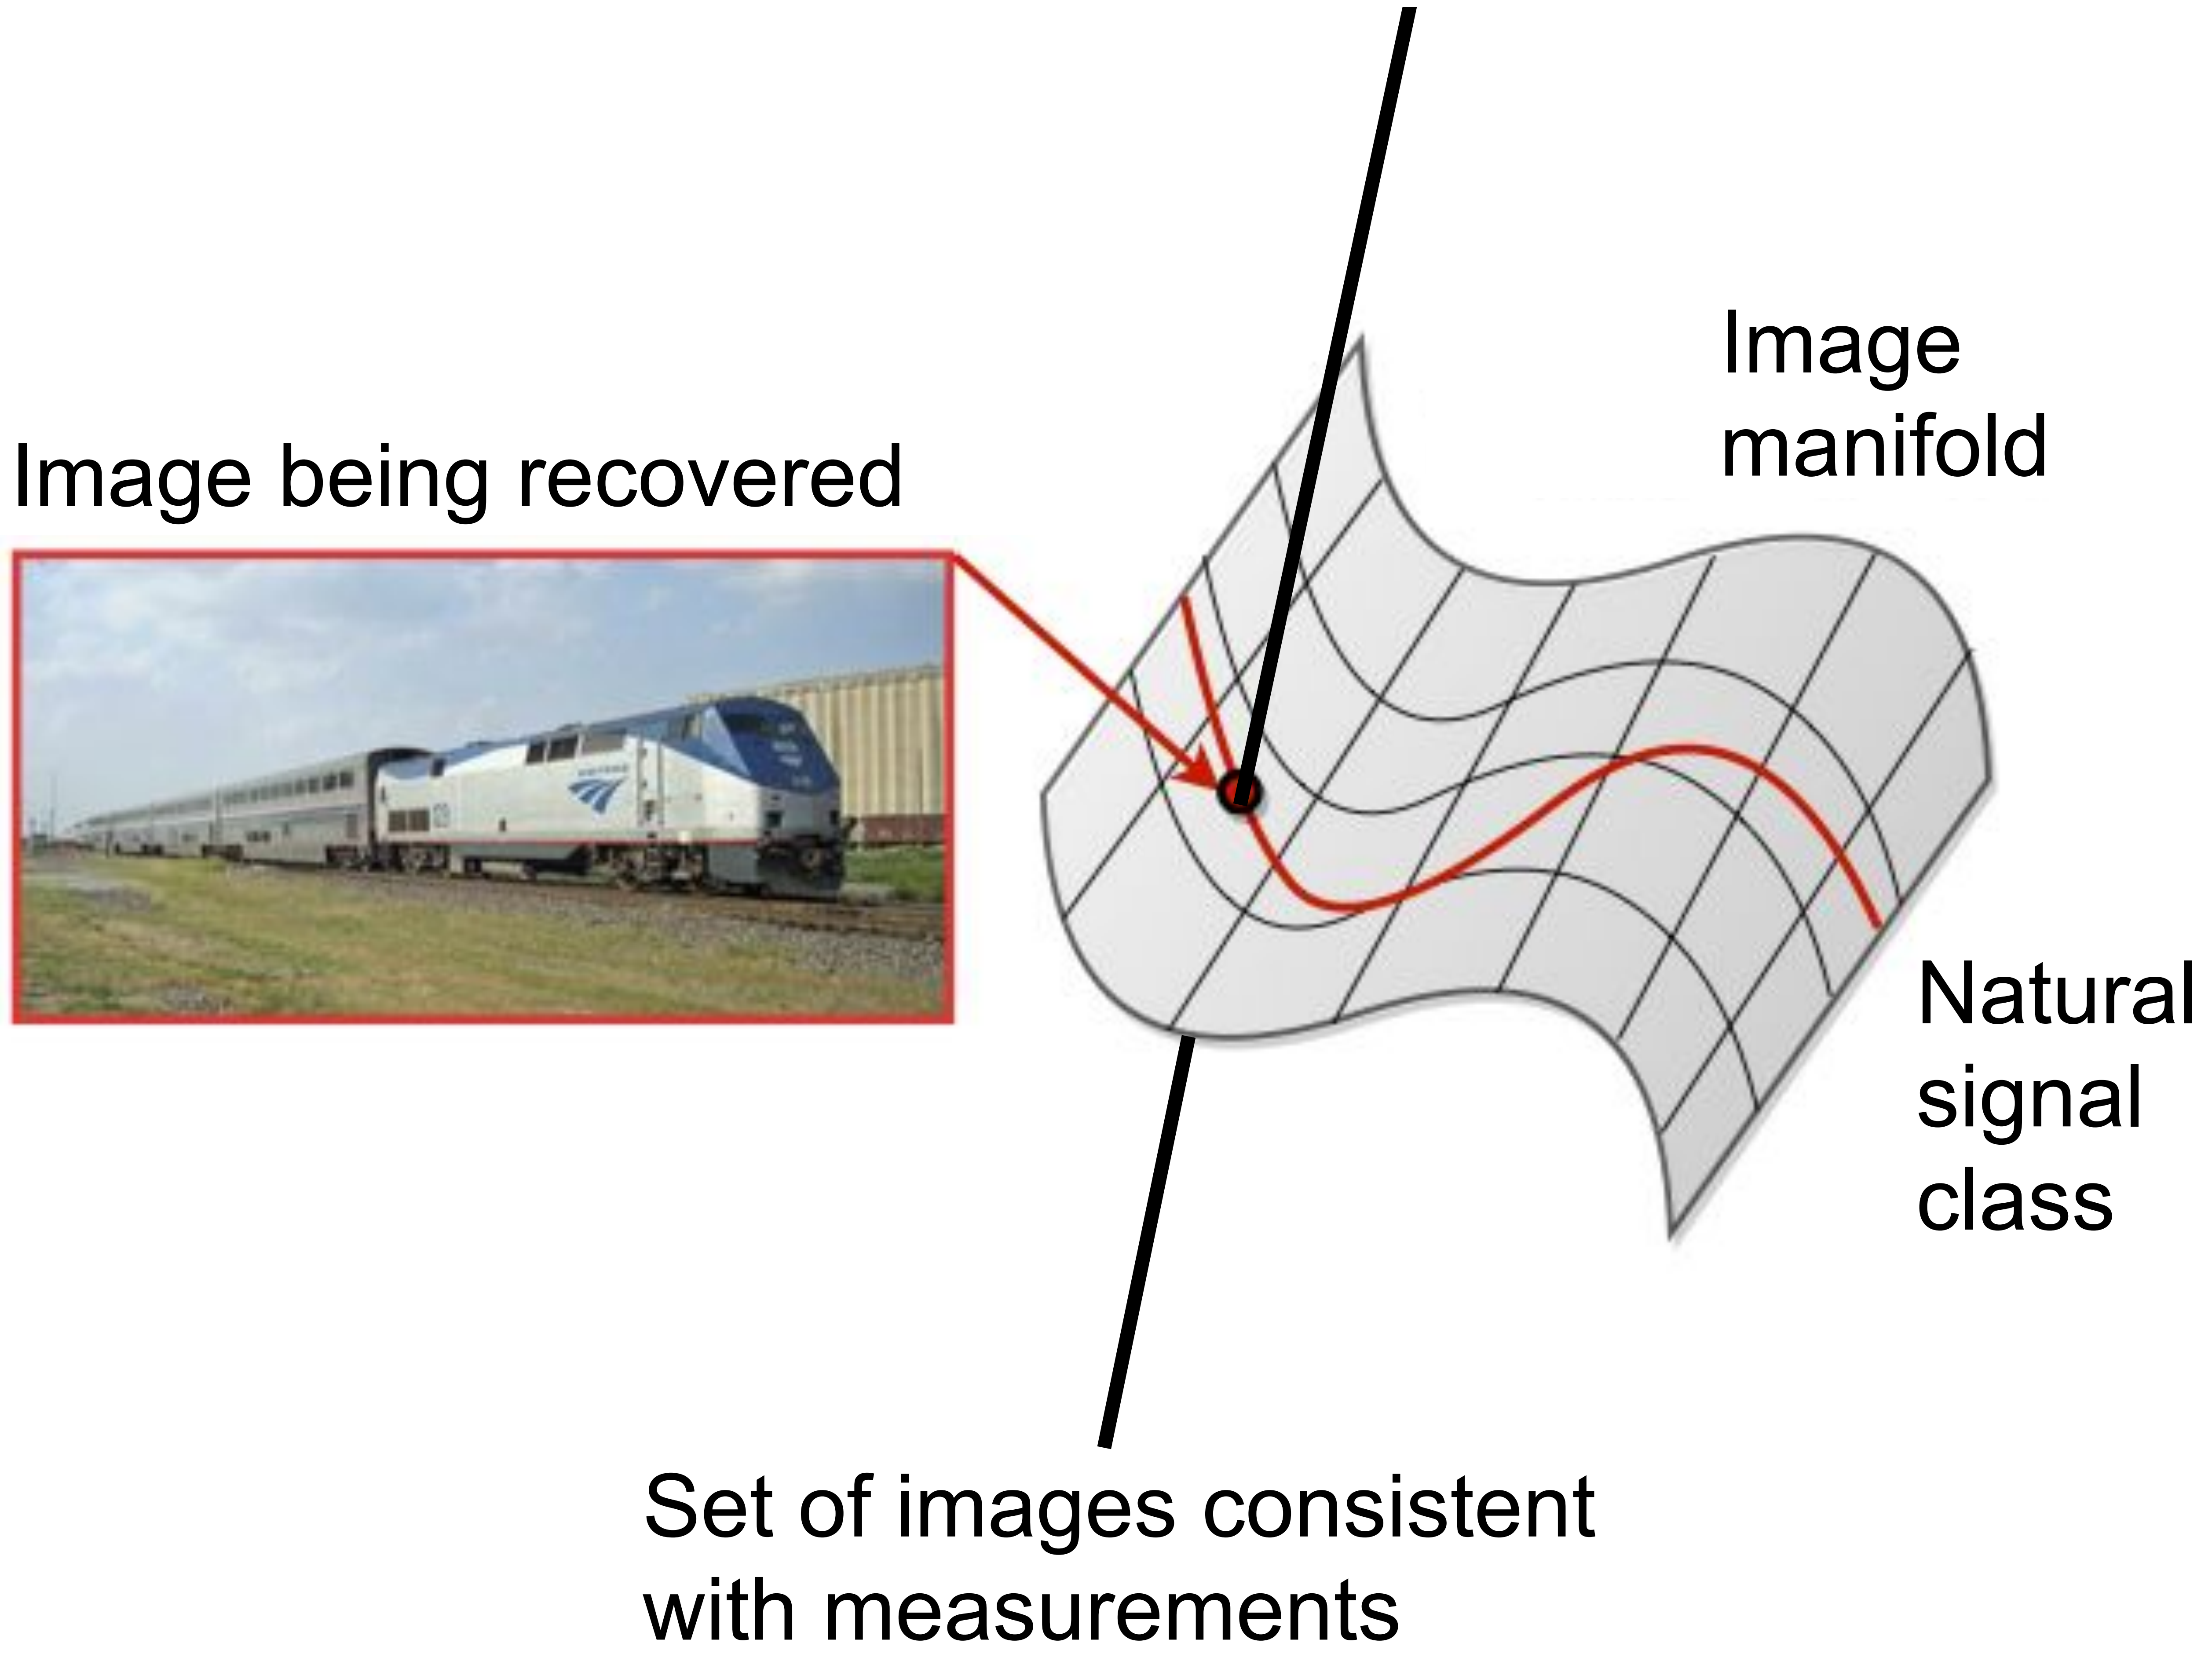

Random generative priors allow rigorous recovery guarantees

Let: $\mathcal{G} : \mathbb{R}^k \to \mathbb{R}^n$

$\mathcal{G}(z) = \text{relu}(W_d \ldots \text{relu}(W_2 \text{relu}(W_1 z)) \ldots)$

Given: $W_i \in \mathbb{R}^{n_i \times n_{i-1}}, A \in \mathbb{R}^{m \times n}, y := A\mathcal{G}(z_0) \in \mathbb{R}^m$

Find: x_0

- **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 phase retrieval from generic measurements is possible at optimal sample complexity

1. $\#$ measurements $= \Omega(k)$, up to log factors
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 in latent space outside of two small neighborhoods of the minimizer and a negative multiple thereof, with high probability.

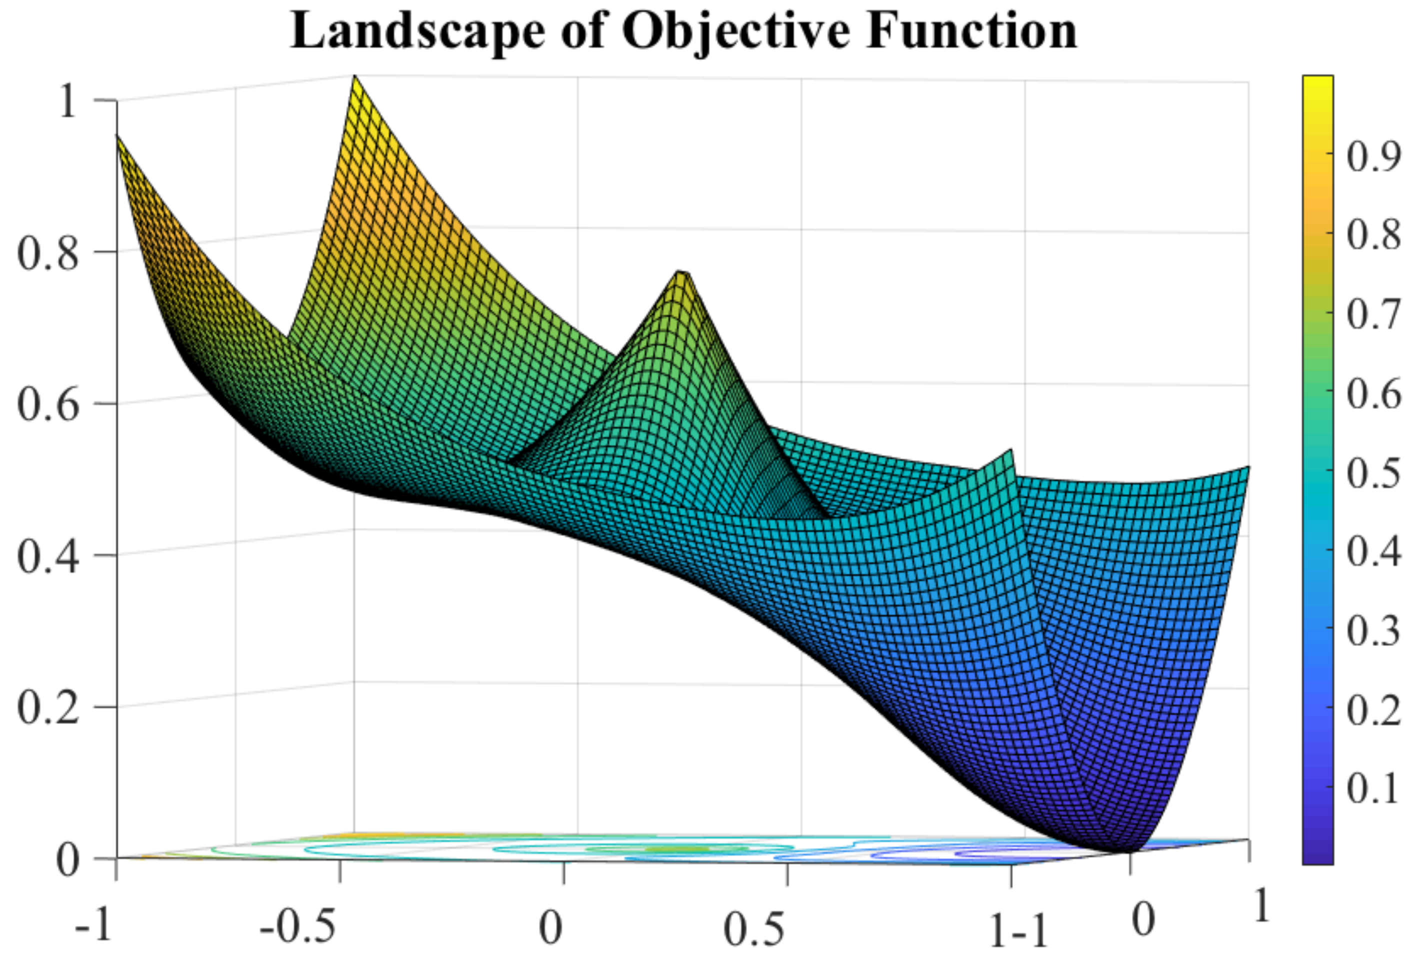

Proof requires concentration of discontinuous matrix-valued random functions

Lemma: Fix ϵ. Let $W \in \mathbb{R}^{n \times k}$ have i.i.d. $\mathcal{N}(0, 1/n)$ entries. If $n > ck \log k$, then with probability at least $1 - 8ne^{-\gamma k}$, we have for all $x, y \neq 0 \in \mathbb{R}^k$,

$$\left\| \frac{1}{n} \sum_{i=1}^{n} 1_{w_i \cdot x > 0} 1_{w_i \cdot y > 0} \cdot w_i w_i^T - \mathbb{E}[\cdots] \right\| \leq \epsilon$$

The constants depend polynomially on ϵ.

Compressive Phrase Retrieval on MNIST

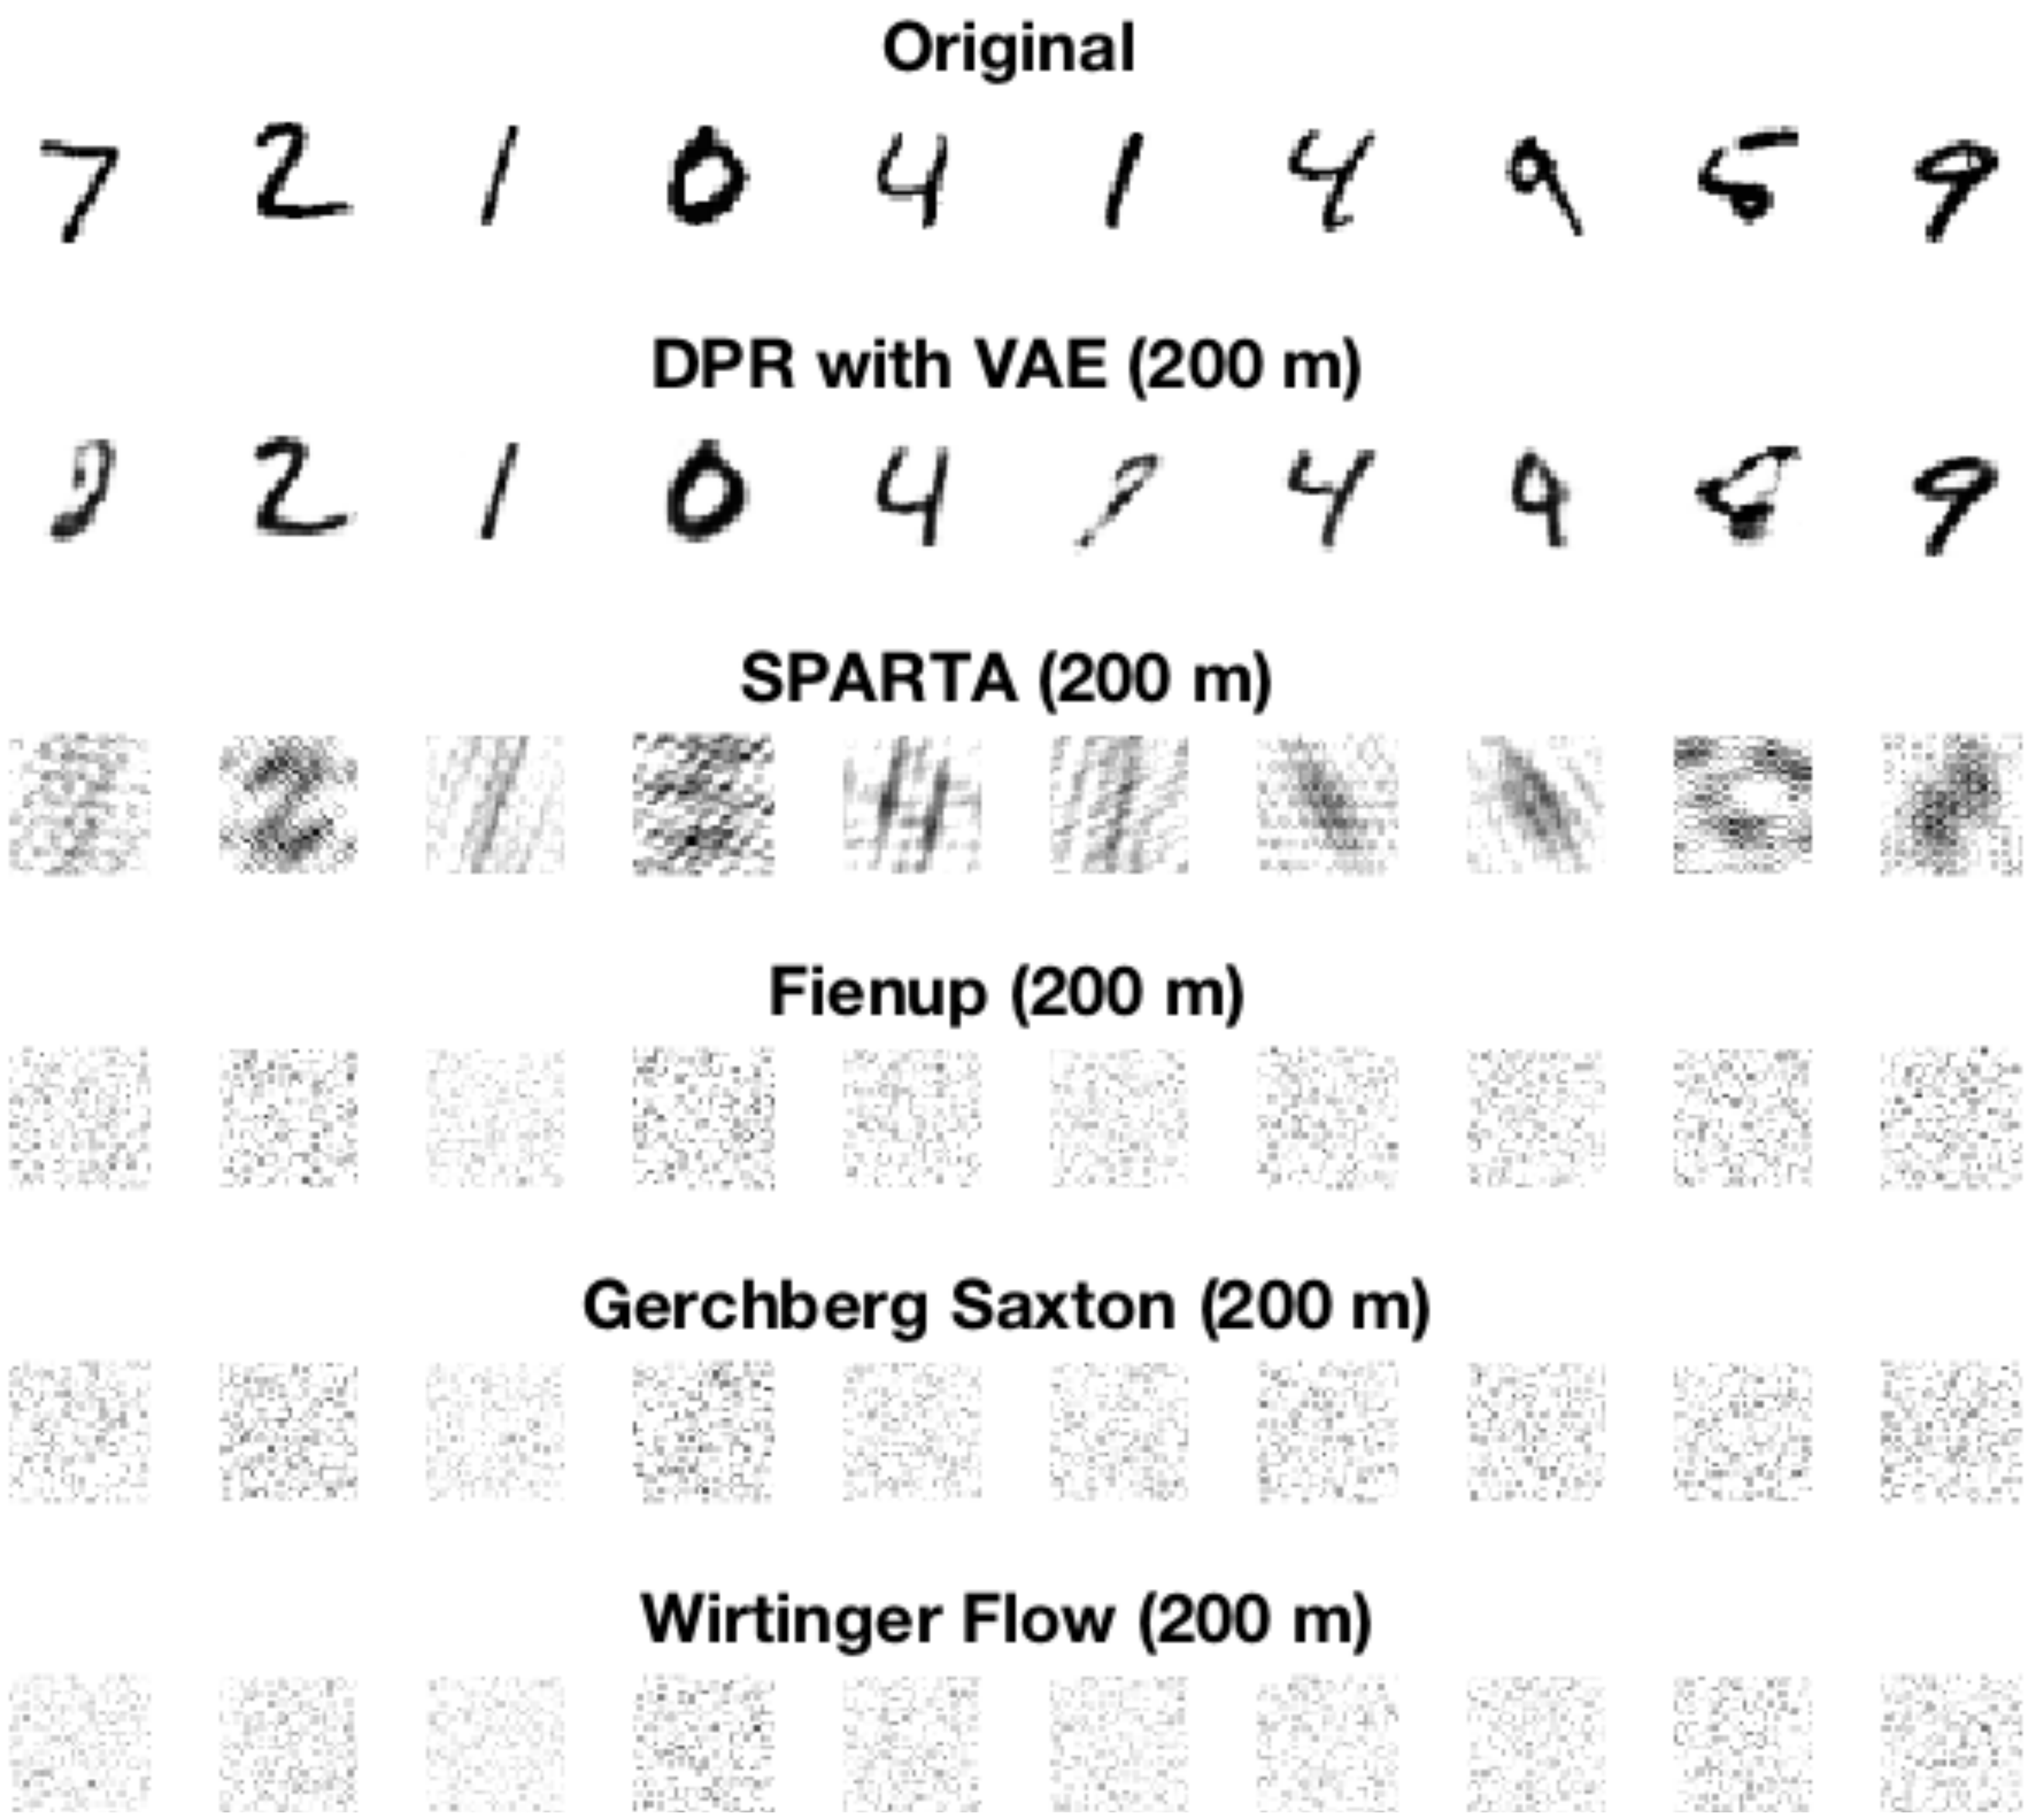

Deep phase retrieval can outperform sparse phase retrieval in the low measurement regime

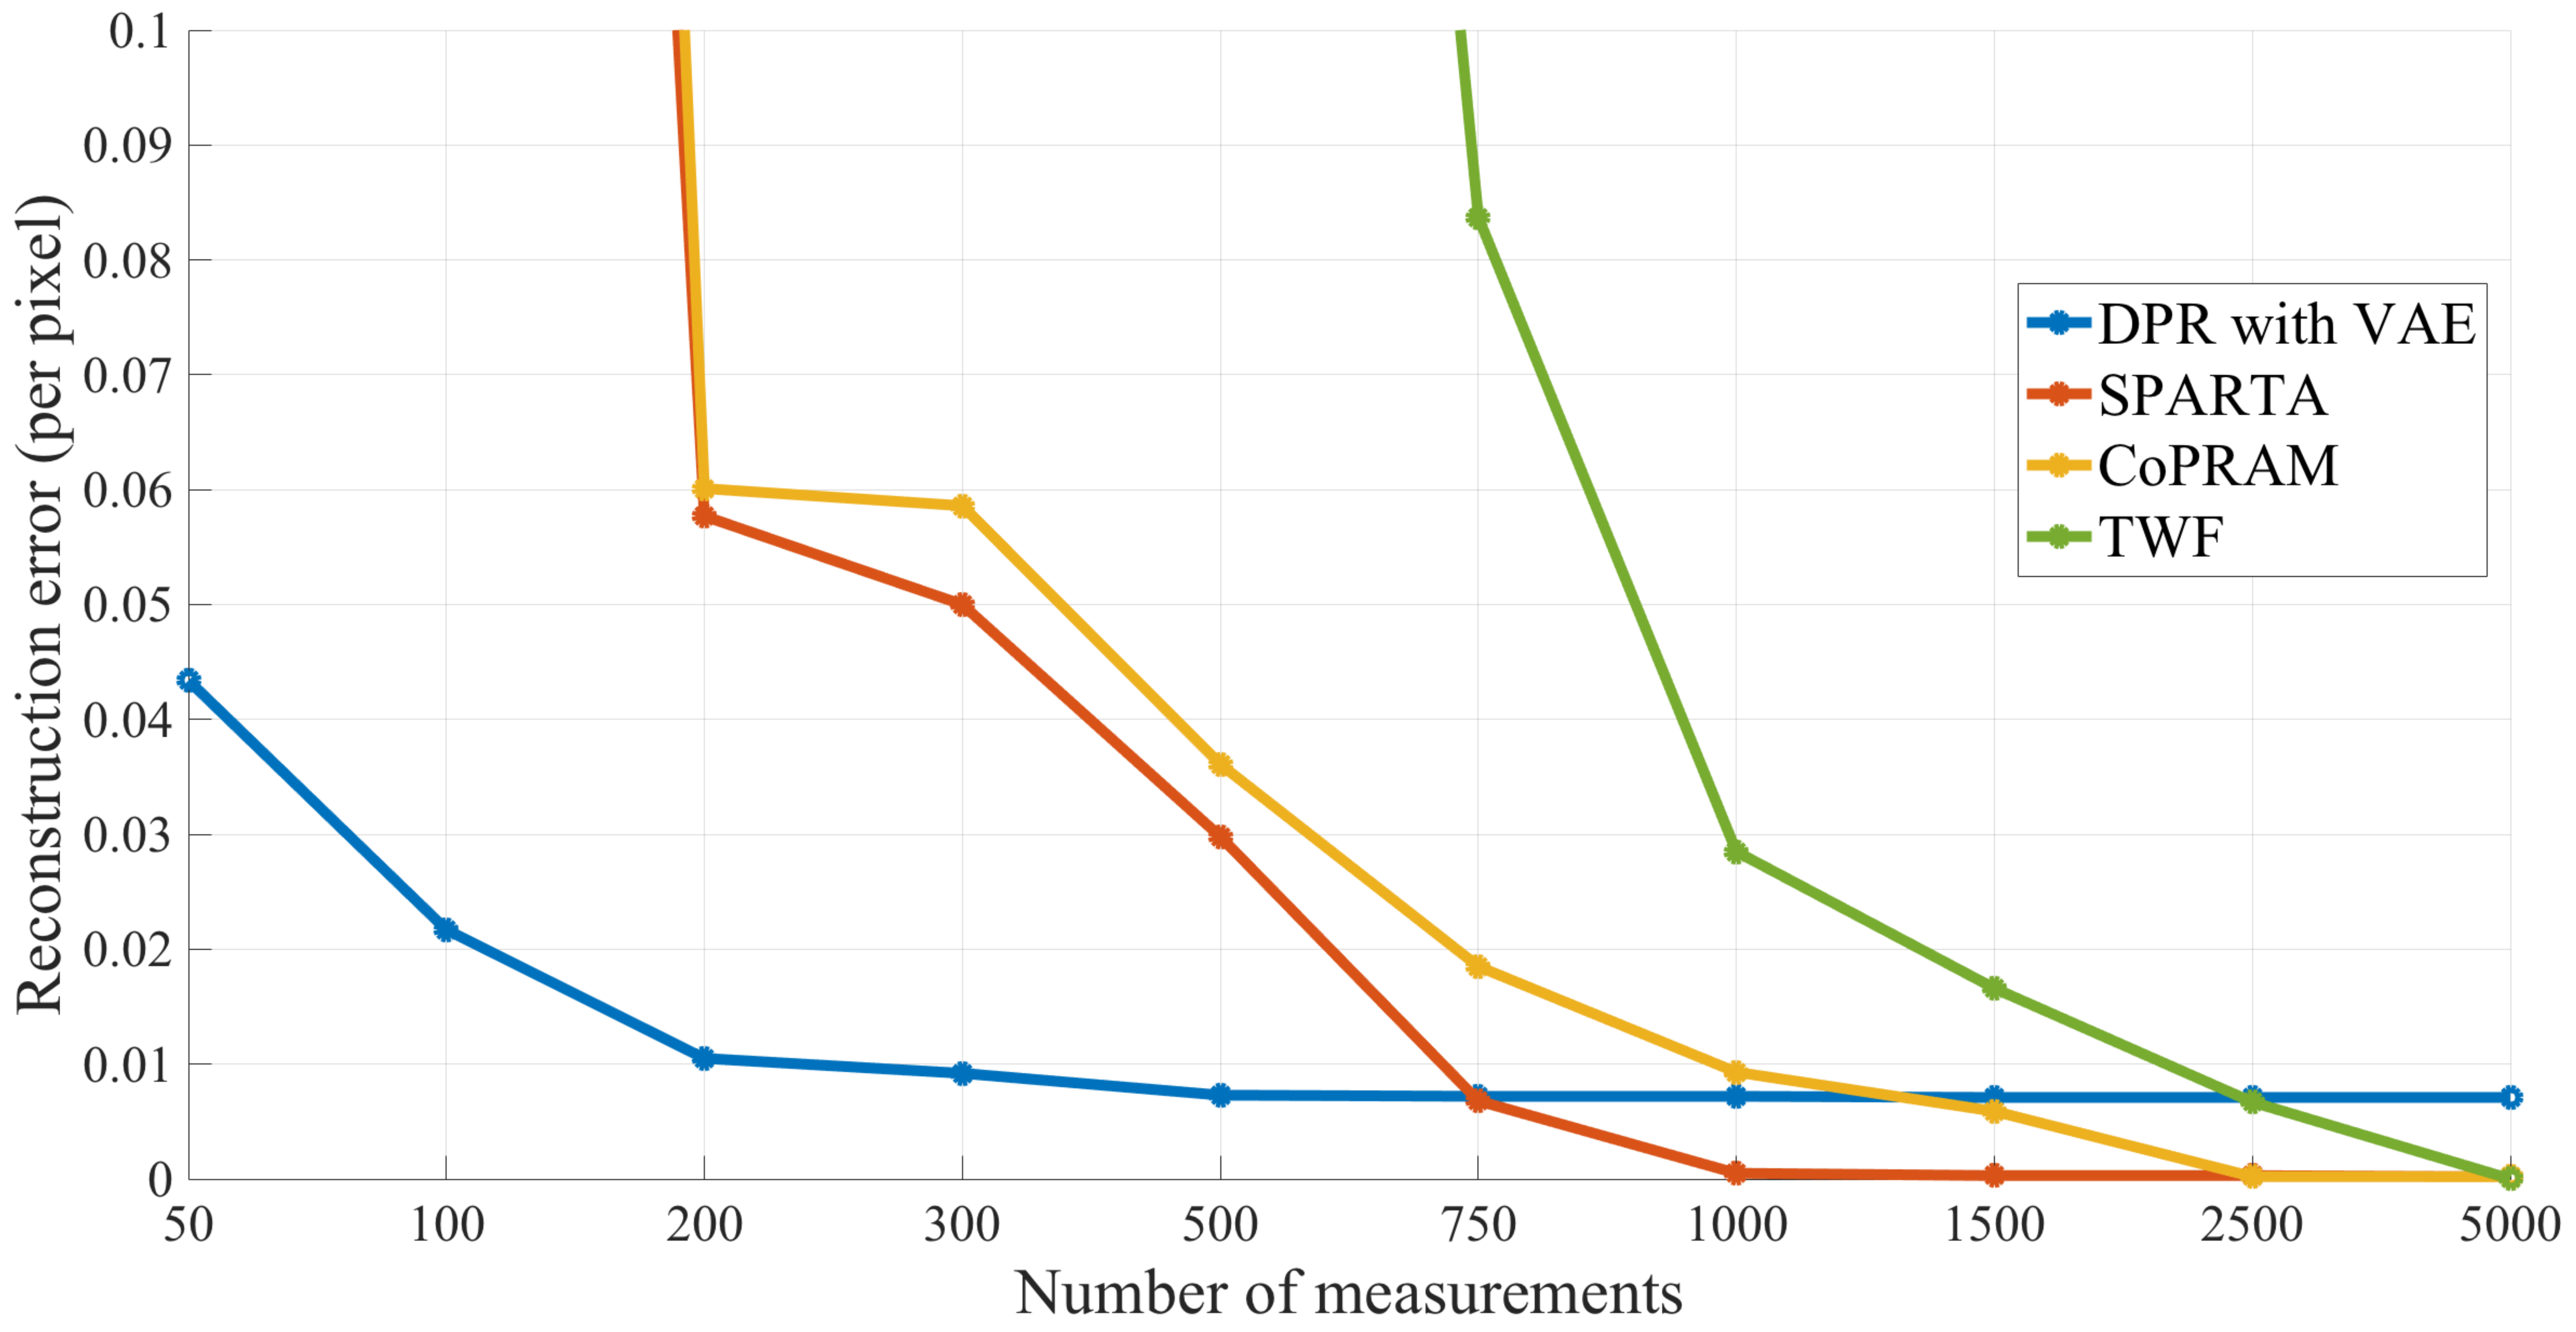

Is there a catch?

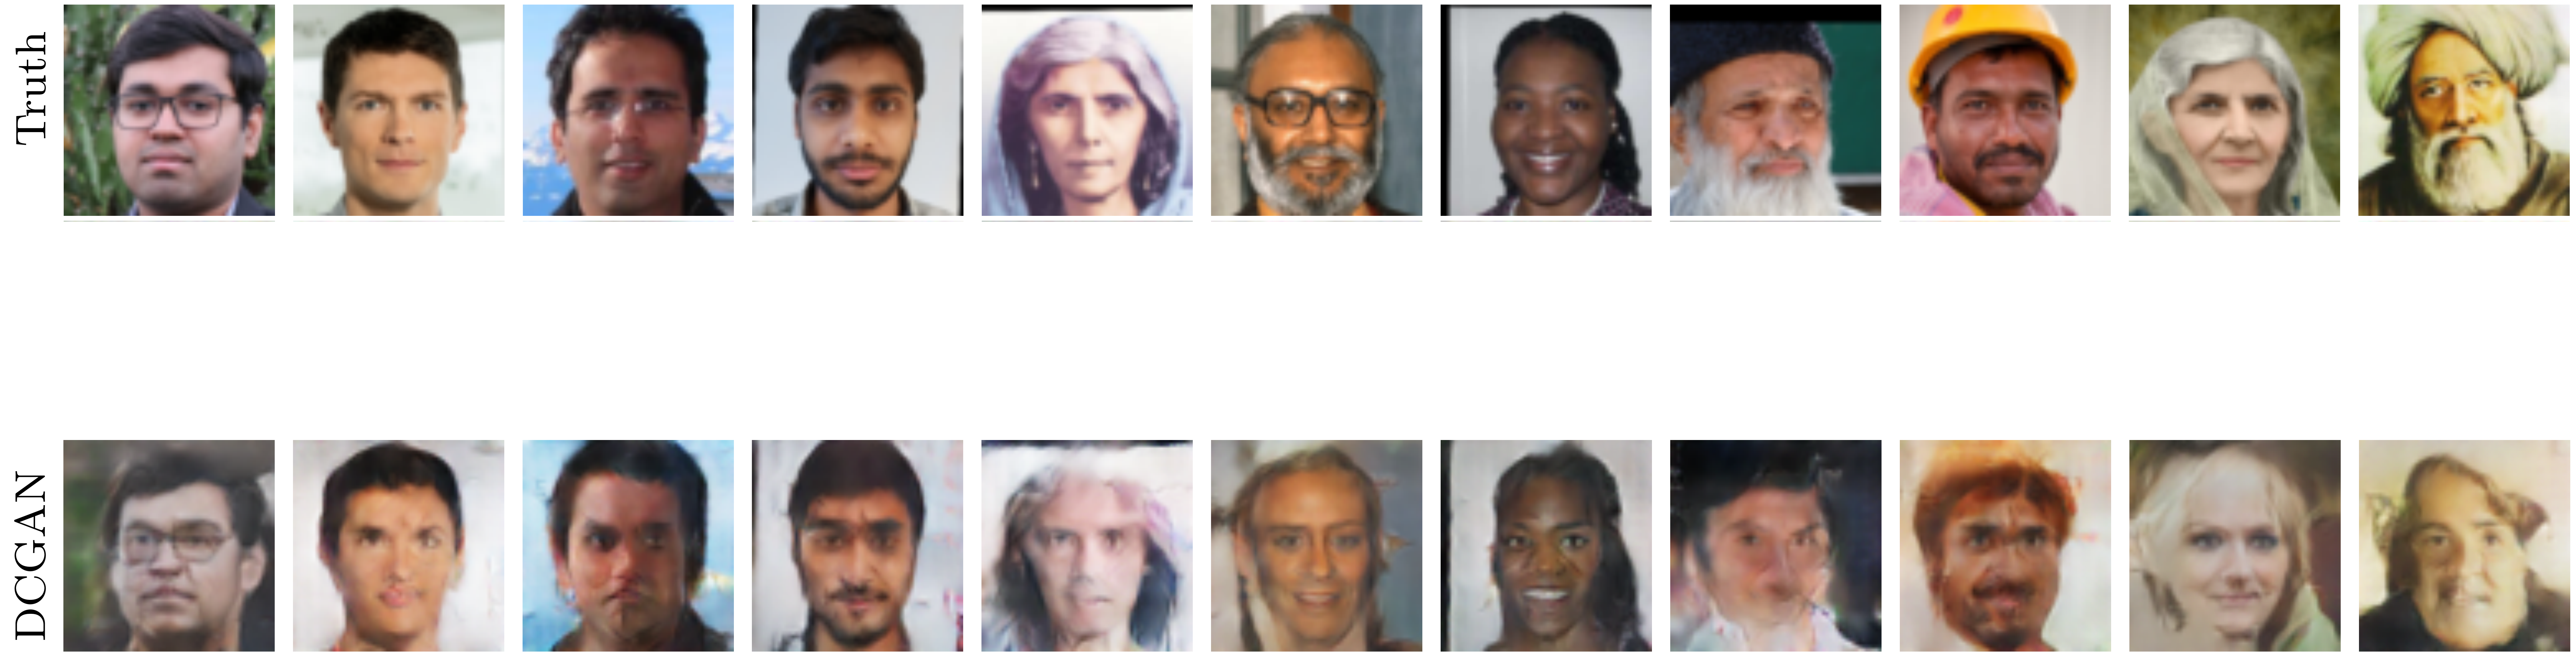

GANs can have significant representation error

GANs can have very poor performance far off distribution

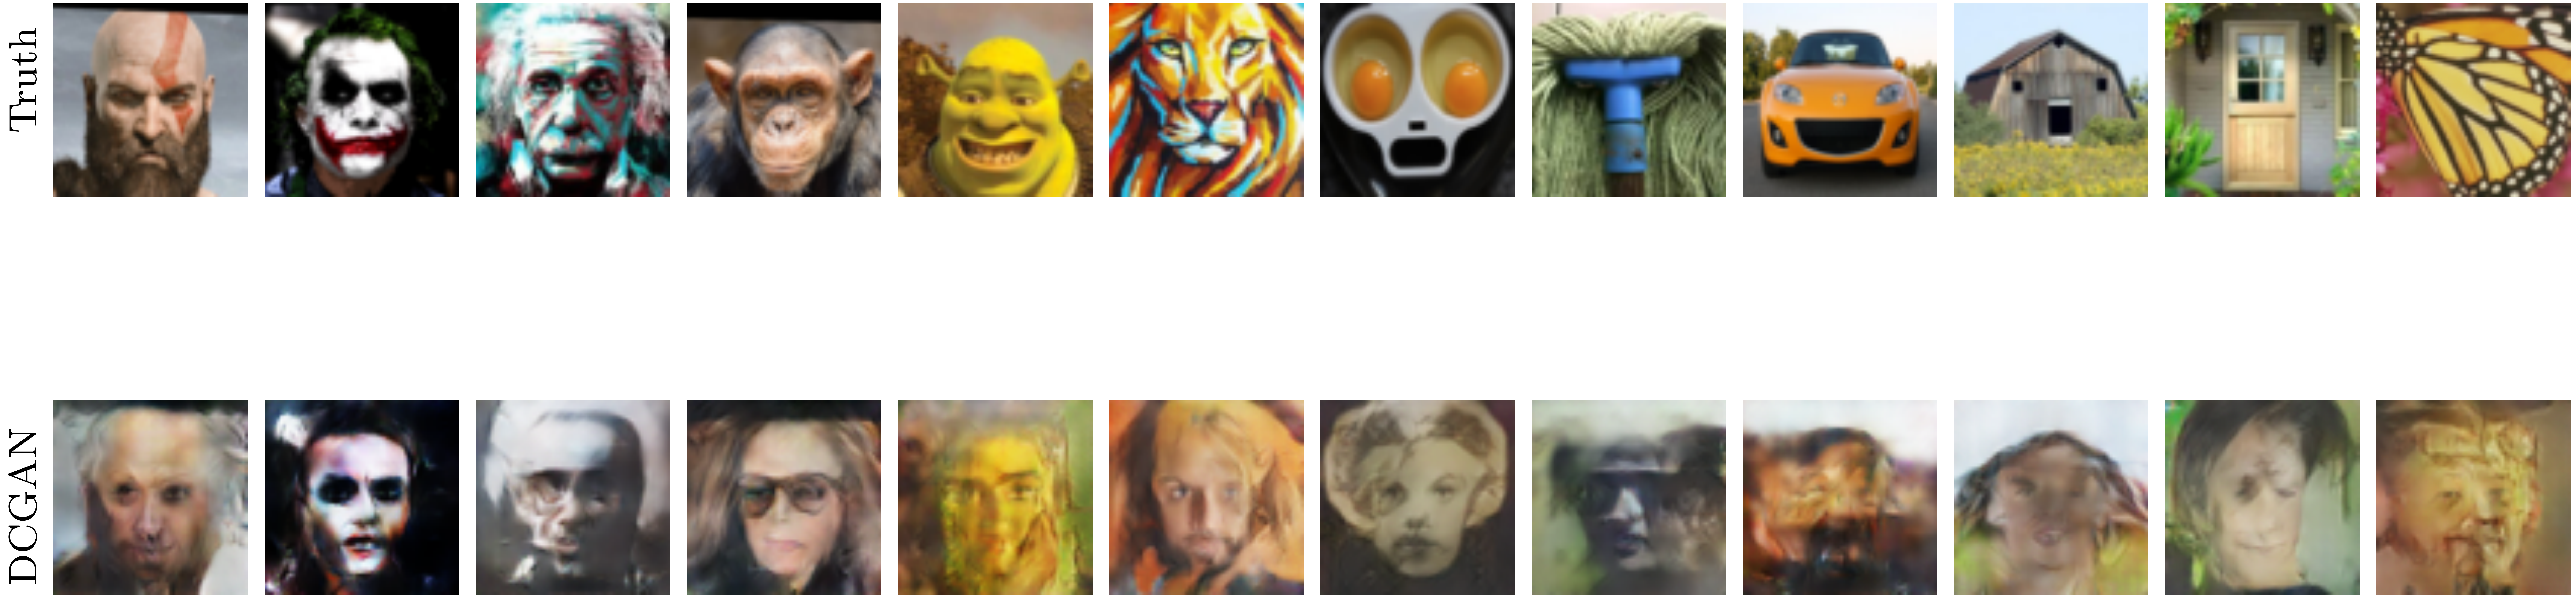

Signal Recovery Under Generative Priors

- Generative priors can be optimally exploited for some nonlinear problems
- Generative priors could provide tighter representations of natural images
- Low dim. nonconvex optimization replaces high dim. convex optimization
- Generative priors may outperform sparsity priors for a variety of problems

References

Deep Compressive Sensing

- Bora et al. 2017 - *ICML*
- Hand and Voroninski 2019 - *IEEE Trans IT*
- Huang et al. 2021 - *J. Fourier Anal. App.*

Other Inverse Problems

- Denoising:
  Heckel et al. 2020 - *Information and Inference*
- Phase Retrieval:
  Hand, Leong, and Voroninski 2018 - *NeurIPS*
- Spiked Matrix Recovery:
  Aubin et al. 2020 - *IEEE Trans IT*
  Cocola et al. 2020 - *NeurIPS, Entropy*
- Blind Demodulation:
  Hand and Joshi 2018 - *NeurIPS*

Generalization of Assumptions

- Convolutional Generators
  Ma, Ayaz, and Karaman 2018 - *NeurIPS*
- Better Expansivity Condition
  Daskalakis et al. 2020 - *NeurIPS*

Normalizing Flow Priors

- Asim et al. 2020 - *ICML*

Review Articles

- Lucas et al. 2018 - *IEEE Sig. Proc. Mag.*
- Ongie et al. 2020 - *IEEE JSAIT*

Workshops

- 2019, 2020, 2021 NeurIPS Workshops