Compressed Sensing using Generative Models

Ashish Bora, Ajil Jalal, Eric Price, Alex Dimakis

TEXAS
The University of Texas at Austin

Compressed Sensing - Introduction

- Goal: estimate a signal $x^* \in \mathbb{R}^n$ from a linear system $y = Ax^* + \eta$.
- Applications to MRIs, IR imaging, oil exploration, etc.
- Let $A \in \mathbb{R}^{m \times n}$. How many measurements m are needed? Naively: $m \geq n$ or else underdetermined; multiple x possible.
- But not all x are plausible/natural:

5 MB

36 MB

- Ideally: m = (information in signal) / (new info. per measurement)
  - Signal is "natural" $\implies$ information in image is small.
  - Measurements "incoherent" $\implies$ most info new.
- Three questions:
  1. How to define "natural"?
  2. For that definition, how to choose A?
  3. Given measurements, what algorithm for recovery?

Standard Compressed Sensing

- "Natural" represented as k-sparsity in well-chosen basis.
- Typical target: find $\widehat{x}$ with

$$\|\widehat{x} - x^*\|_2 \lesssim \min_{k\text{-sparse } x} \|x - x^*\|_2 + \|\eta\|_2$$

- Possible when A is i.i.d. Gaussian with $m = O(k \log n)$ rows [Candès-Romber-Tao '06]. Can be efficiently computed using LASSO.
- But is sparsity really the right structure? Or can we study the millions of MRIs we take every year to learn a better structure?

Our approach

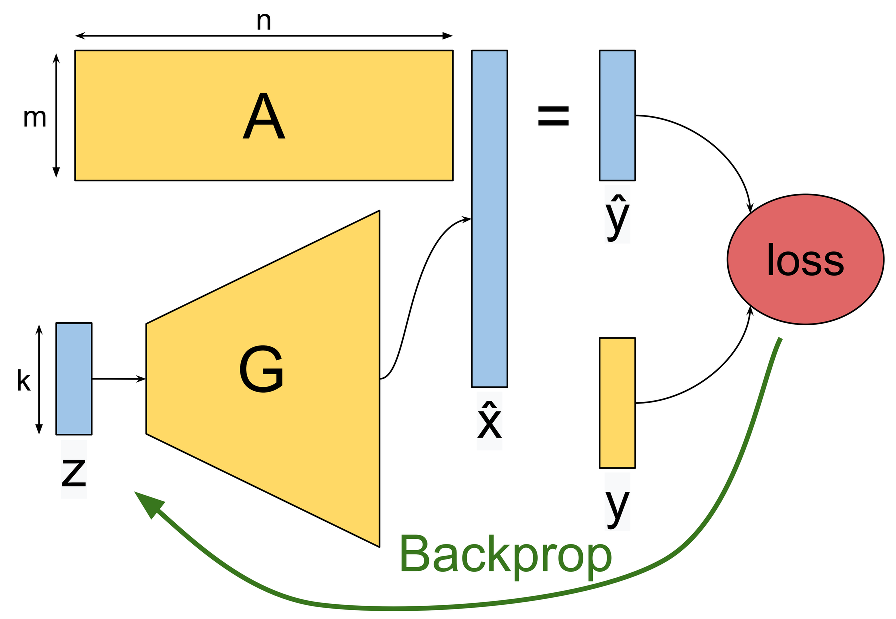

- $loss(z) = \|AG(z) - y\|_2^2$
- Get gradients of loss with respect to z by backpropagation.
- Optimize wrt z iteratively using gradient descent to get $\widehat{z}$
- Finally, output $\widehat{x} = G(\widehat{z})$

Highly nonconvex, so no proof of convergence. But approximate minimization implies approximate recovery, and we can check the final error.

Our result

We replace sparsity with a more powerful notion: **appearing in the range of a generator** $G : \mathbb{R}^k \to \mathbb{R}^n$. The target becomes:

$$\|\widehat{x} - x^*\|_2 \lesssim \min_{x = G(z)} \|x - x^*\|_2 + \|\eta\|_2.$$

We show this is possible for random Gaussian A and

- $m = O(kd \log n)$ for depth-d ReLU-based neural networks.
- $m \approx O(k \log(L/\delta))$ for general L-Lipschitz functions G, with additive error δ.

Well-trained GAN or VAE generative models can represent images with much smaller k than sparse representations.

Key Property: a new eigenvalue condition

- A satisfies the normal REC, if for all approximately sparse vectors x, $\|Ax\|_2 \geq \gamma \|x\|_2$.
- It satisfies our "Set-Restricted Eigenvalue Condition" (S-REC) for $S \subseteq \mathbb{R}^n$ if, for all $x_1, x_2 \in S$ we have

$$\|A(x_1 - x_2)\|_2 \geq \gamma \|x_1 - x_2\|_2 - \delta.$$

- With enough measurements, A satisfies the S-REC for $S = \text{range}(G)$, which implies recovery.

Qualitative Results

Reconstruction:

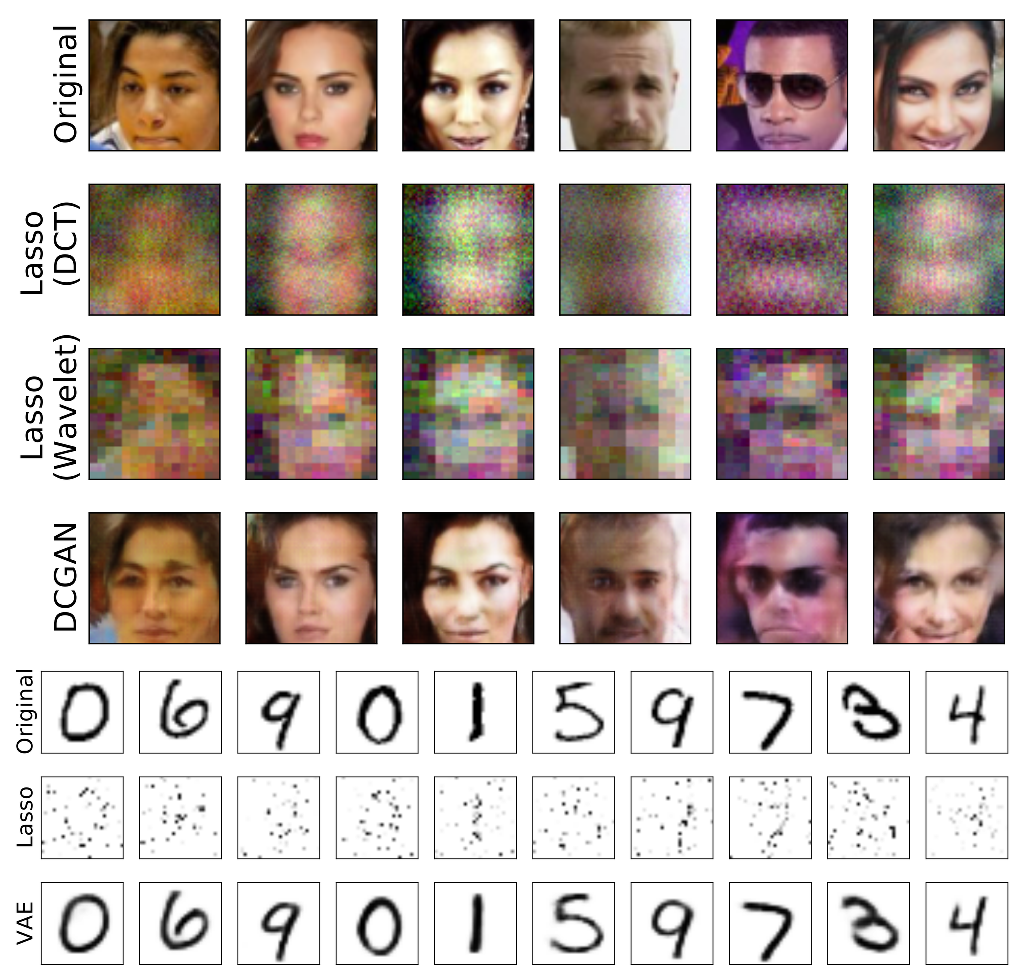

Super-resolution:

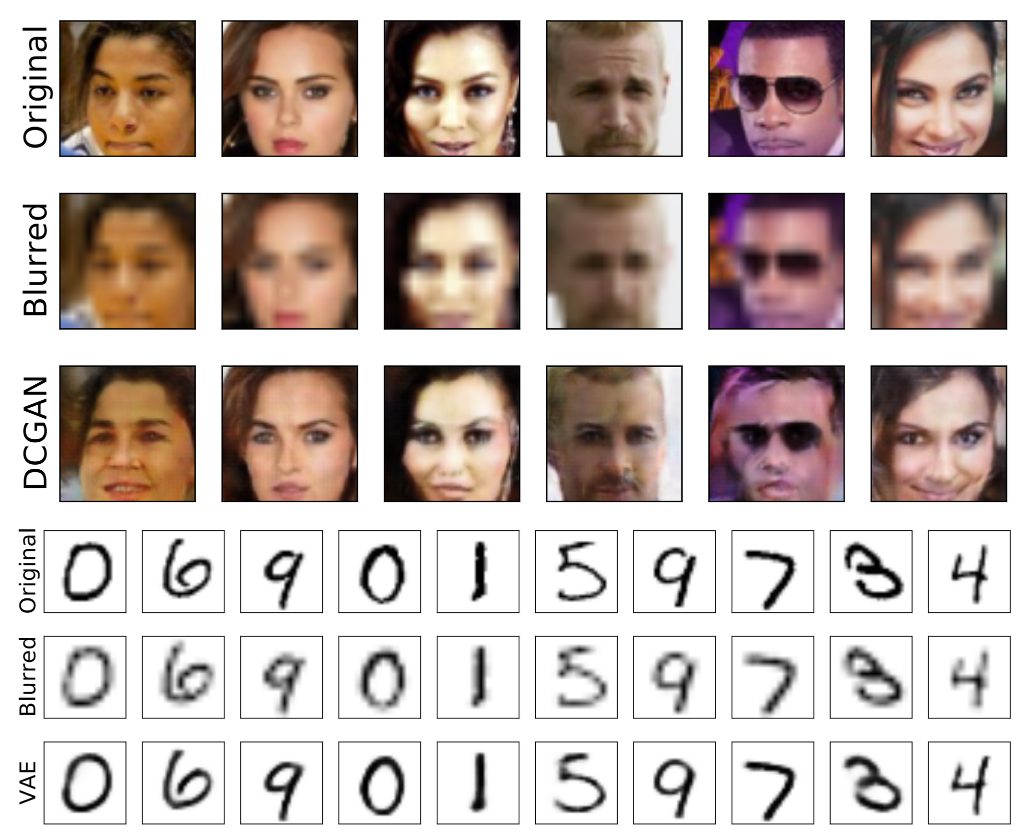

Inpainting:

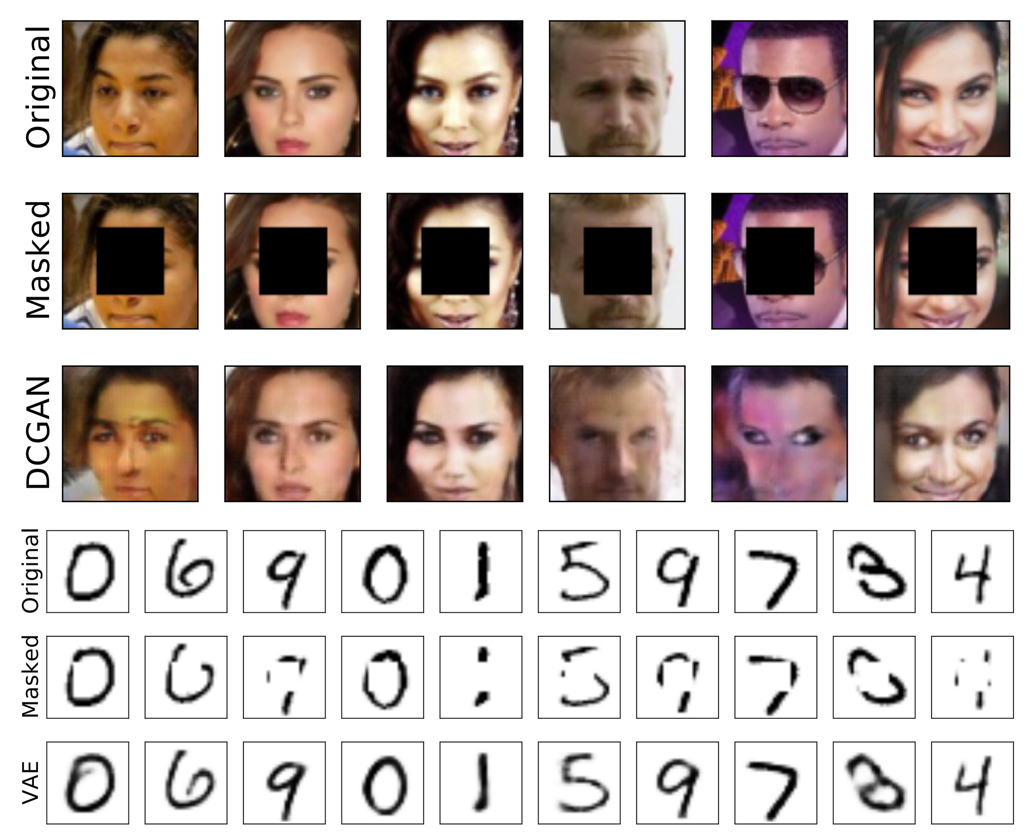

Quantitative Results

MNIST with VAE

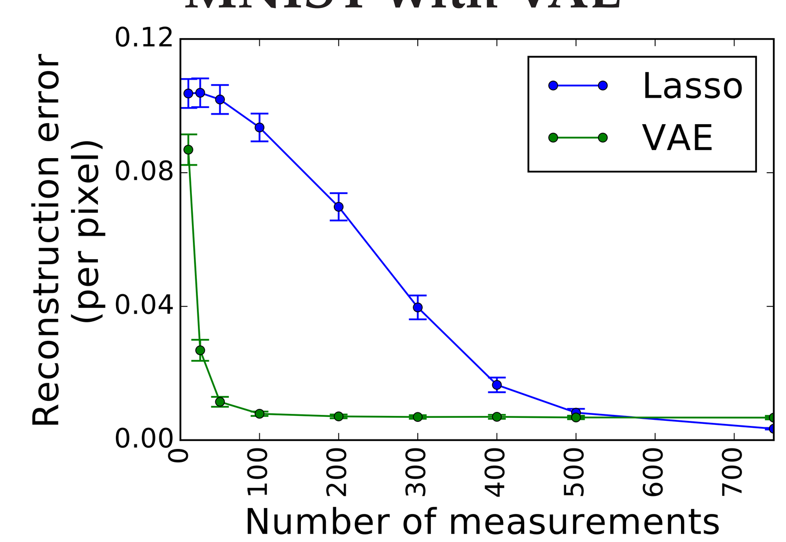

celebA with GAN

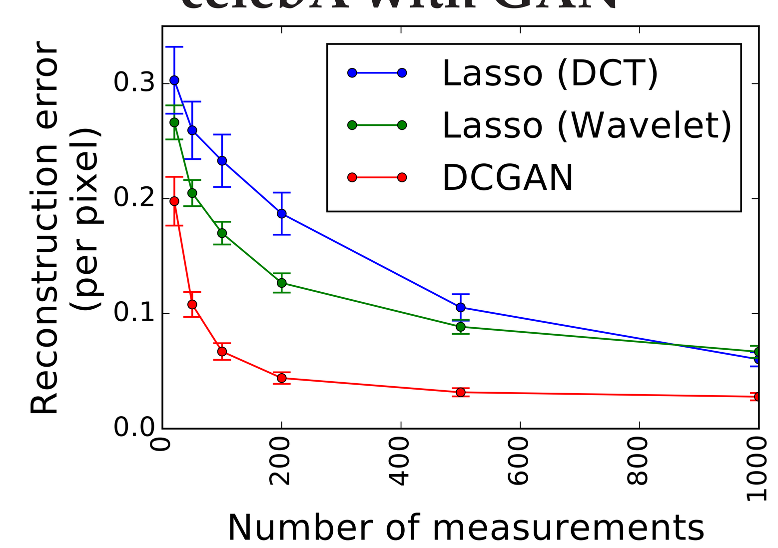

- For fixed G, error saturates. Larger m should use higher capacity model with larger k.
- Total Error = Representation + Measurement + Optimization. We find that (a) Optimization error ≈ 0, and (b) Representation error $\gg$ Measurement error.