Invertible Generative Models for Inverse Problems

Mitigating Representation Error and Dataset Bias

M. Asim, M. Daniels, O. Leong, P. Hand, A. Ahmed

Inverse Problems with Generative Models as Image Priors

Inverse Problems with Generative Models as Image Priors

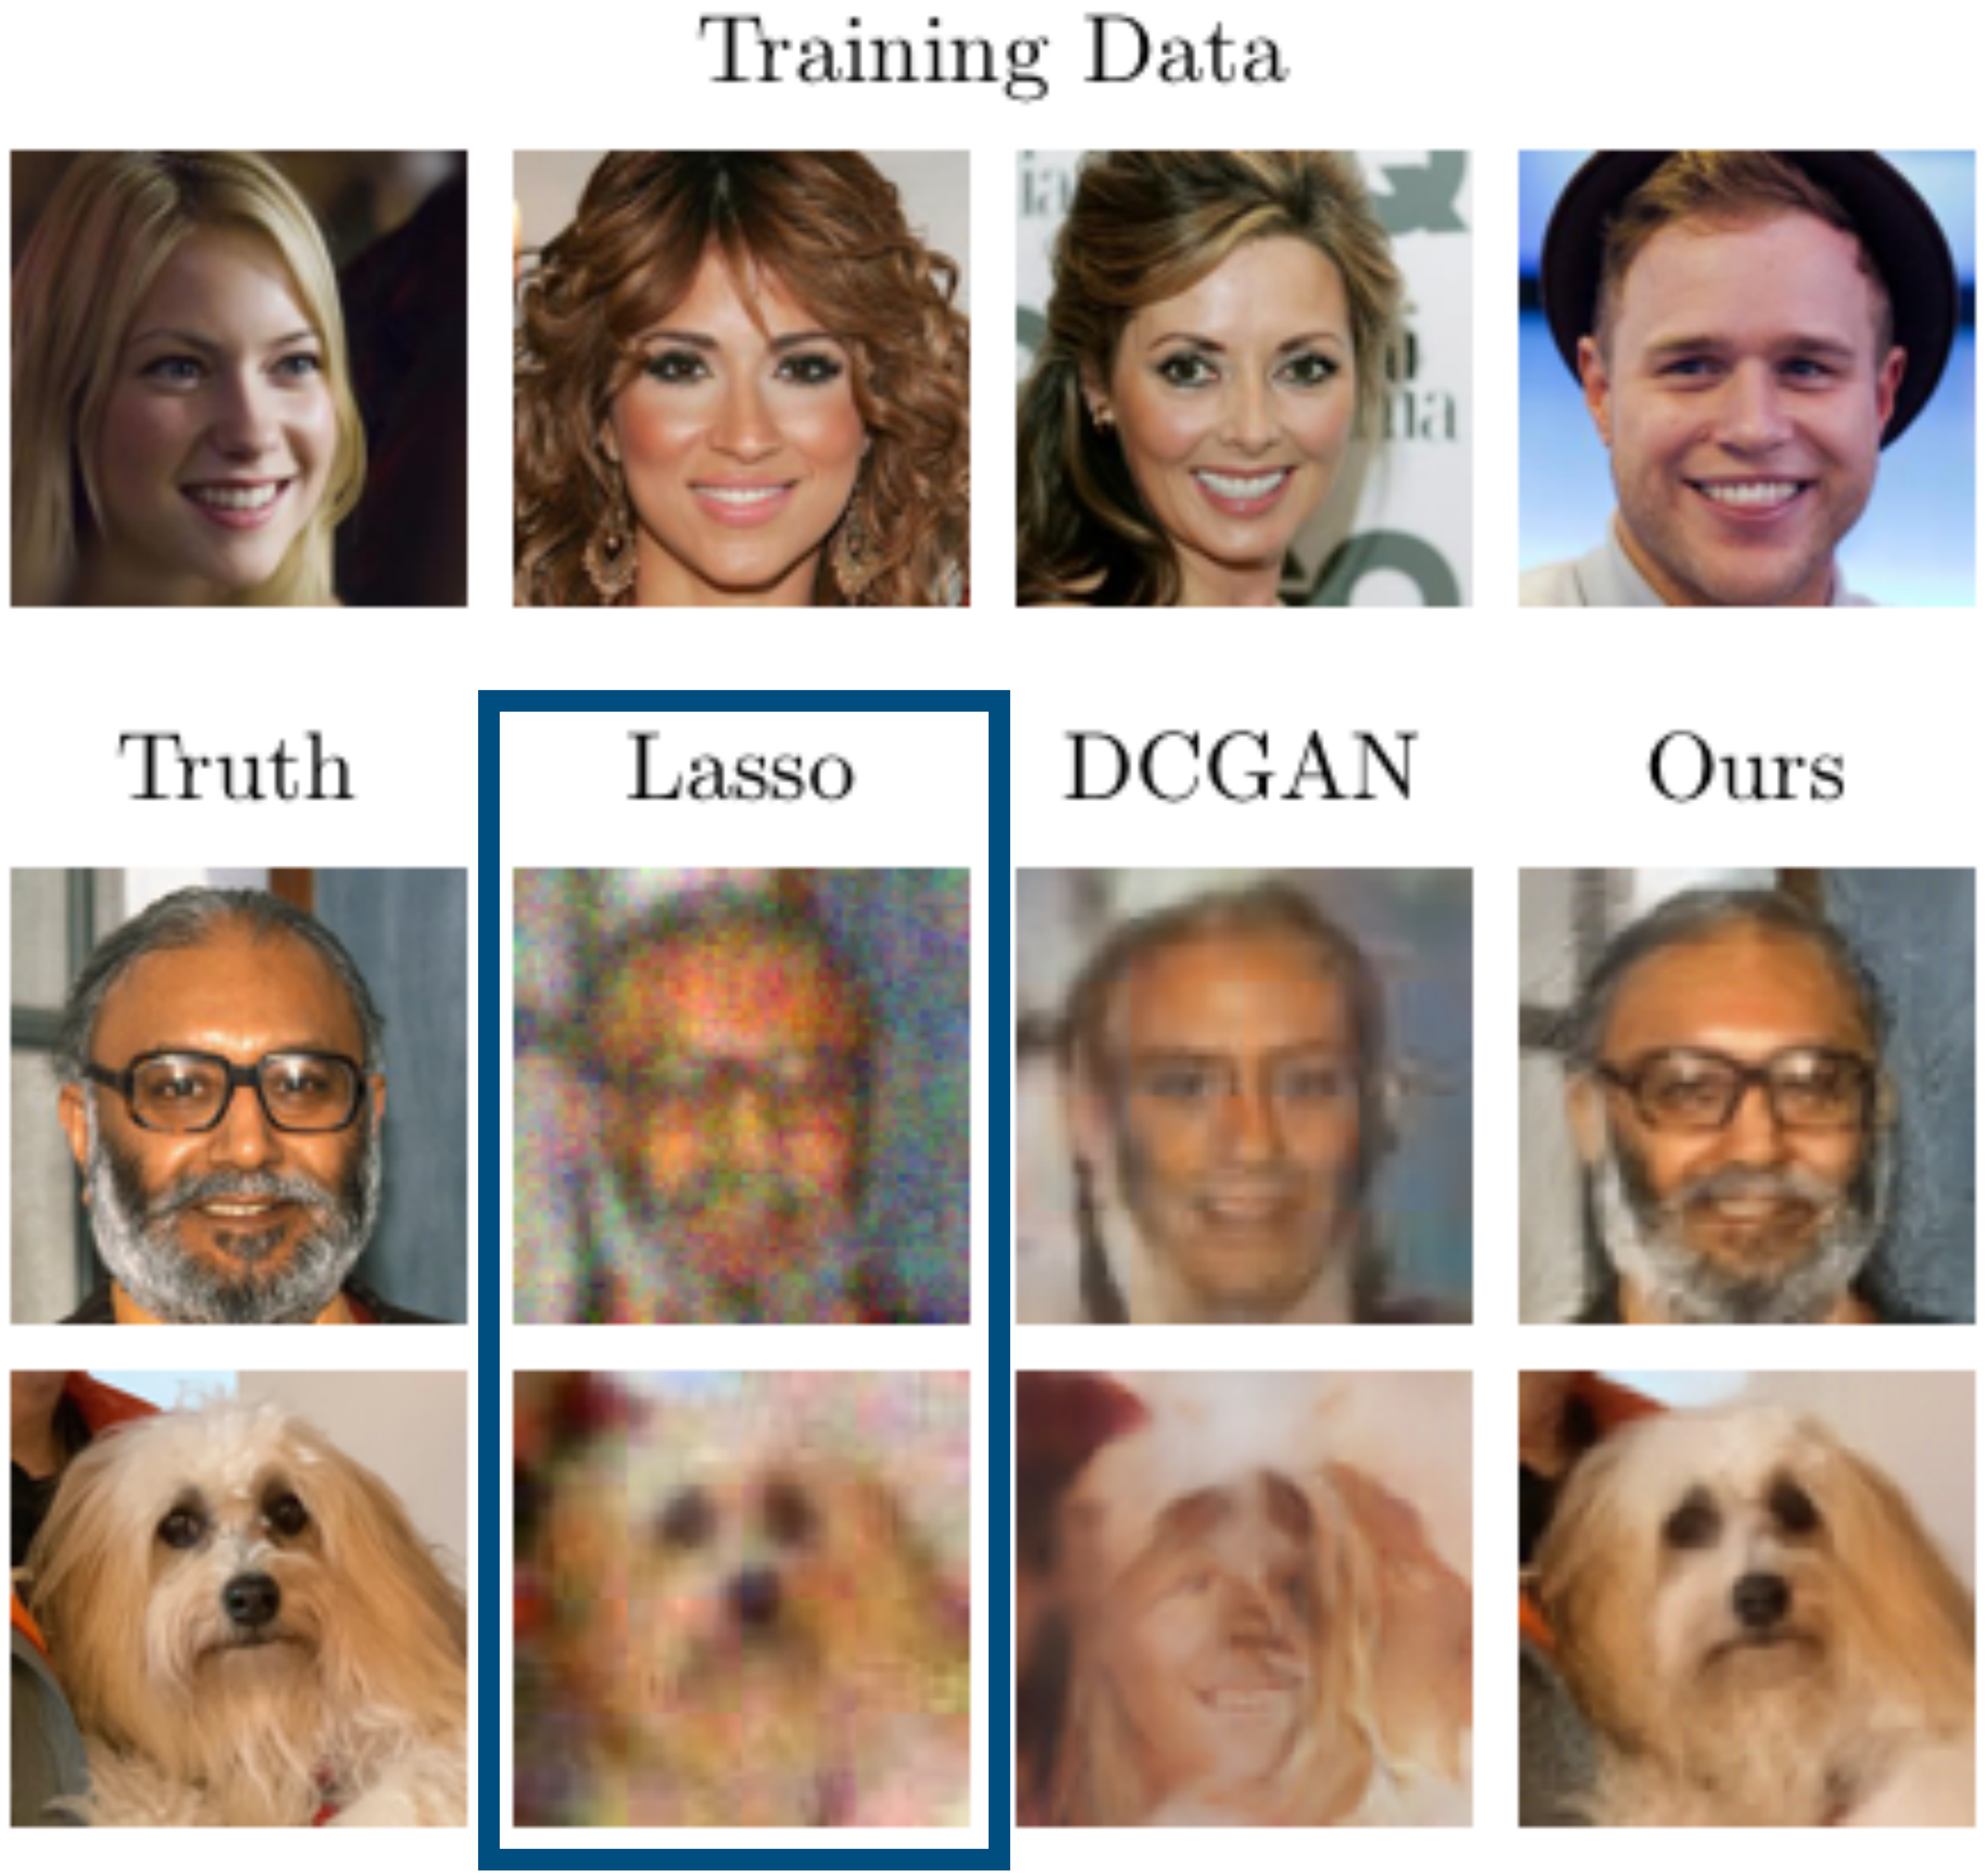

Inverse Problems with Generative Models as Image Priors

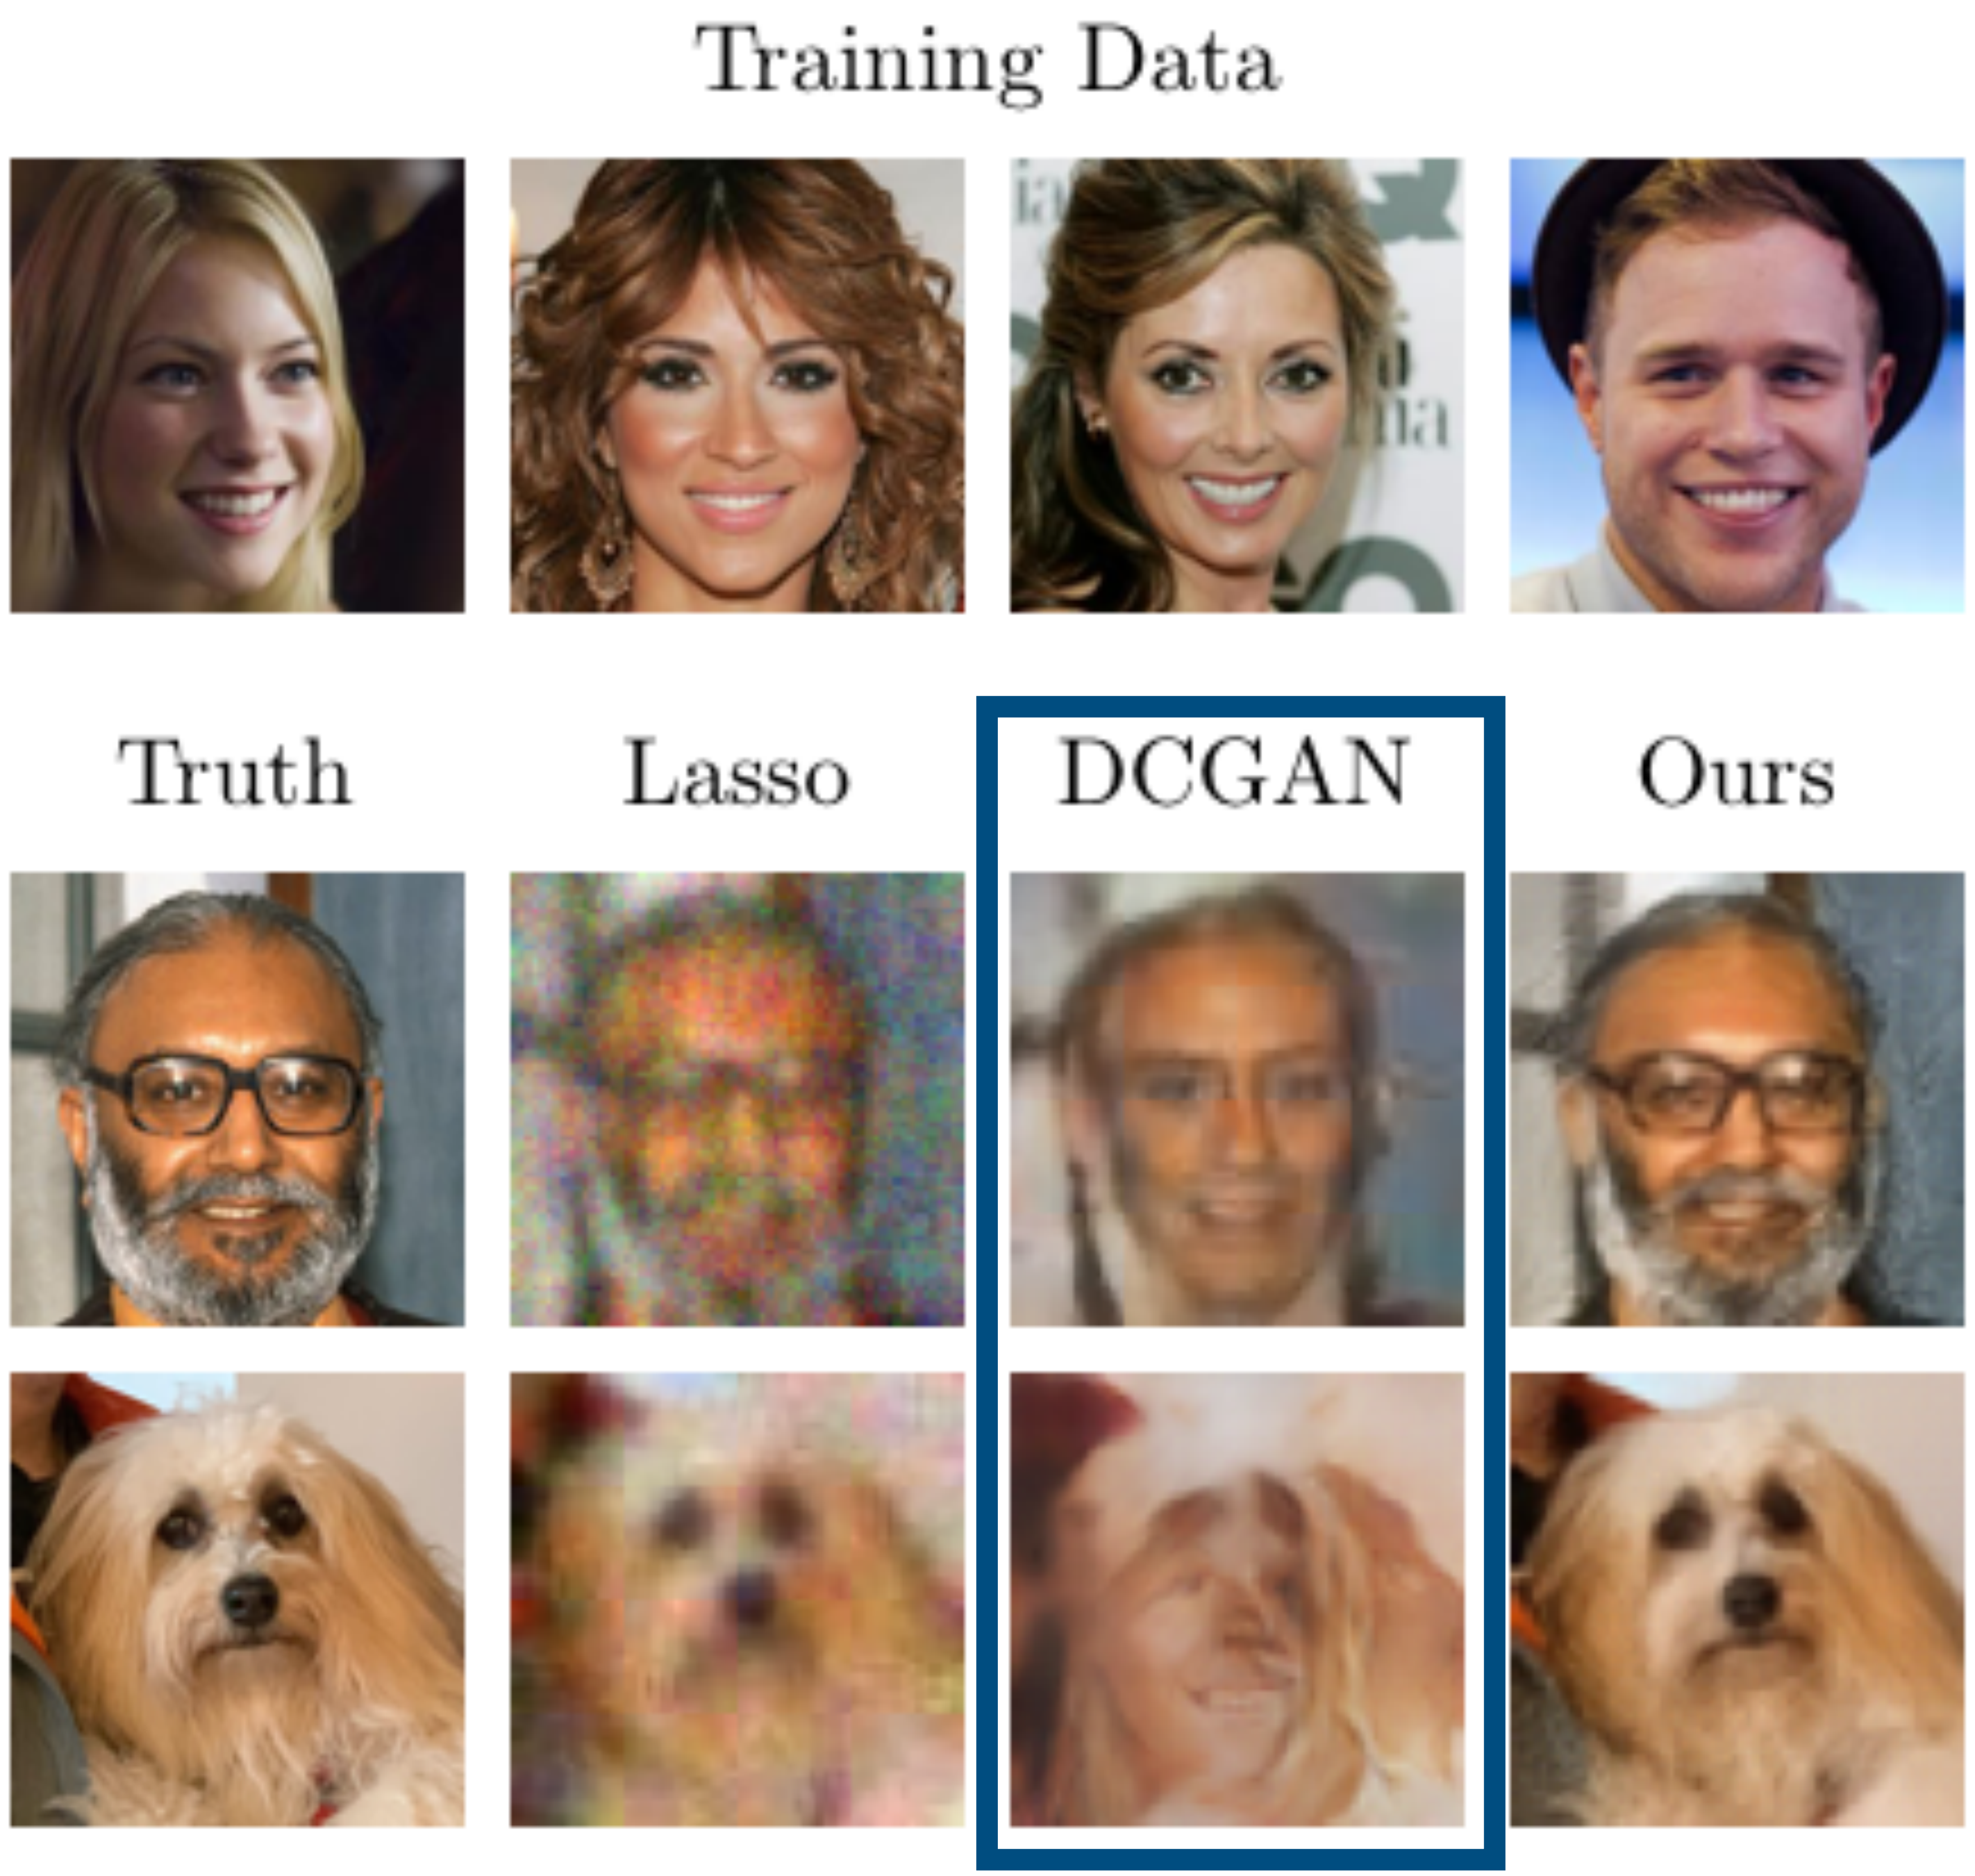

Inverse Problems with Generative Models as Image Priors

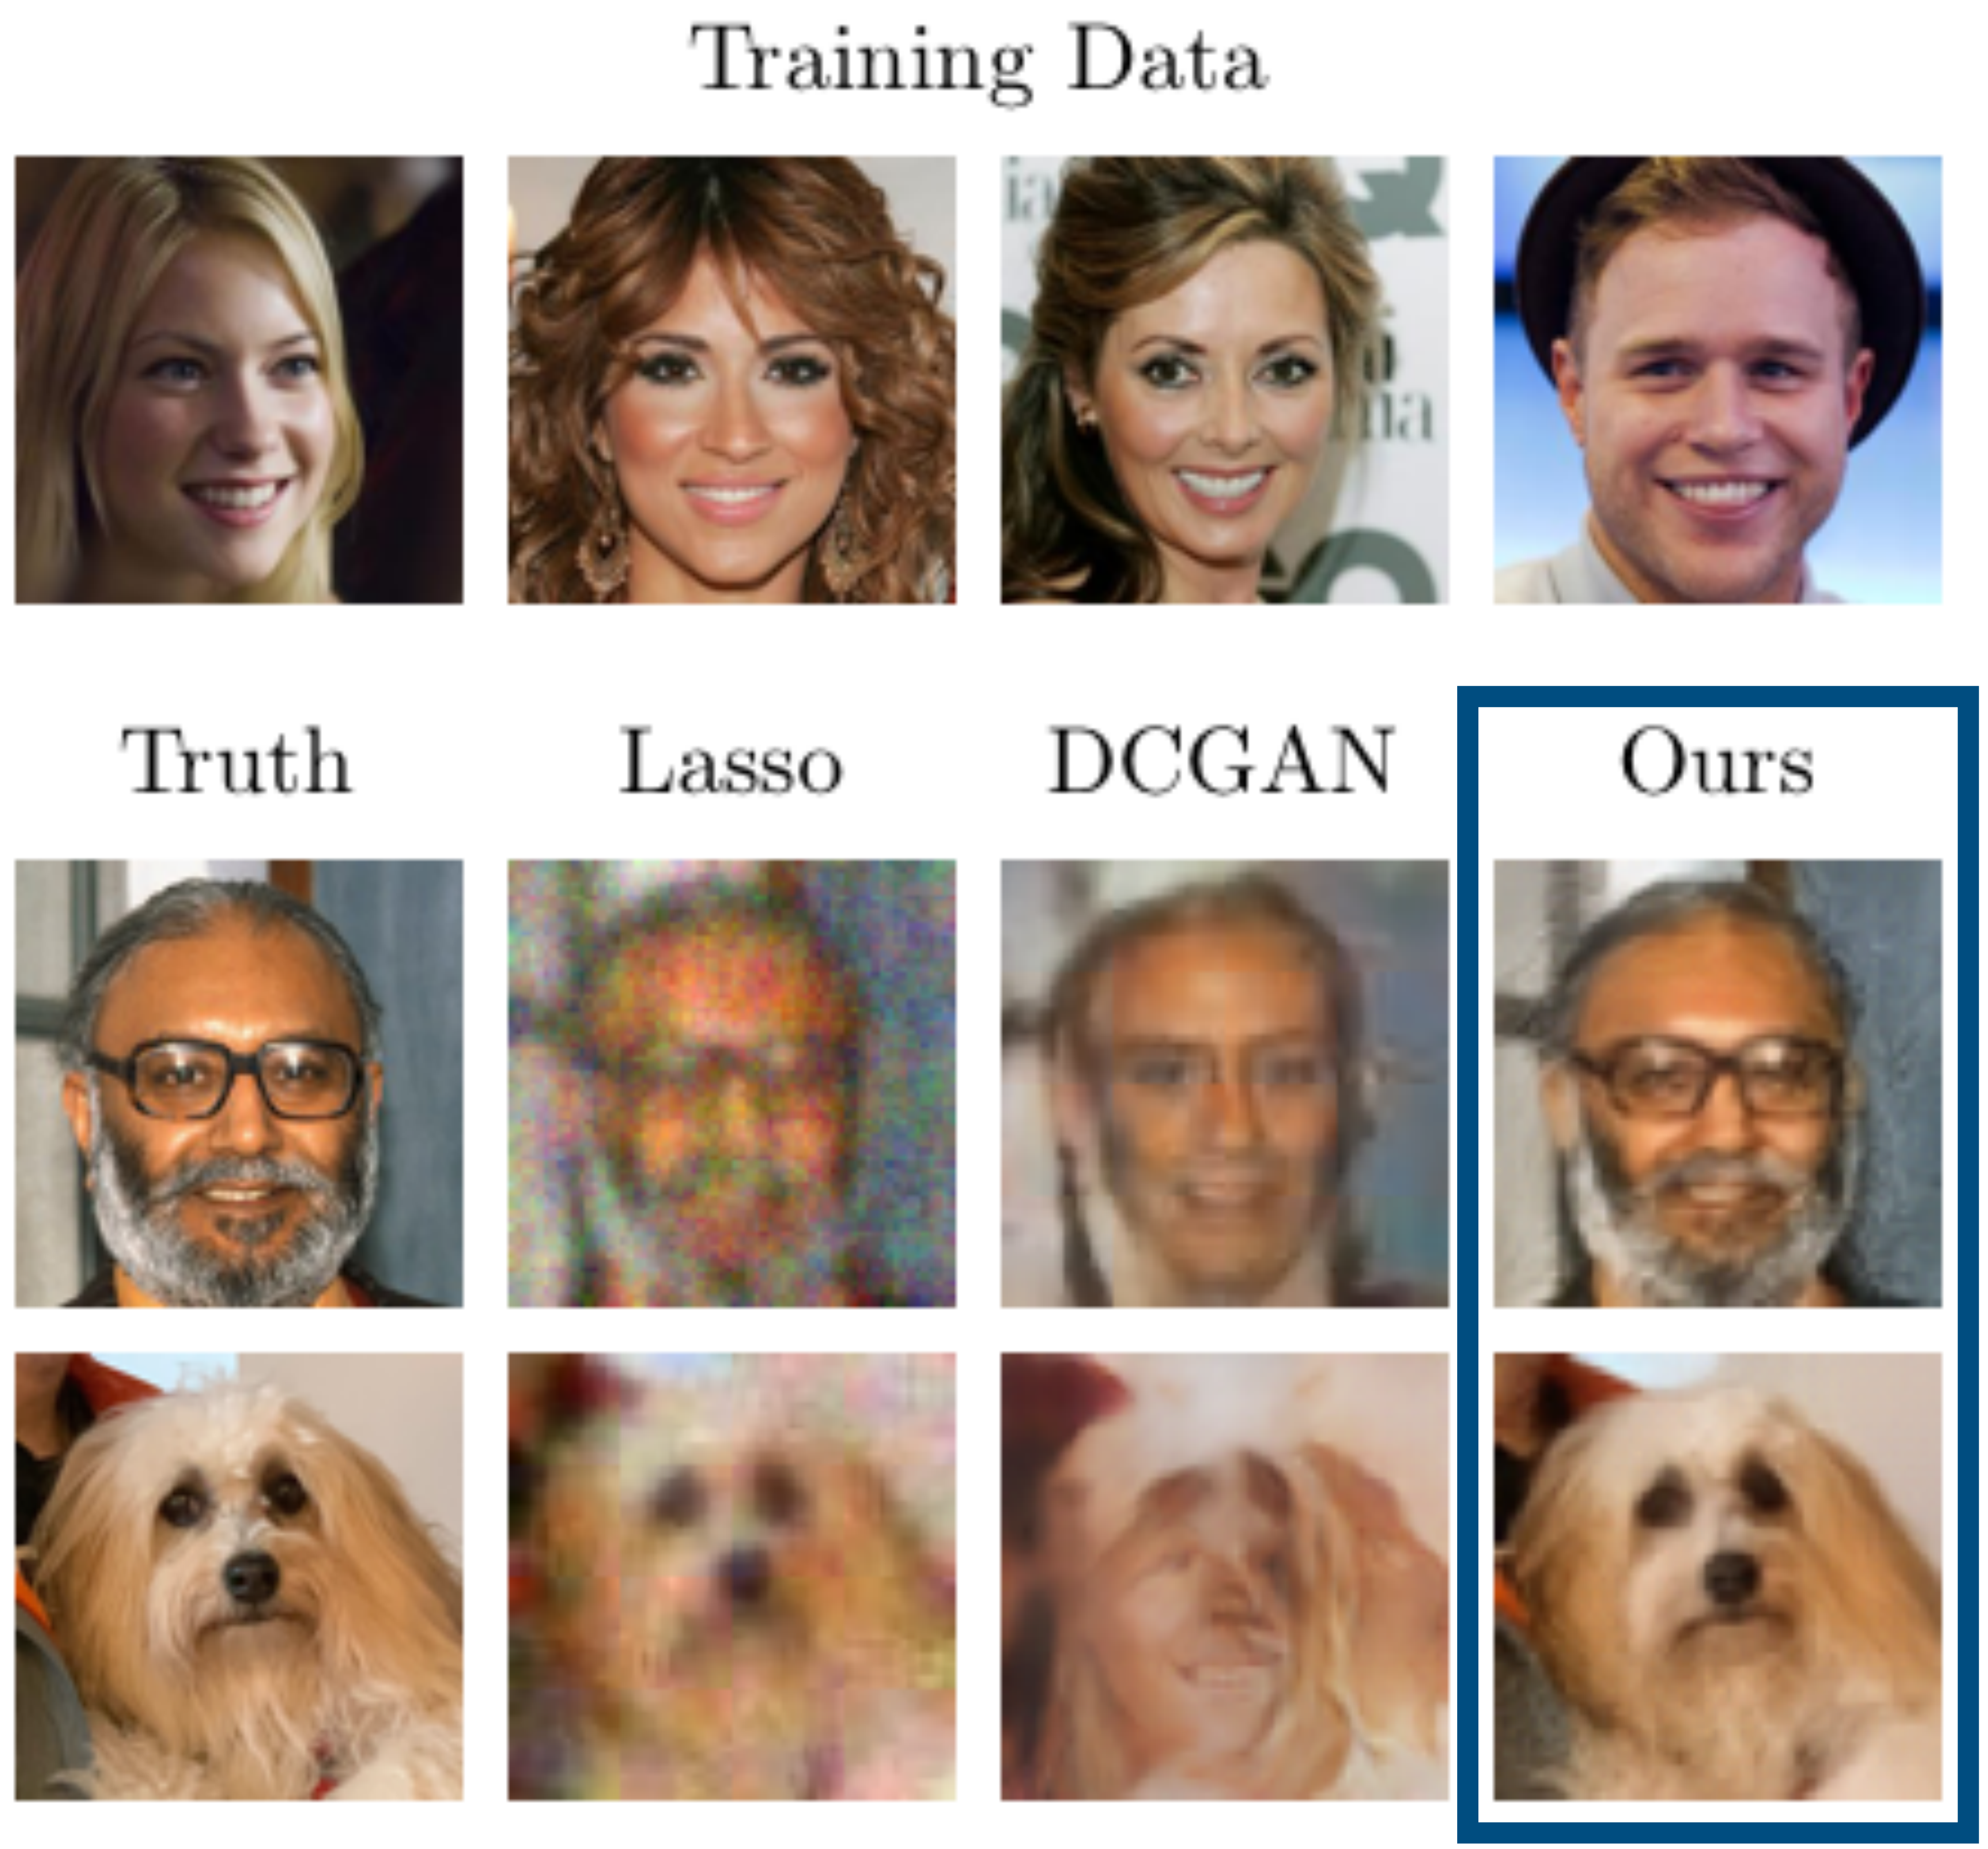

Contributions

1. Trained INN priors provide SOTA performance in a variety of inverse problems

2. Trained INN priors exhibit strong performance on out-of-distribution images

3. Theoretical guarantees in the case of linear invertible model

Linear Inverse Problems in Imaging

Measurement matrix $A \in \mathbb{R}^{m \times n}$

m noisy measurements

$$y = Ax_0 + \eta$$

Recover x_0

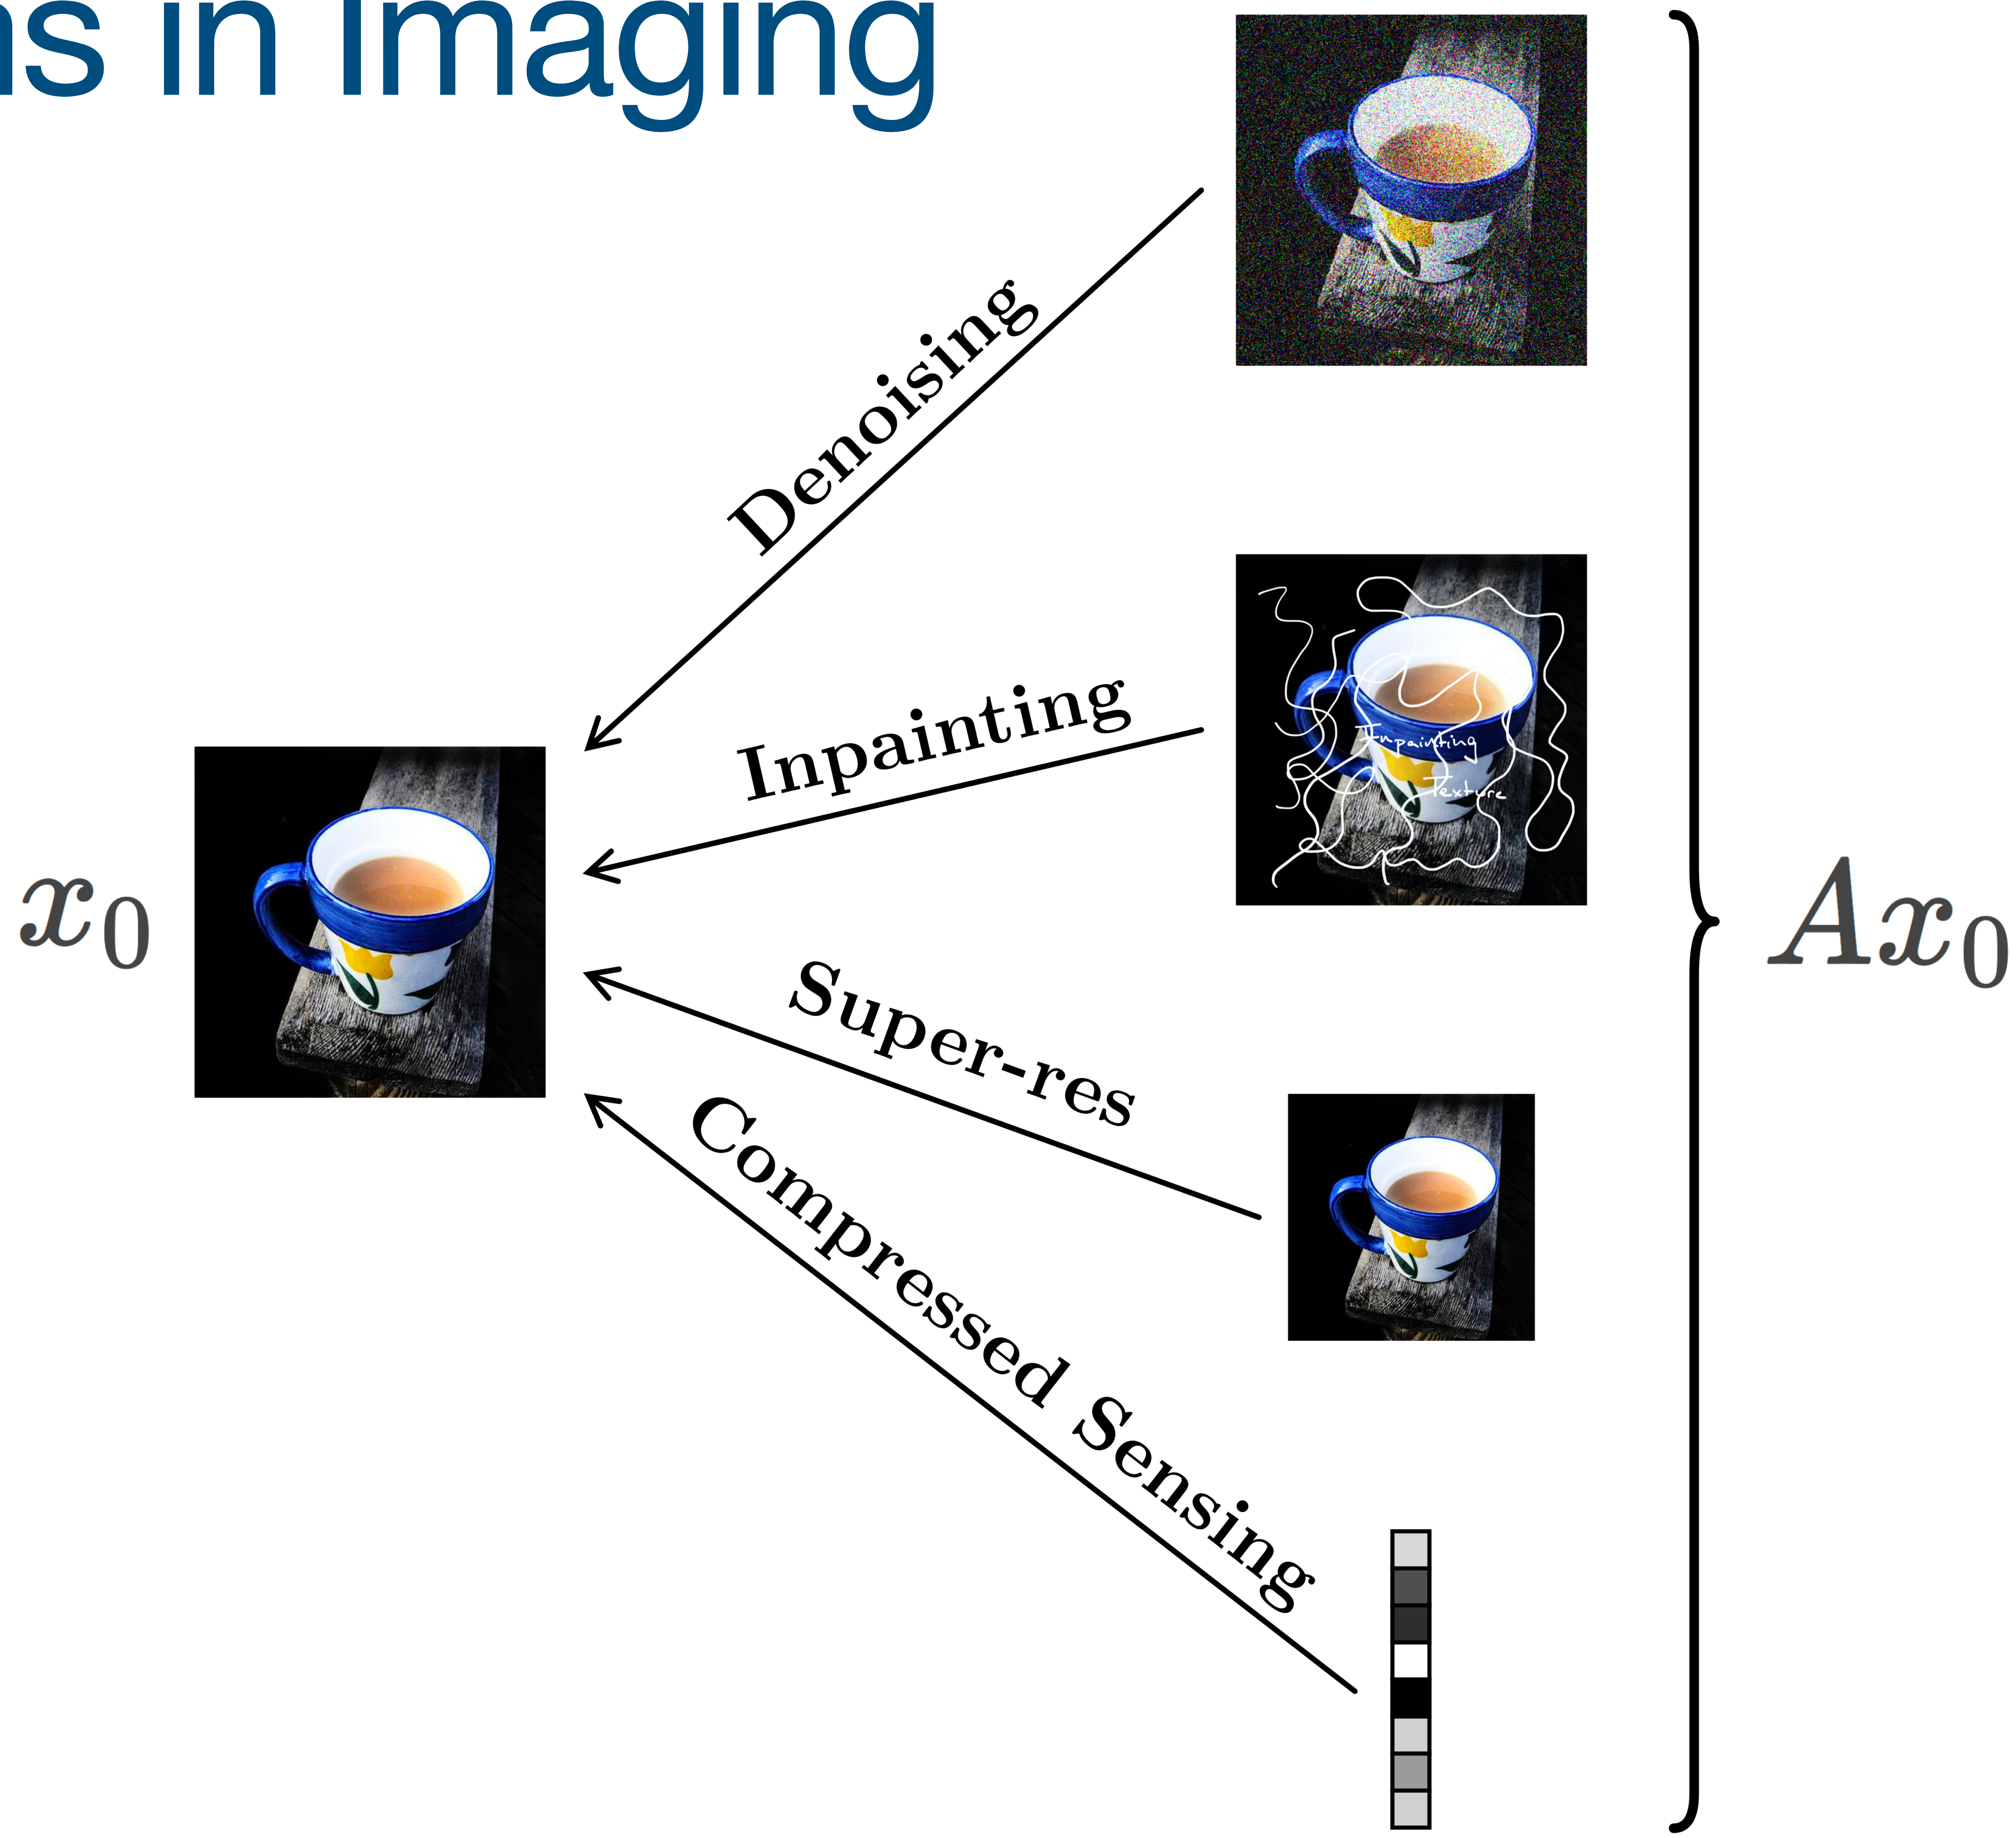

Invertible Generative Models via Normalizing Flows

- Learned invertible map
- Maps Gaussian to signal distribution
- Signal is a composition of Flow steps
- Admits exact calculation of image likelihood

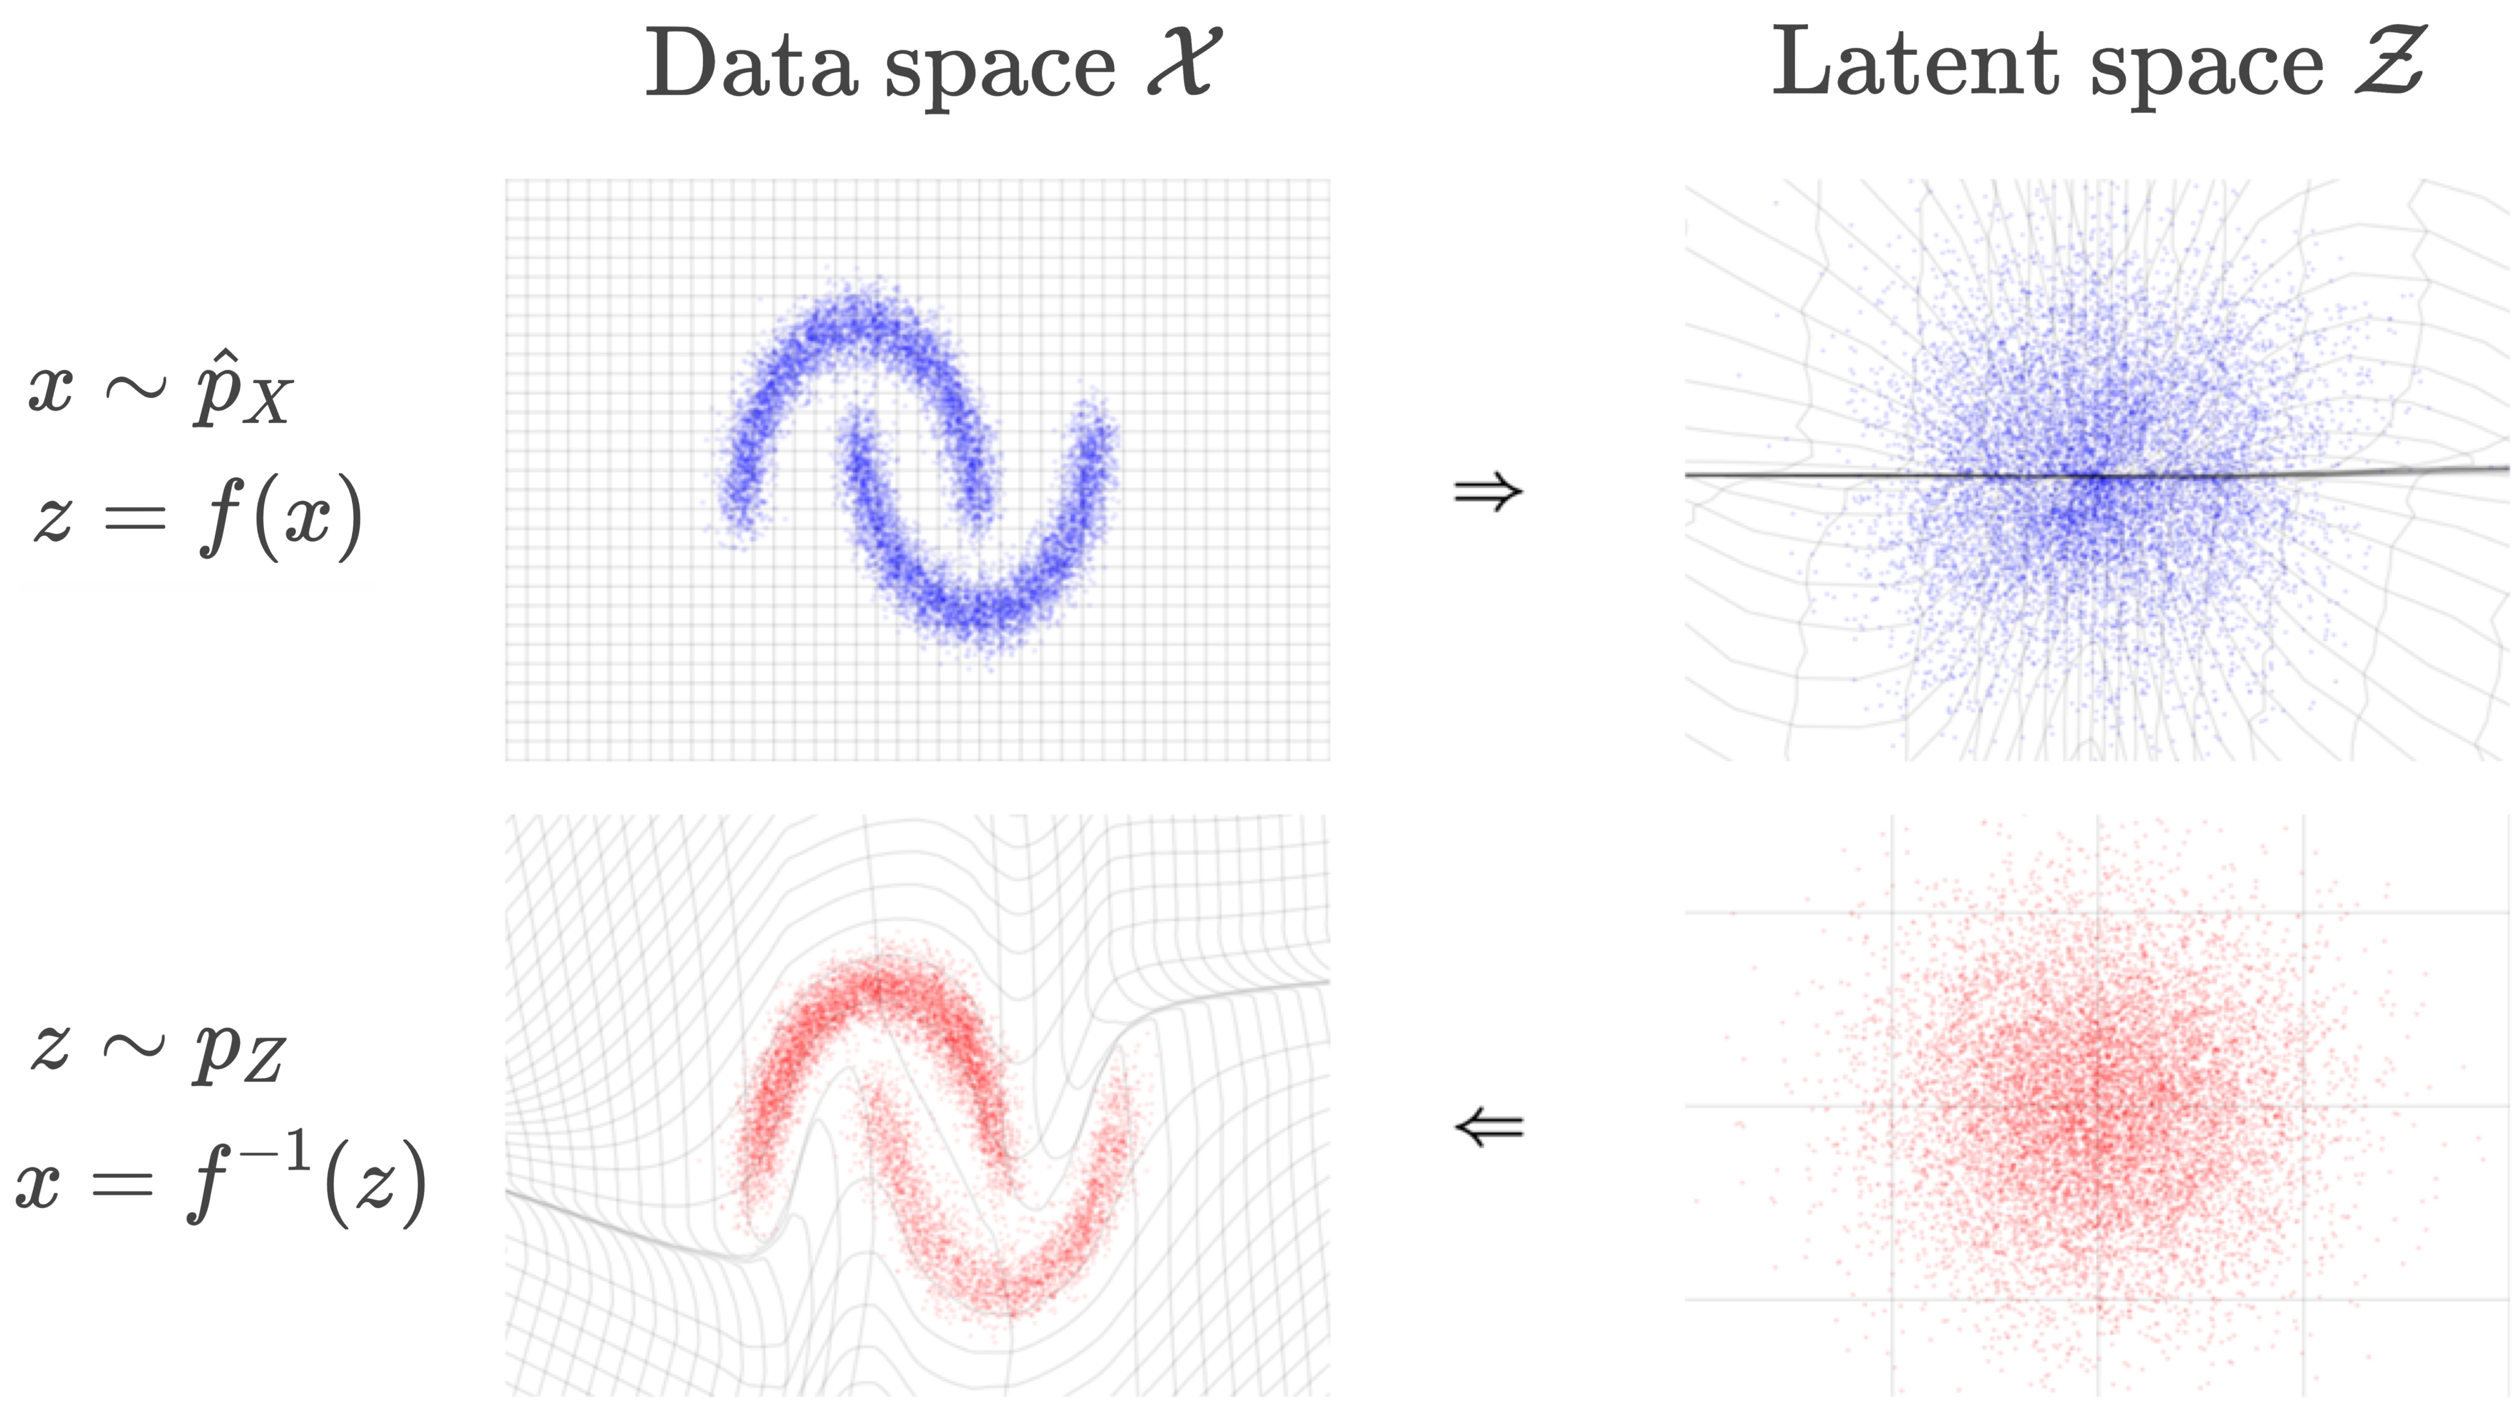

Fig 1. RealNVP (Dinh, Sohl-Dickstein, Bengio)

Central Architectural Element: affine coupling layer

Affine coupling layer:

1. Split input activations

$$x = (x_1, x_2)$$

2. Compute learned affine transform

$$s, t = f_\theta(x_2)$$

3. Apply the transformation

$$y_1 = s \odot x_1 + t$$

Has a tractable Jacobian determinant

Examples: RealNVP, GLOW

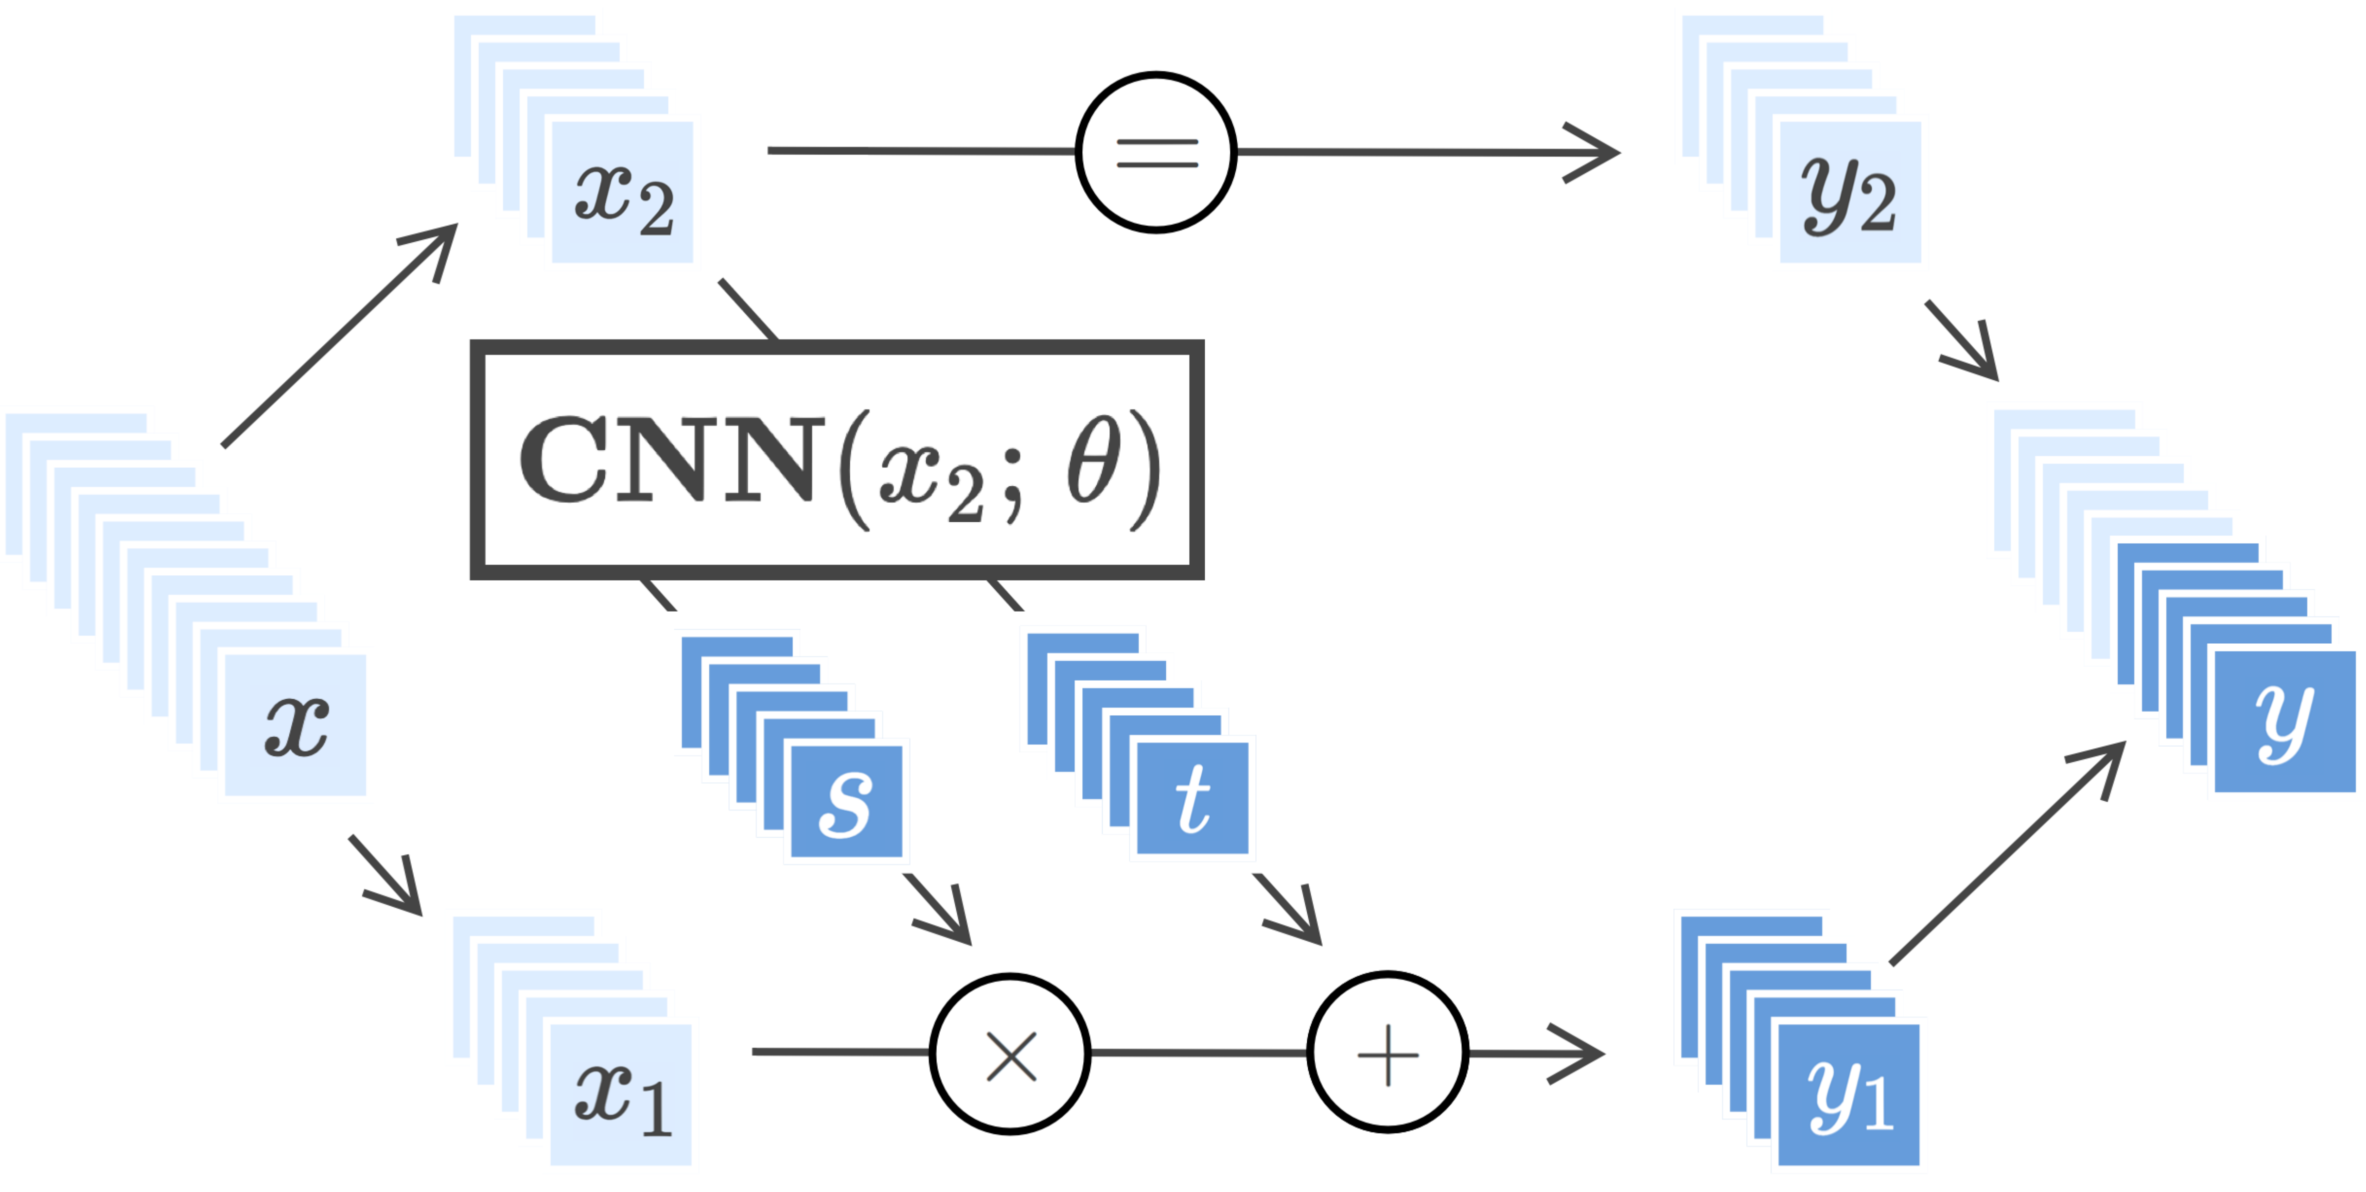

Fig 2. RealNVP (Dinh, Sohl-Dickstein, Bengio)

Formulation for Denoising

Given:

1. Noisy measurements of all pixels:

   $$y = x_0 + \eta, \ \eta \sim \mathcal{N}(0, \sigma^2 I_n)$$

2. Trained INN:

   $$G : \mathbb{R}^n \to \mathbb{R}^n$$

Find: x_0

MLE formulation over x -space:

$$\min_{x \in \mathbb{R}^n} \|x - y\|^2 - \gamma \log p_G(x)$$

Proxy in z -space:

$$\min_{z \in \mathbb{R}^n} \|G(z) - y\|^2 + \gamma \|z\|^2$$

INNs can outperform BM3D in denoising

Given:

1. Noisy measurements of all pixels:

   $$y = x_0 + \eta,\ \eta \sim \mathcal{N}(0, \sigma^2 I_n)$$

2. Trained INN:

   $$G : \mathbb{R}^n \to \mathbb{R}^n$$

Find: x_0

$$\min_{z \in \mathbb{R}^n} \|G(z) - y\|^2 + \gamma \|z\|^2$$

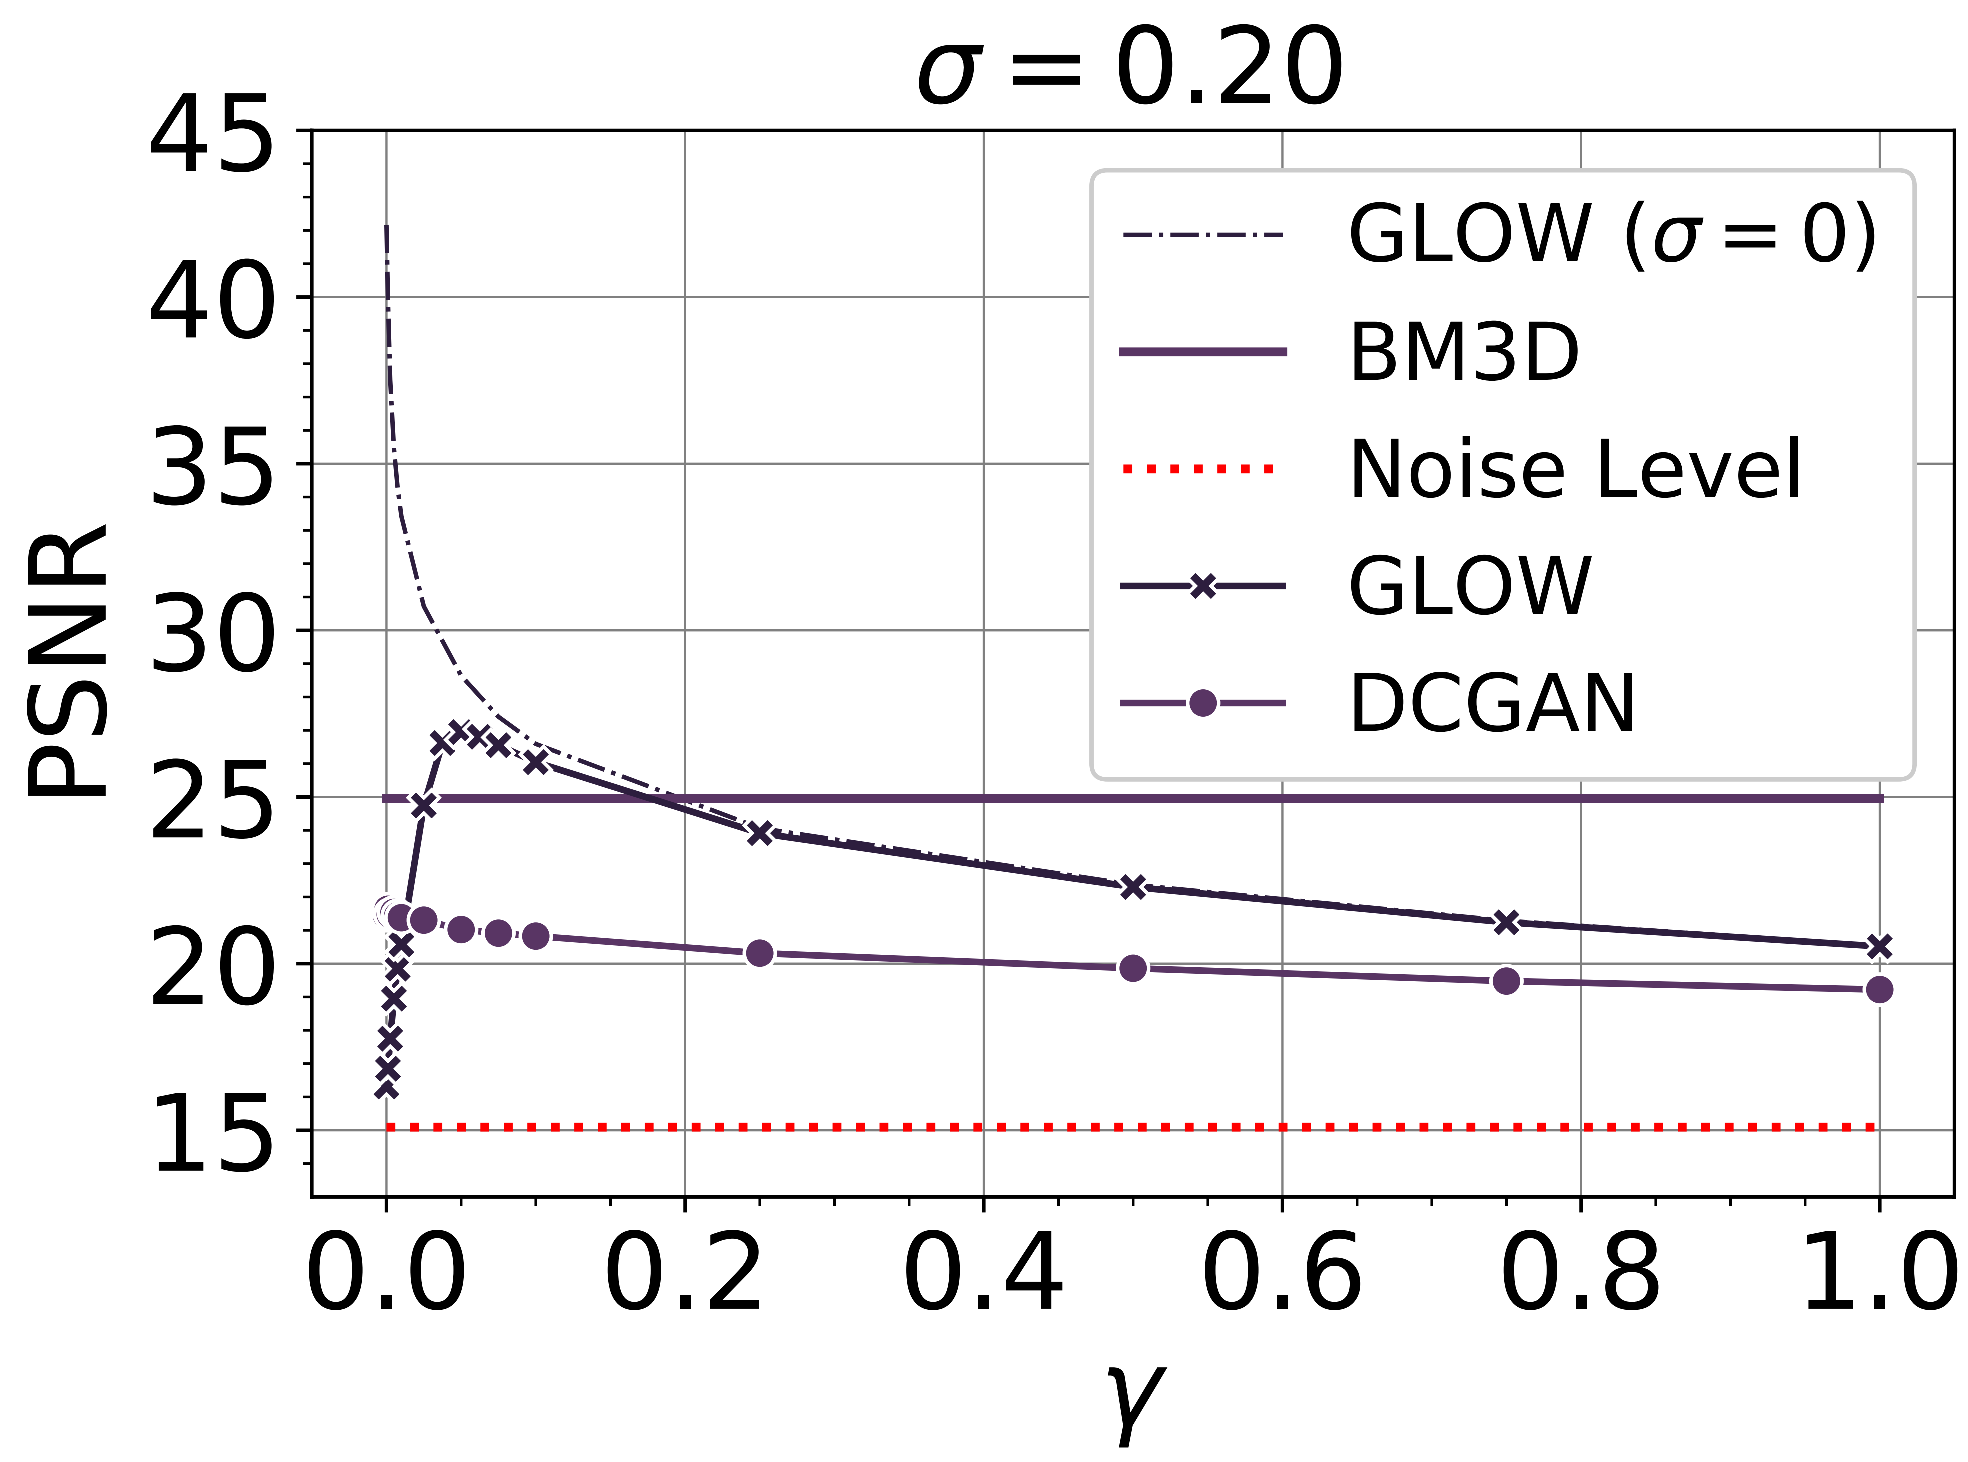

Formulation for Compressed Sensing

Given: $y = Ax_0 + \eta$

$$A \in \mathbb{R}^{m \times n},\ A_{ij} \sim \mathcal{N}(0, 1/m),\ m < n$$

Find: $\hat{z}$ s.t. $G(\hat{z}) \approx x_0$

Solve via optimization in z -space:

$$\min_{z \in \mathbb{R}^n} ||AG(z) - y||^2$$

Compressed Sensing

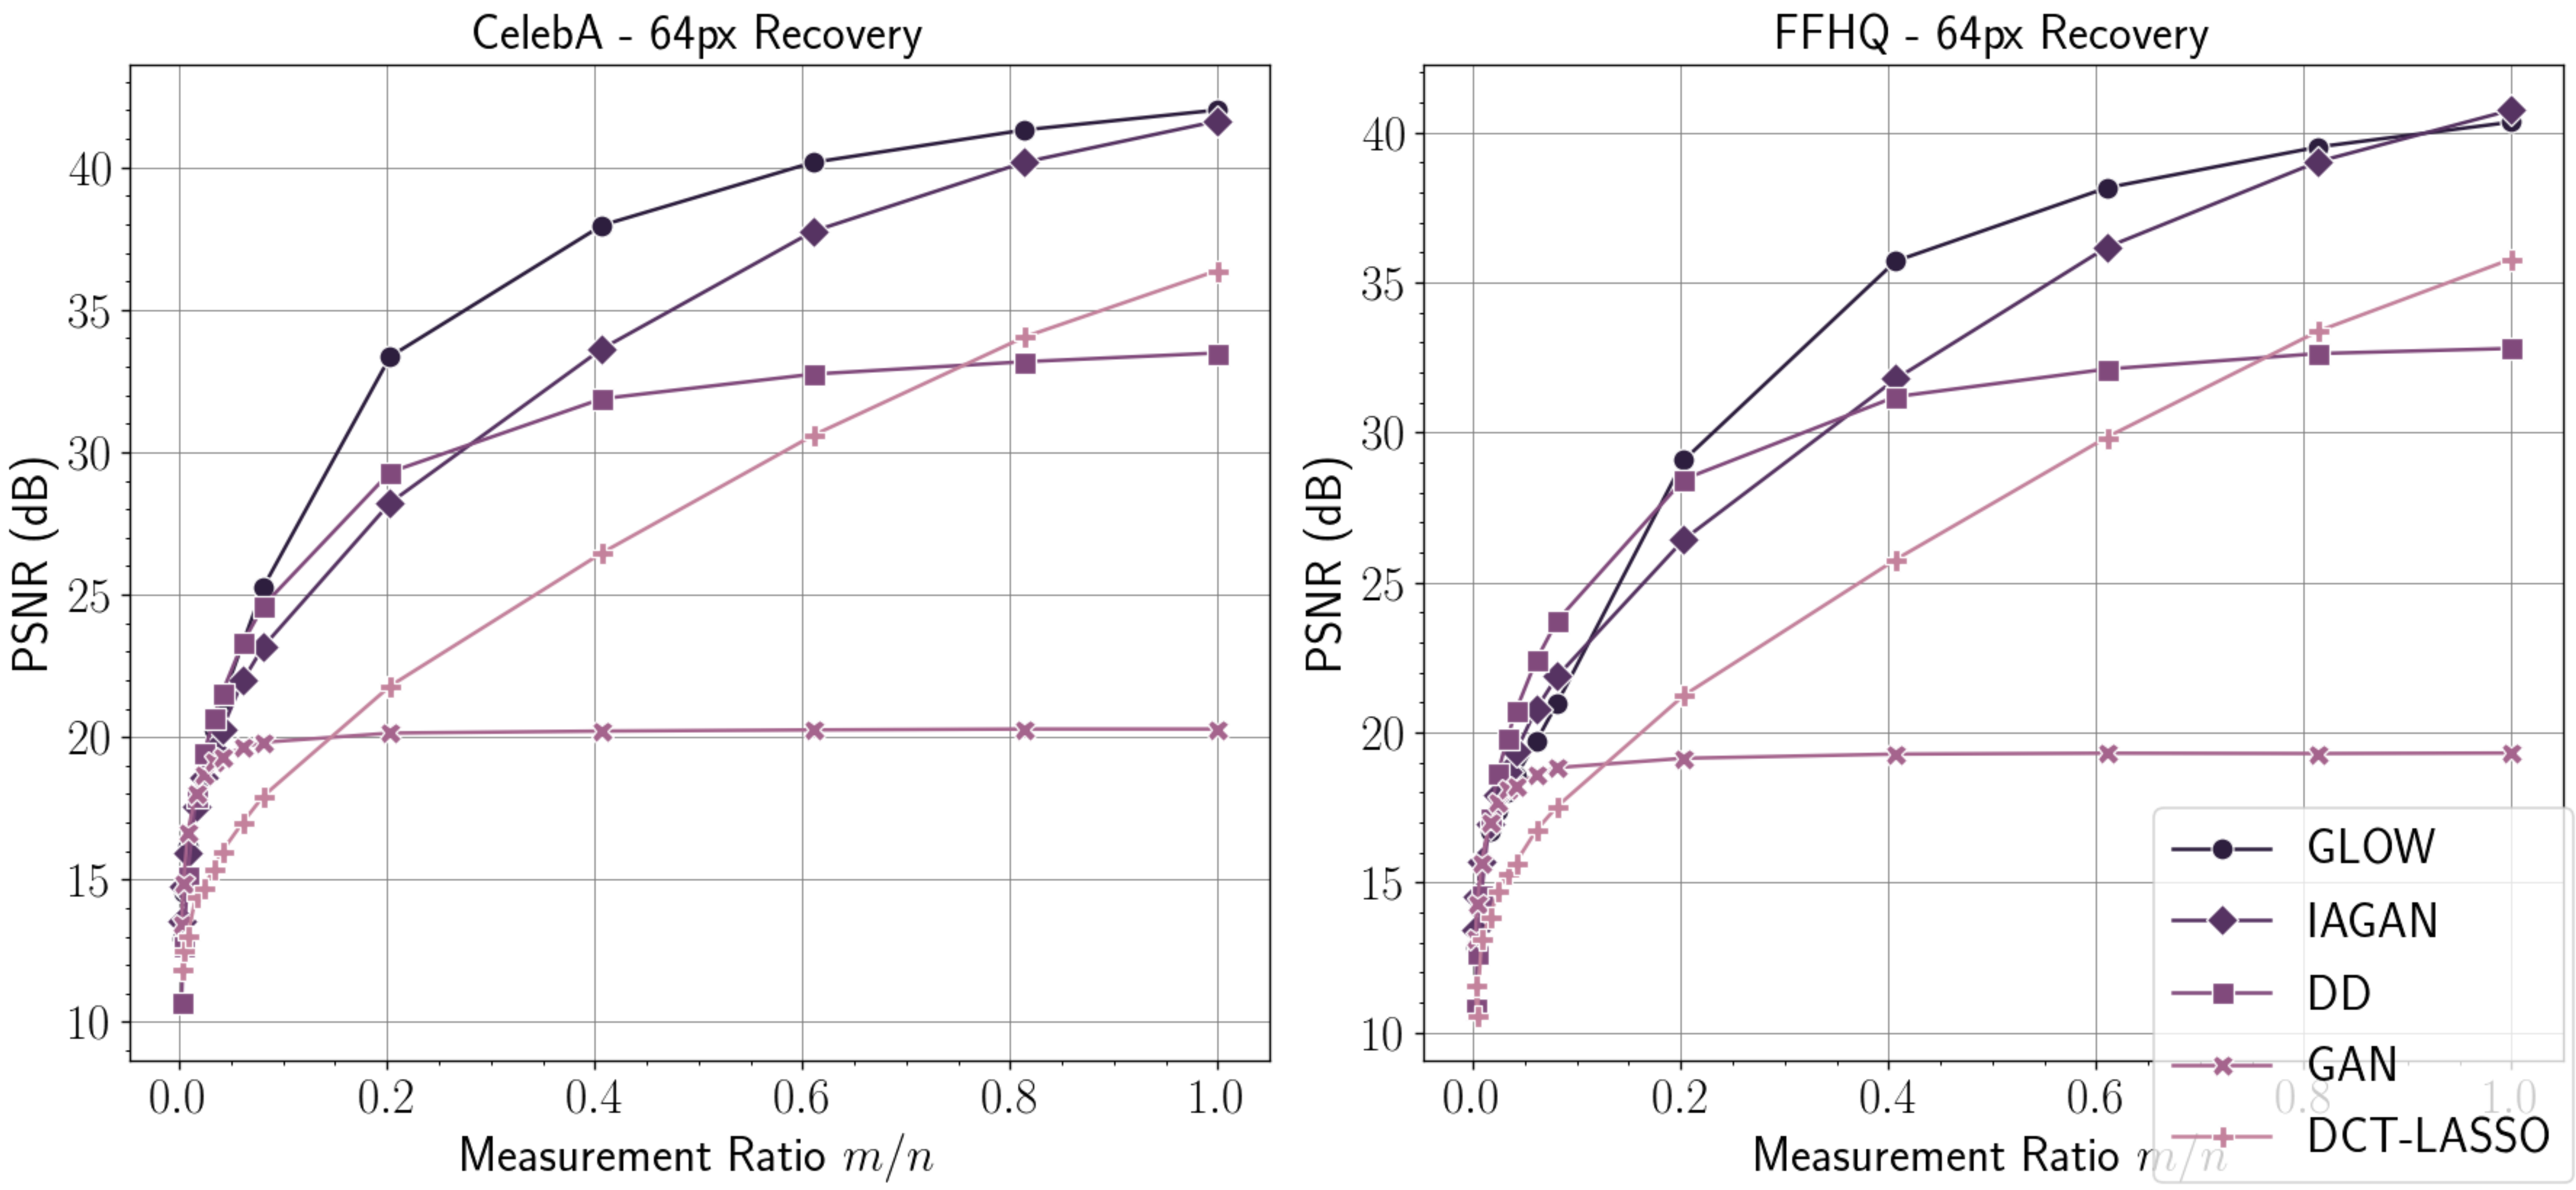

INNs exhibit strong OOD performance

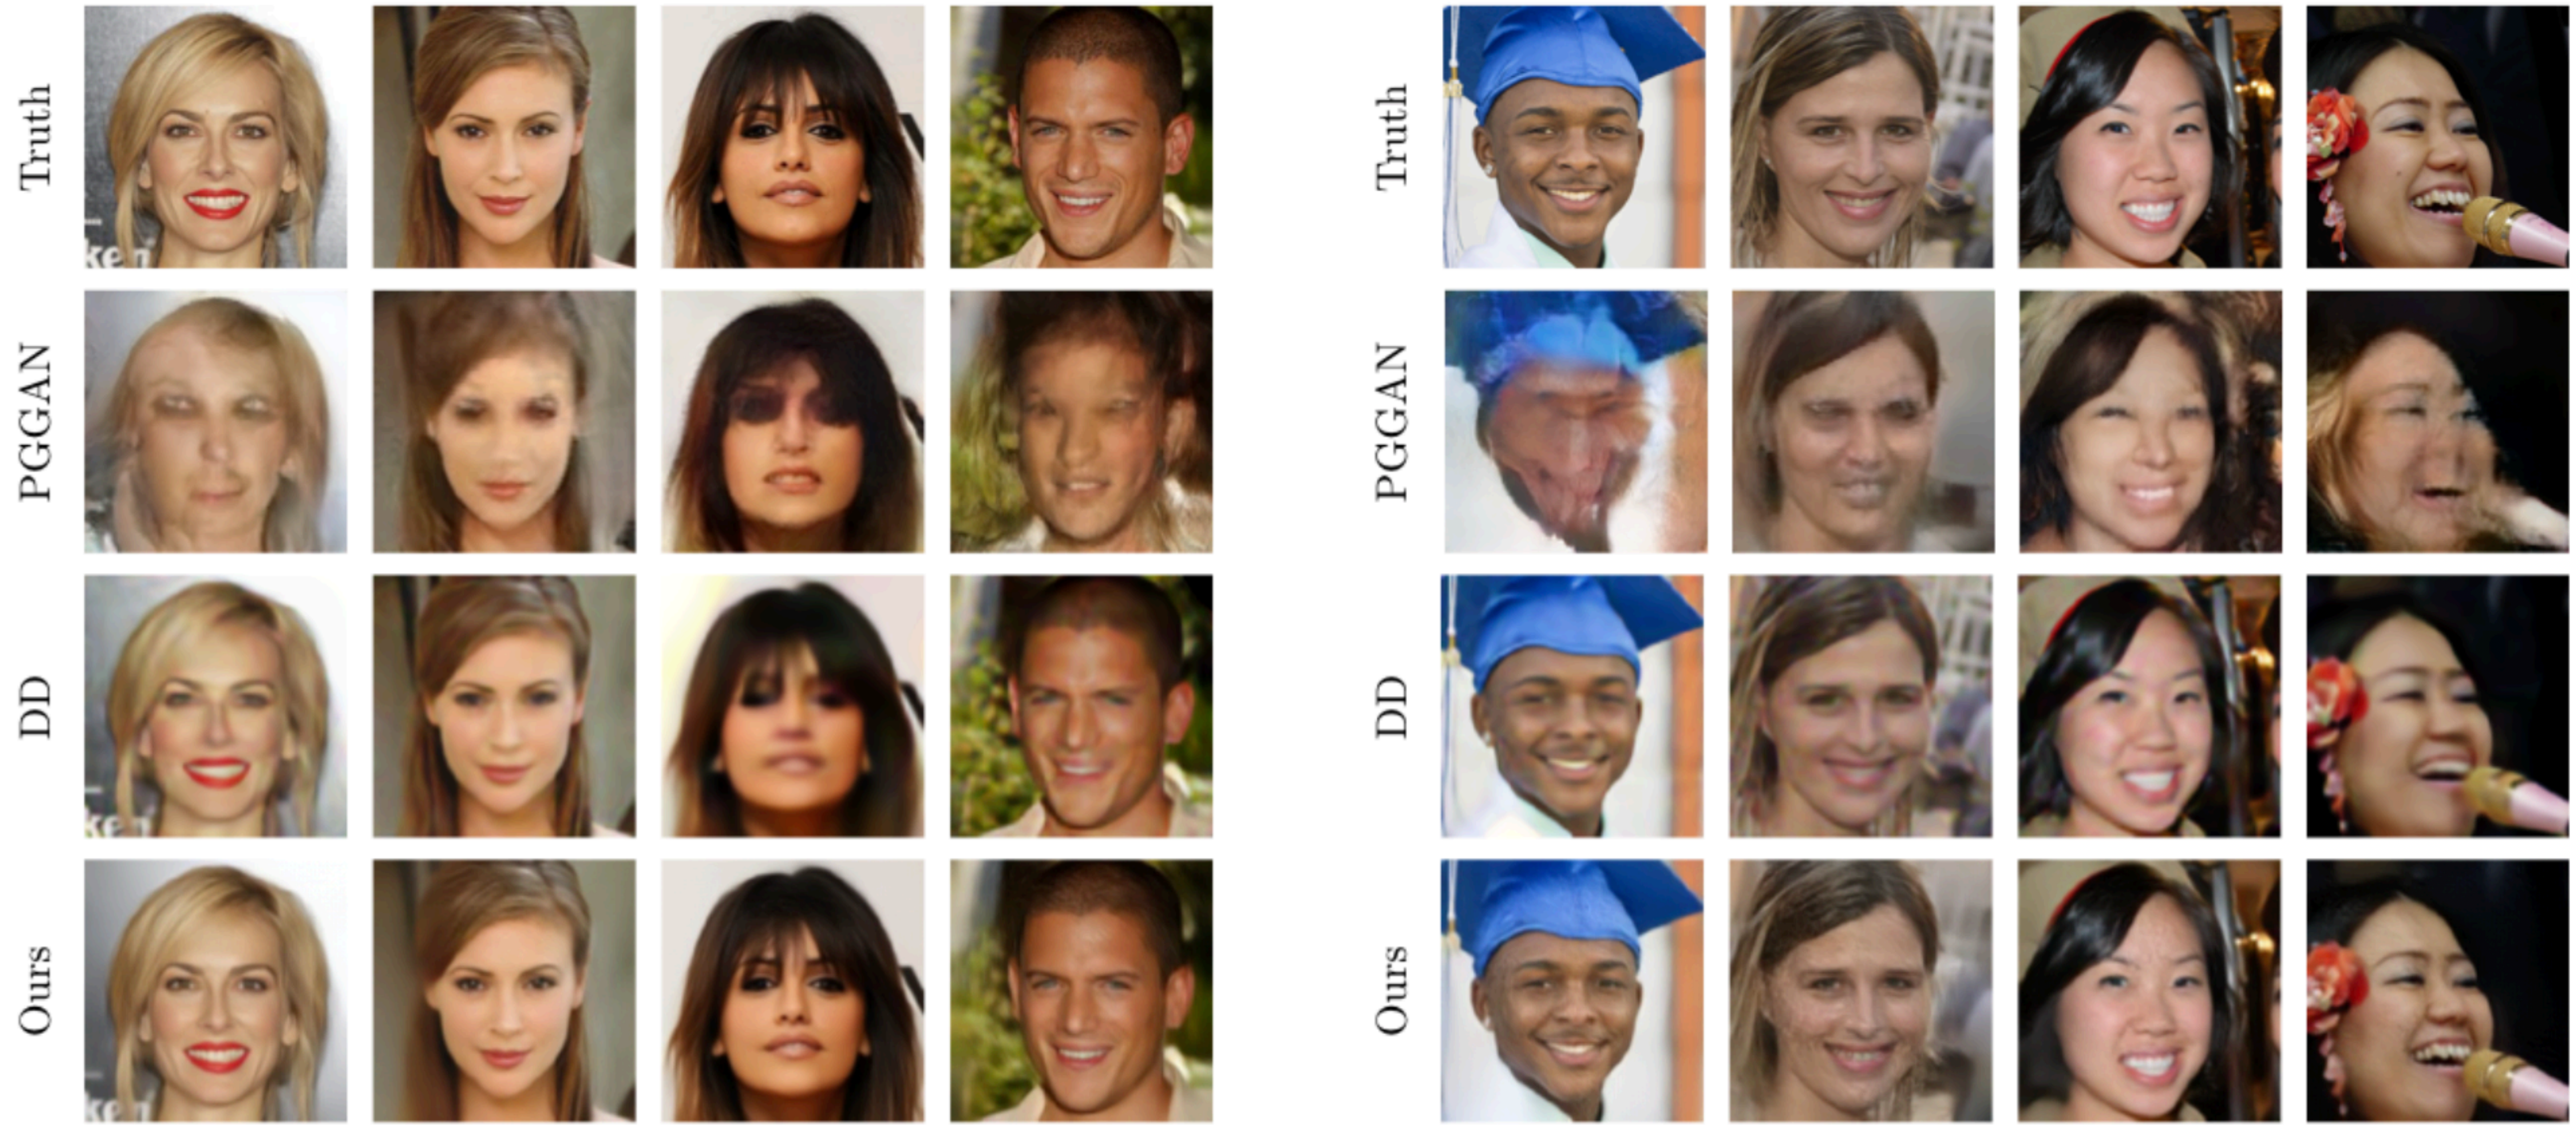

INNs exhibit strong OOD performance

Strong OOD Performance on Semantic Inpainting

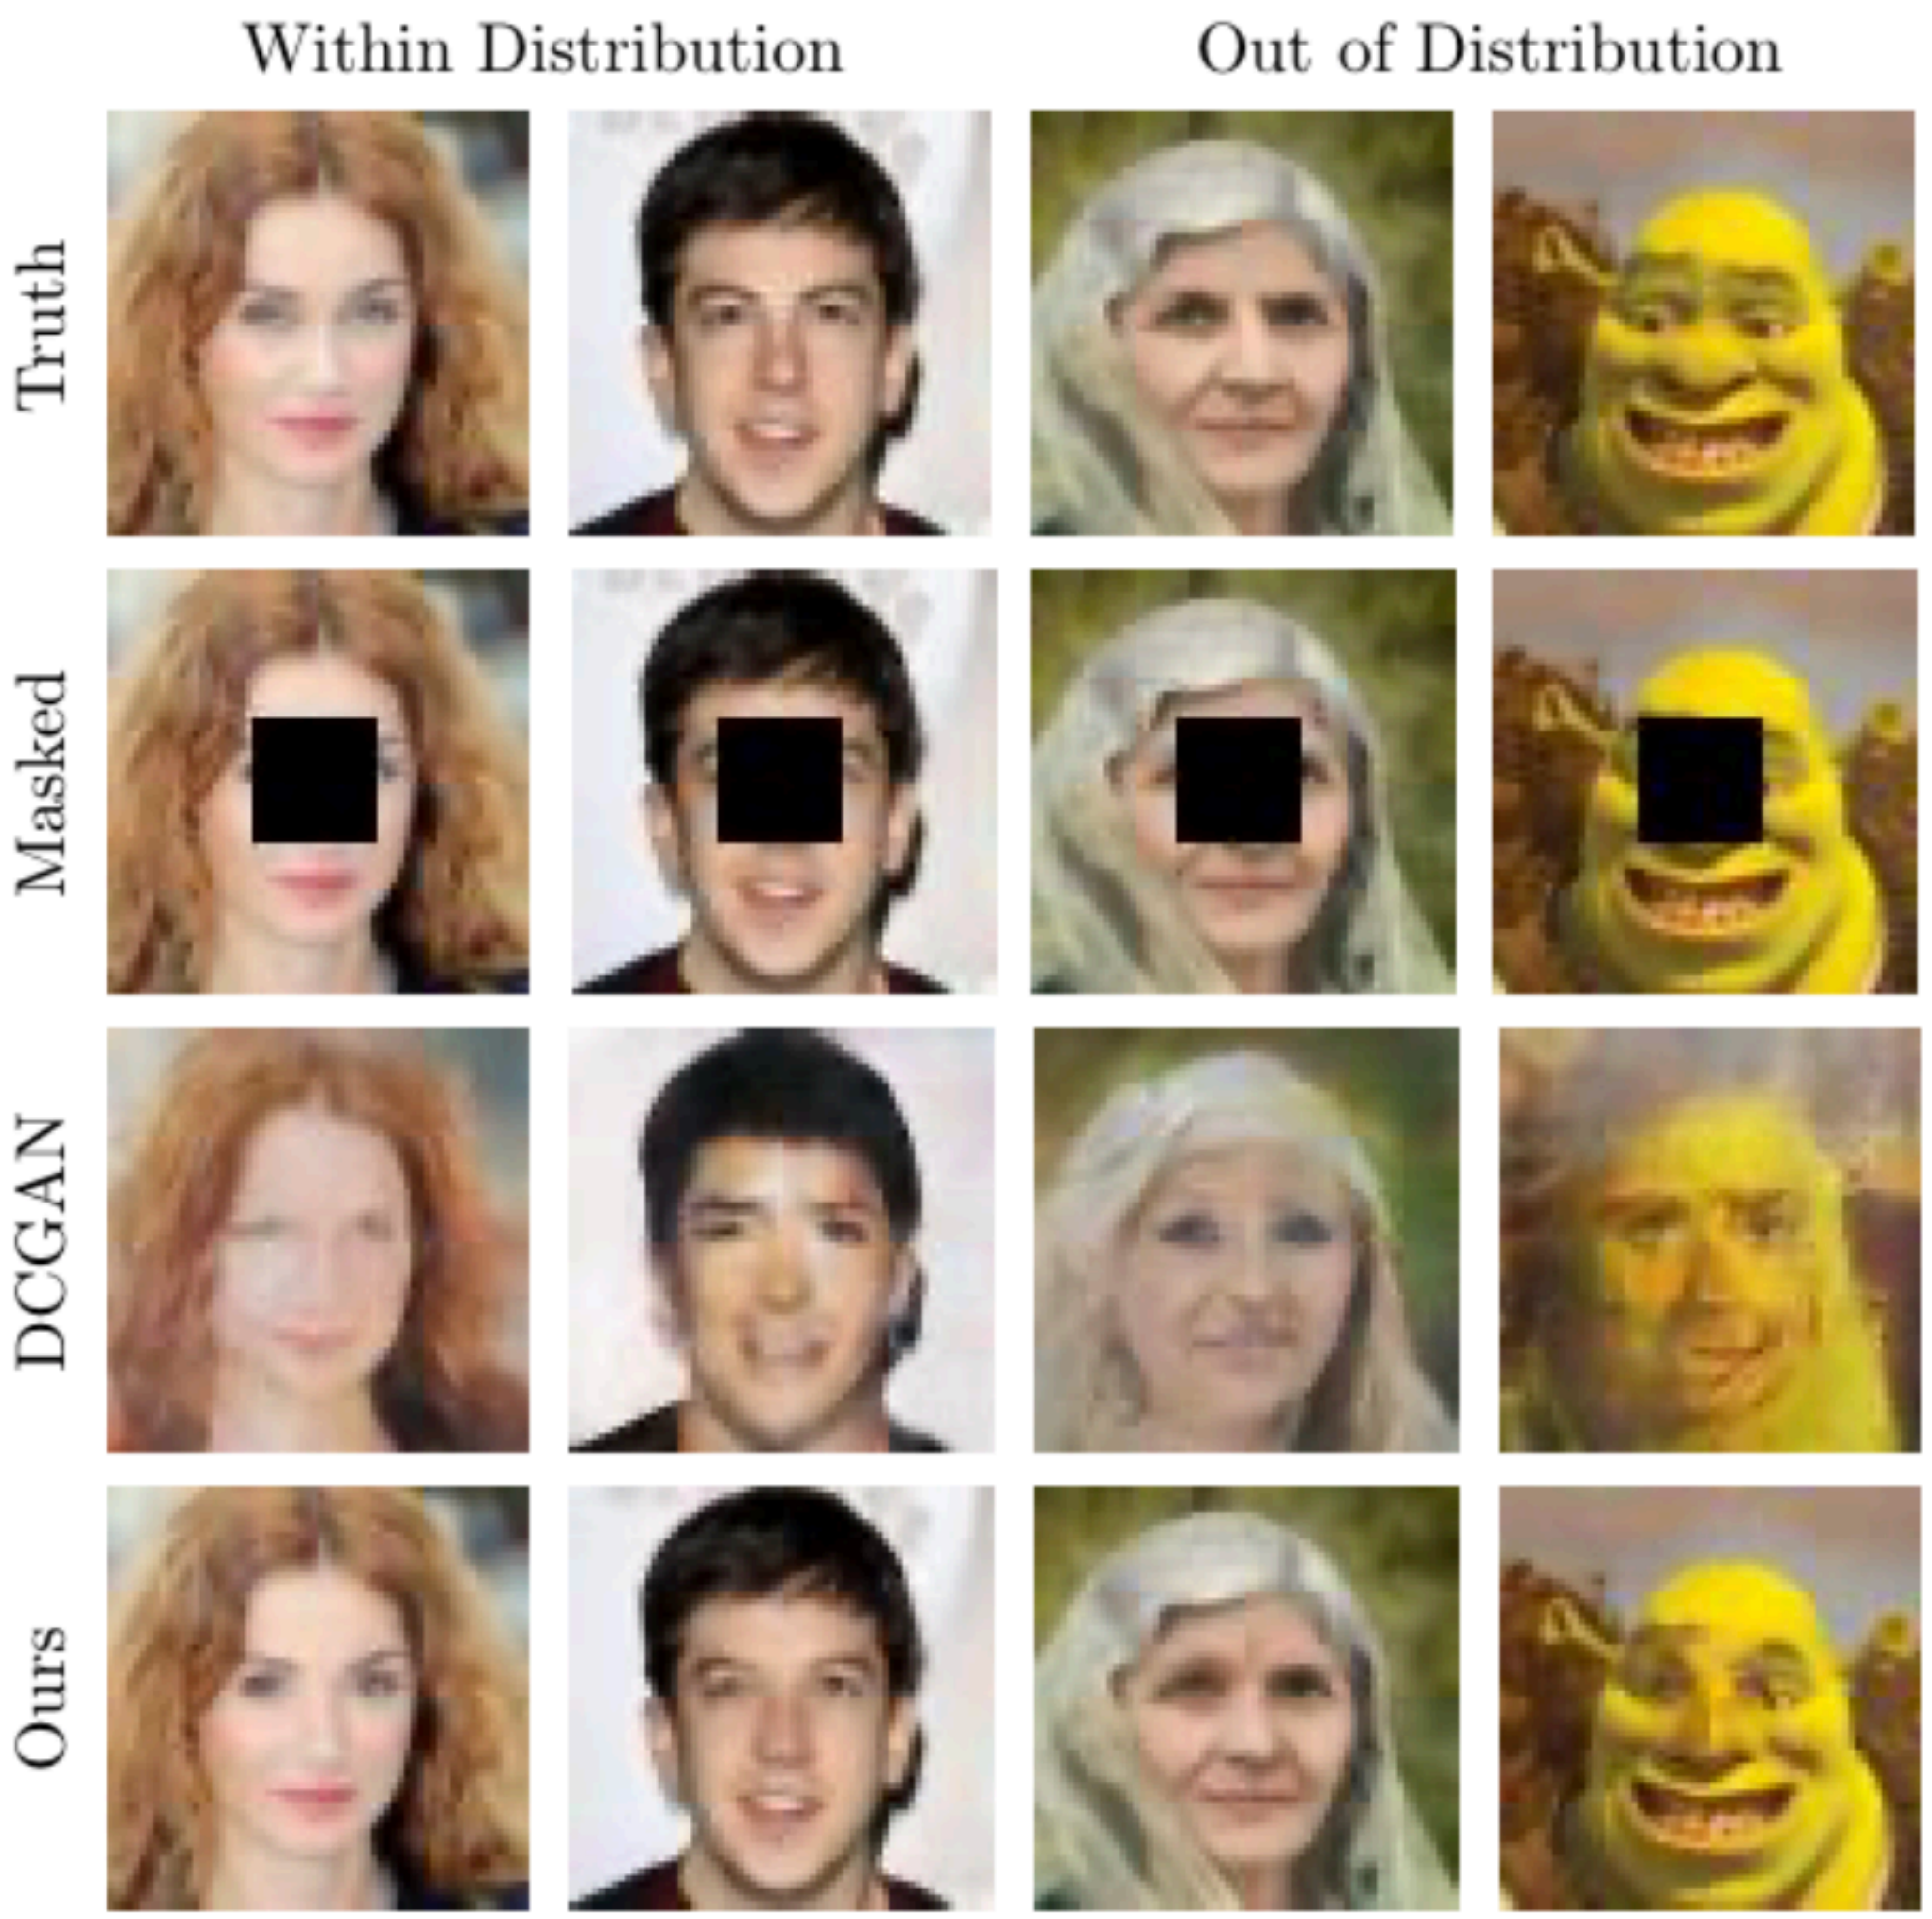

Theory for Linear Invertible Model

Theorem: Let $G \in \mathbb{R}^{n \times n}$ with $\sigma_{\min} > 0$. Given *m* Gaussian measurements Ax_0, the MLE estimator

$$\hat{x} := \arg\max_{x \in \mathbb{R}^n} p_G(x) \text{ s.t. } Ax = Ax_0$$

obeys

$$\sum_{i>m} \sigma_i^2 \leqslant \mathbb{E}_A \mathbb{E}_{x_0} \|\hat{x} - x_0\|^2 \leqslant m \sum_{i>m-2} \sigma_i^2 .$$

Discussion

Why do INNs perform so well OOD?

Invertibility guarantees zero representation error

Where does regularization occur?

Explicitly by penalization or implicitly by initialization + optimization

When is regularization helpful in CS?

$$\min_{z \in \mathbb{R}^n} \|AG(z) - y\|^2 + \gamma \|z\|^2$$

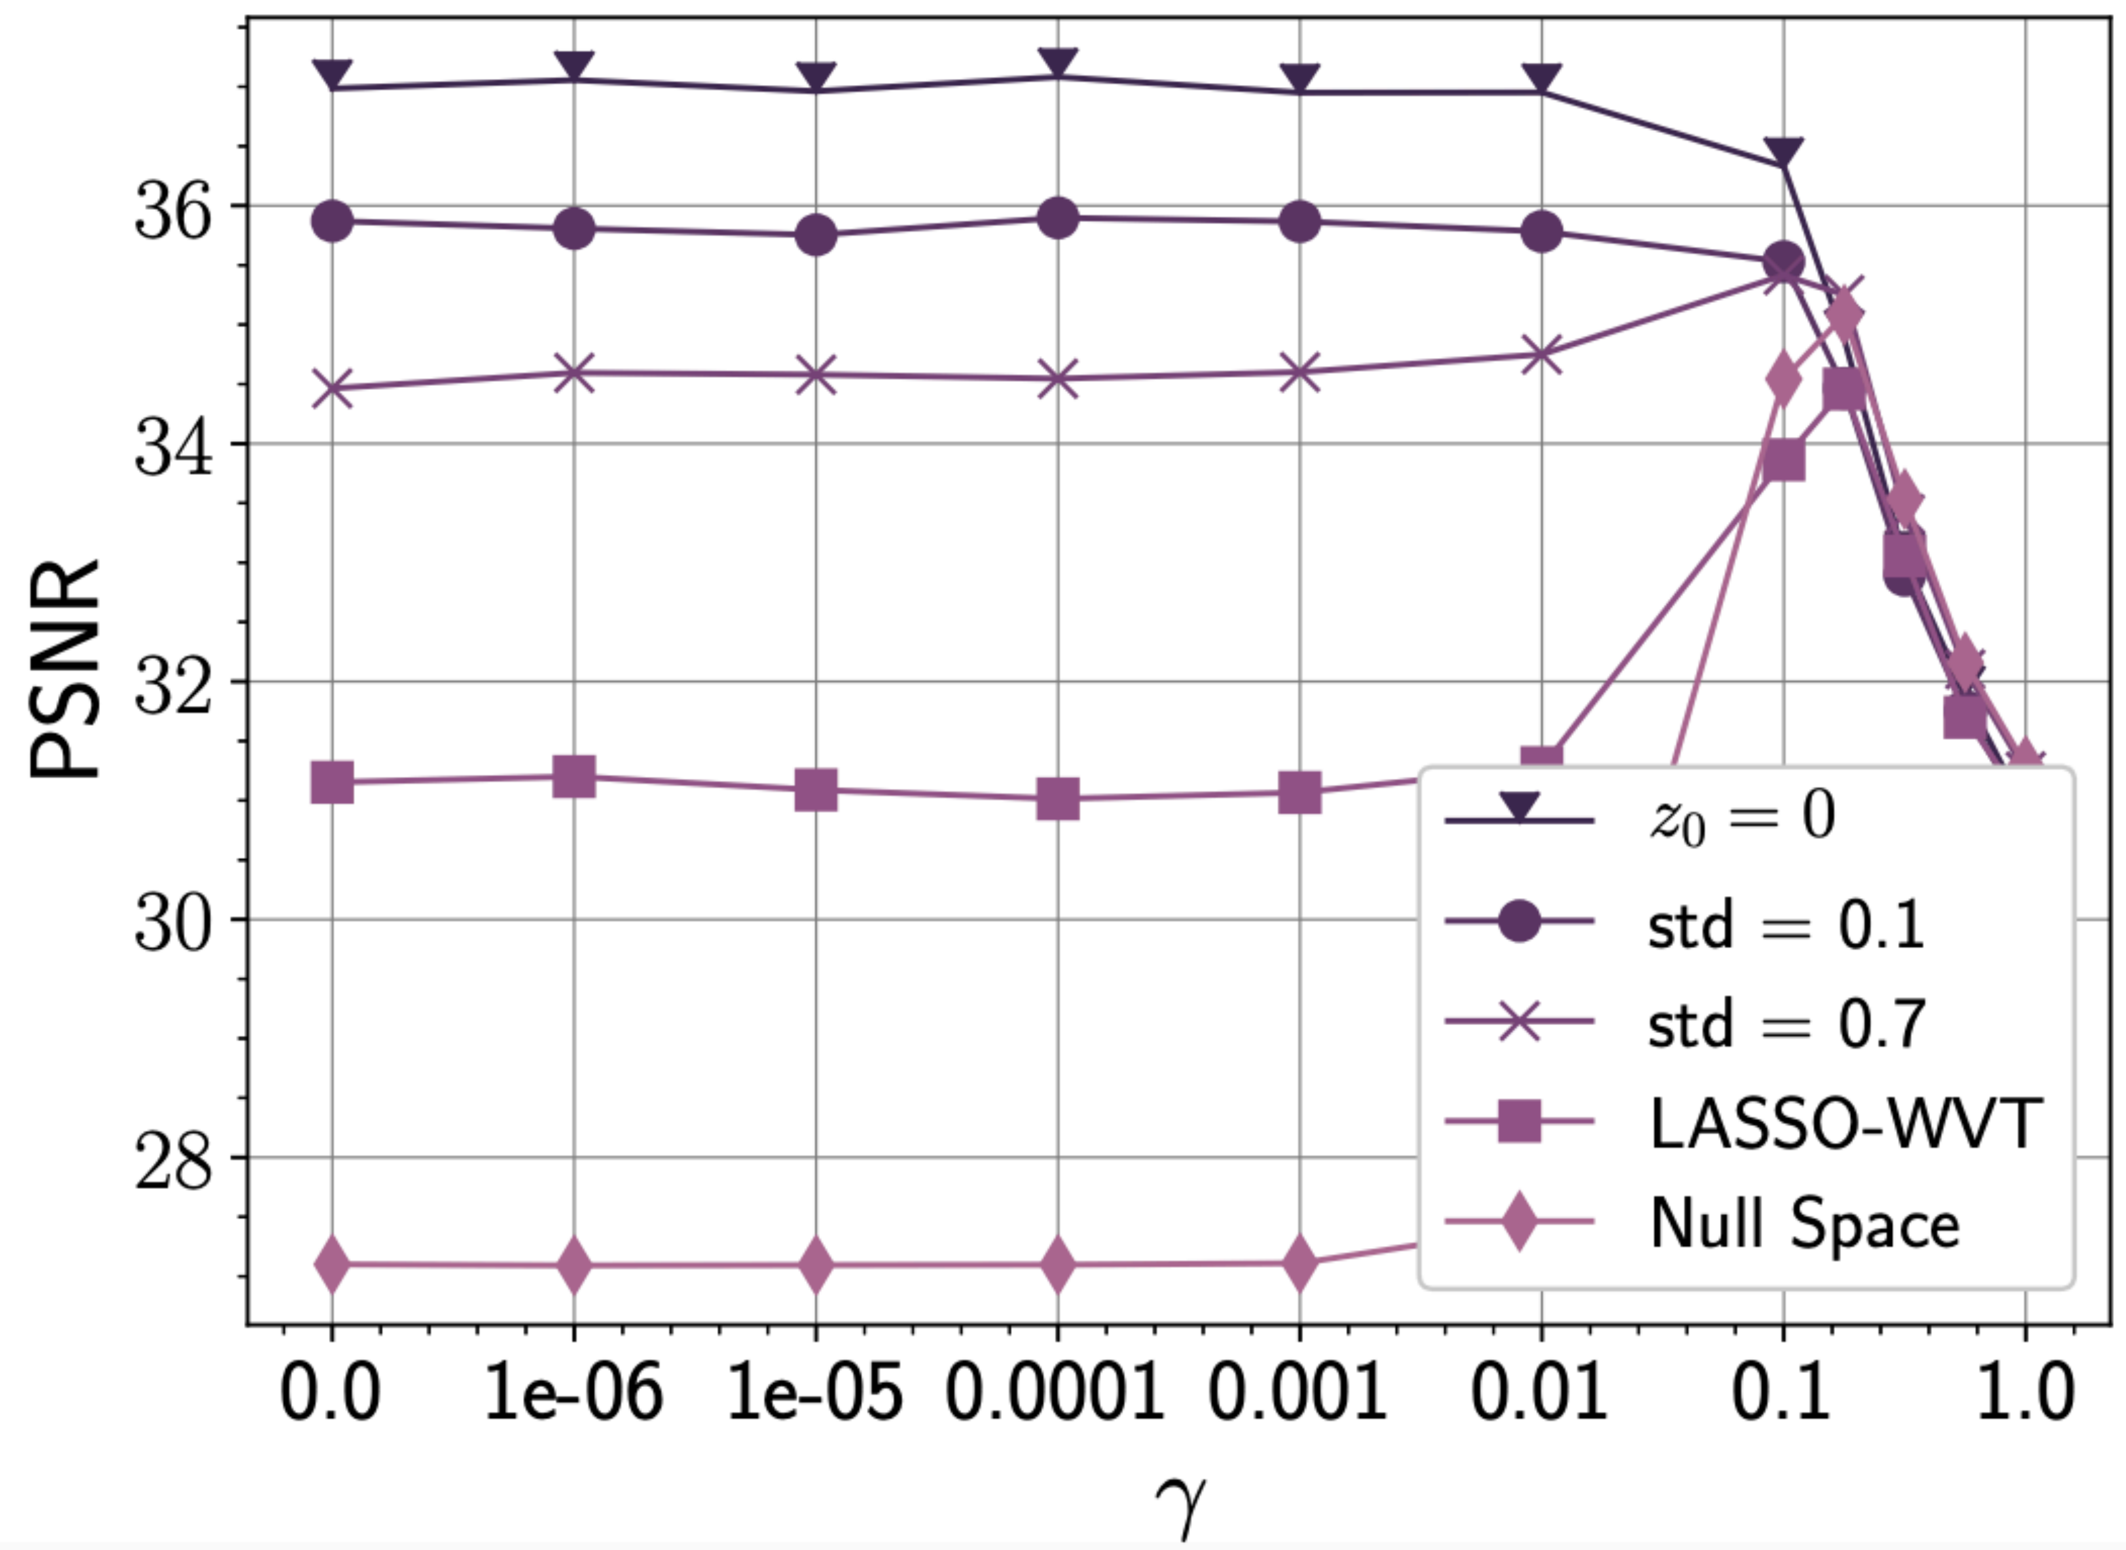

High likelihood init

↓

Regularization by init + opt alg

Low likelihood init

↓

Explicit regularization needed

Why is likelihood in latent space a good proxy?

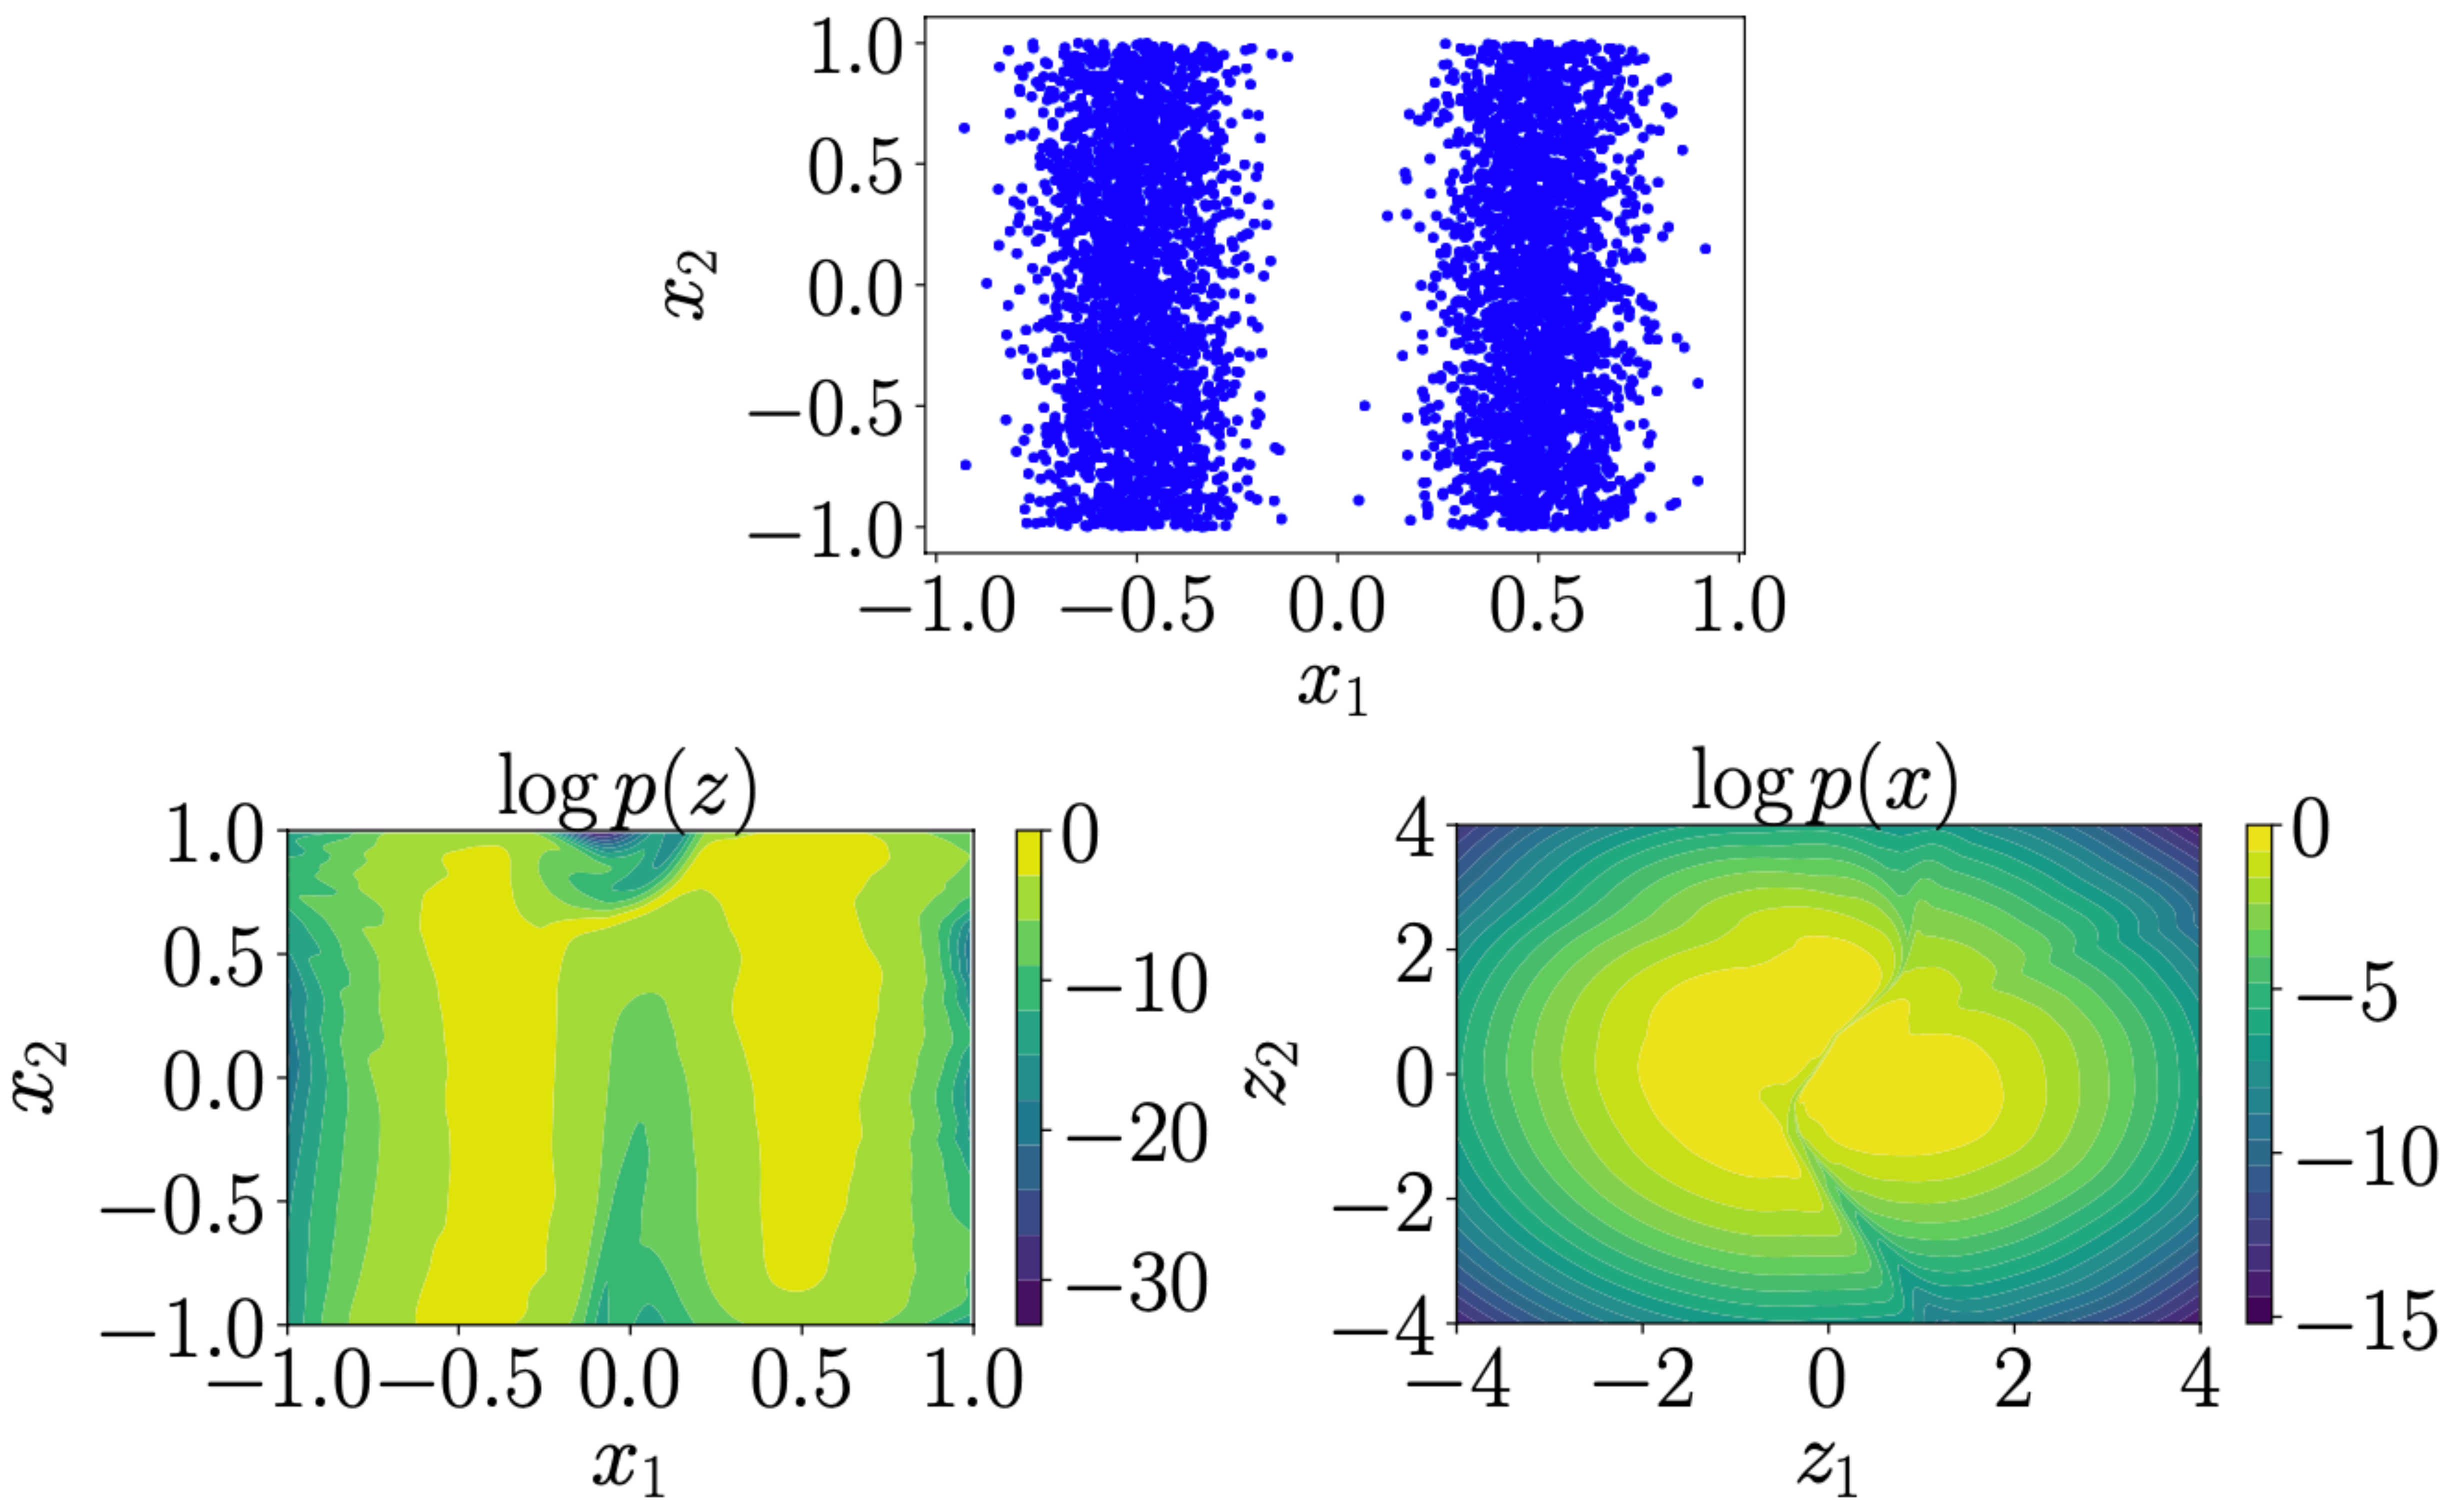

High likelihood regions in latent space generally correspond to high likelihood regions in image space

Why is likelihood in latent space a good proxy?

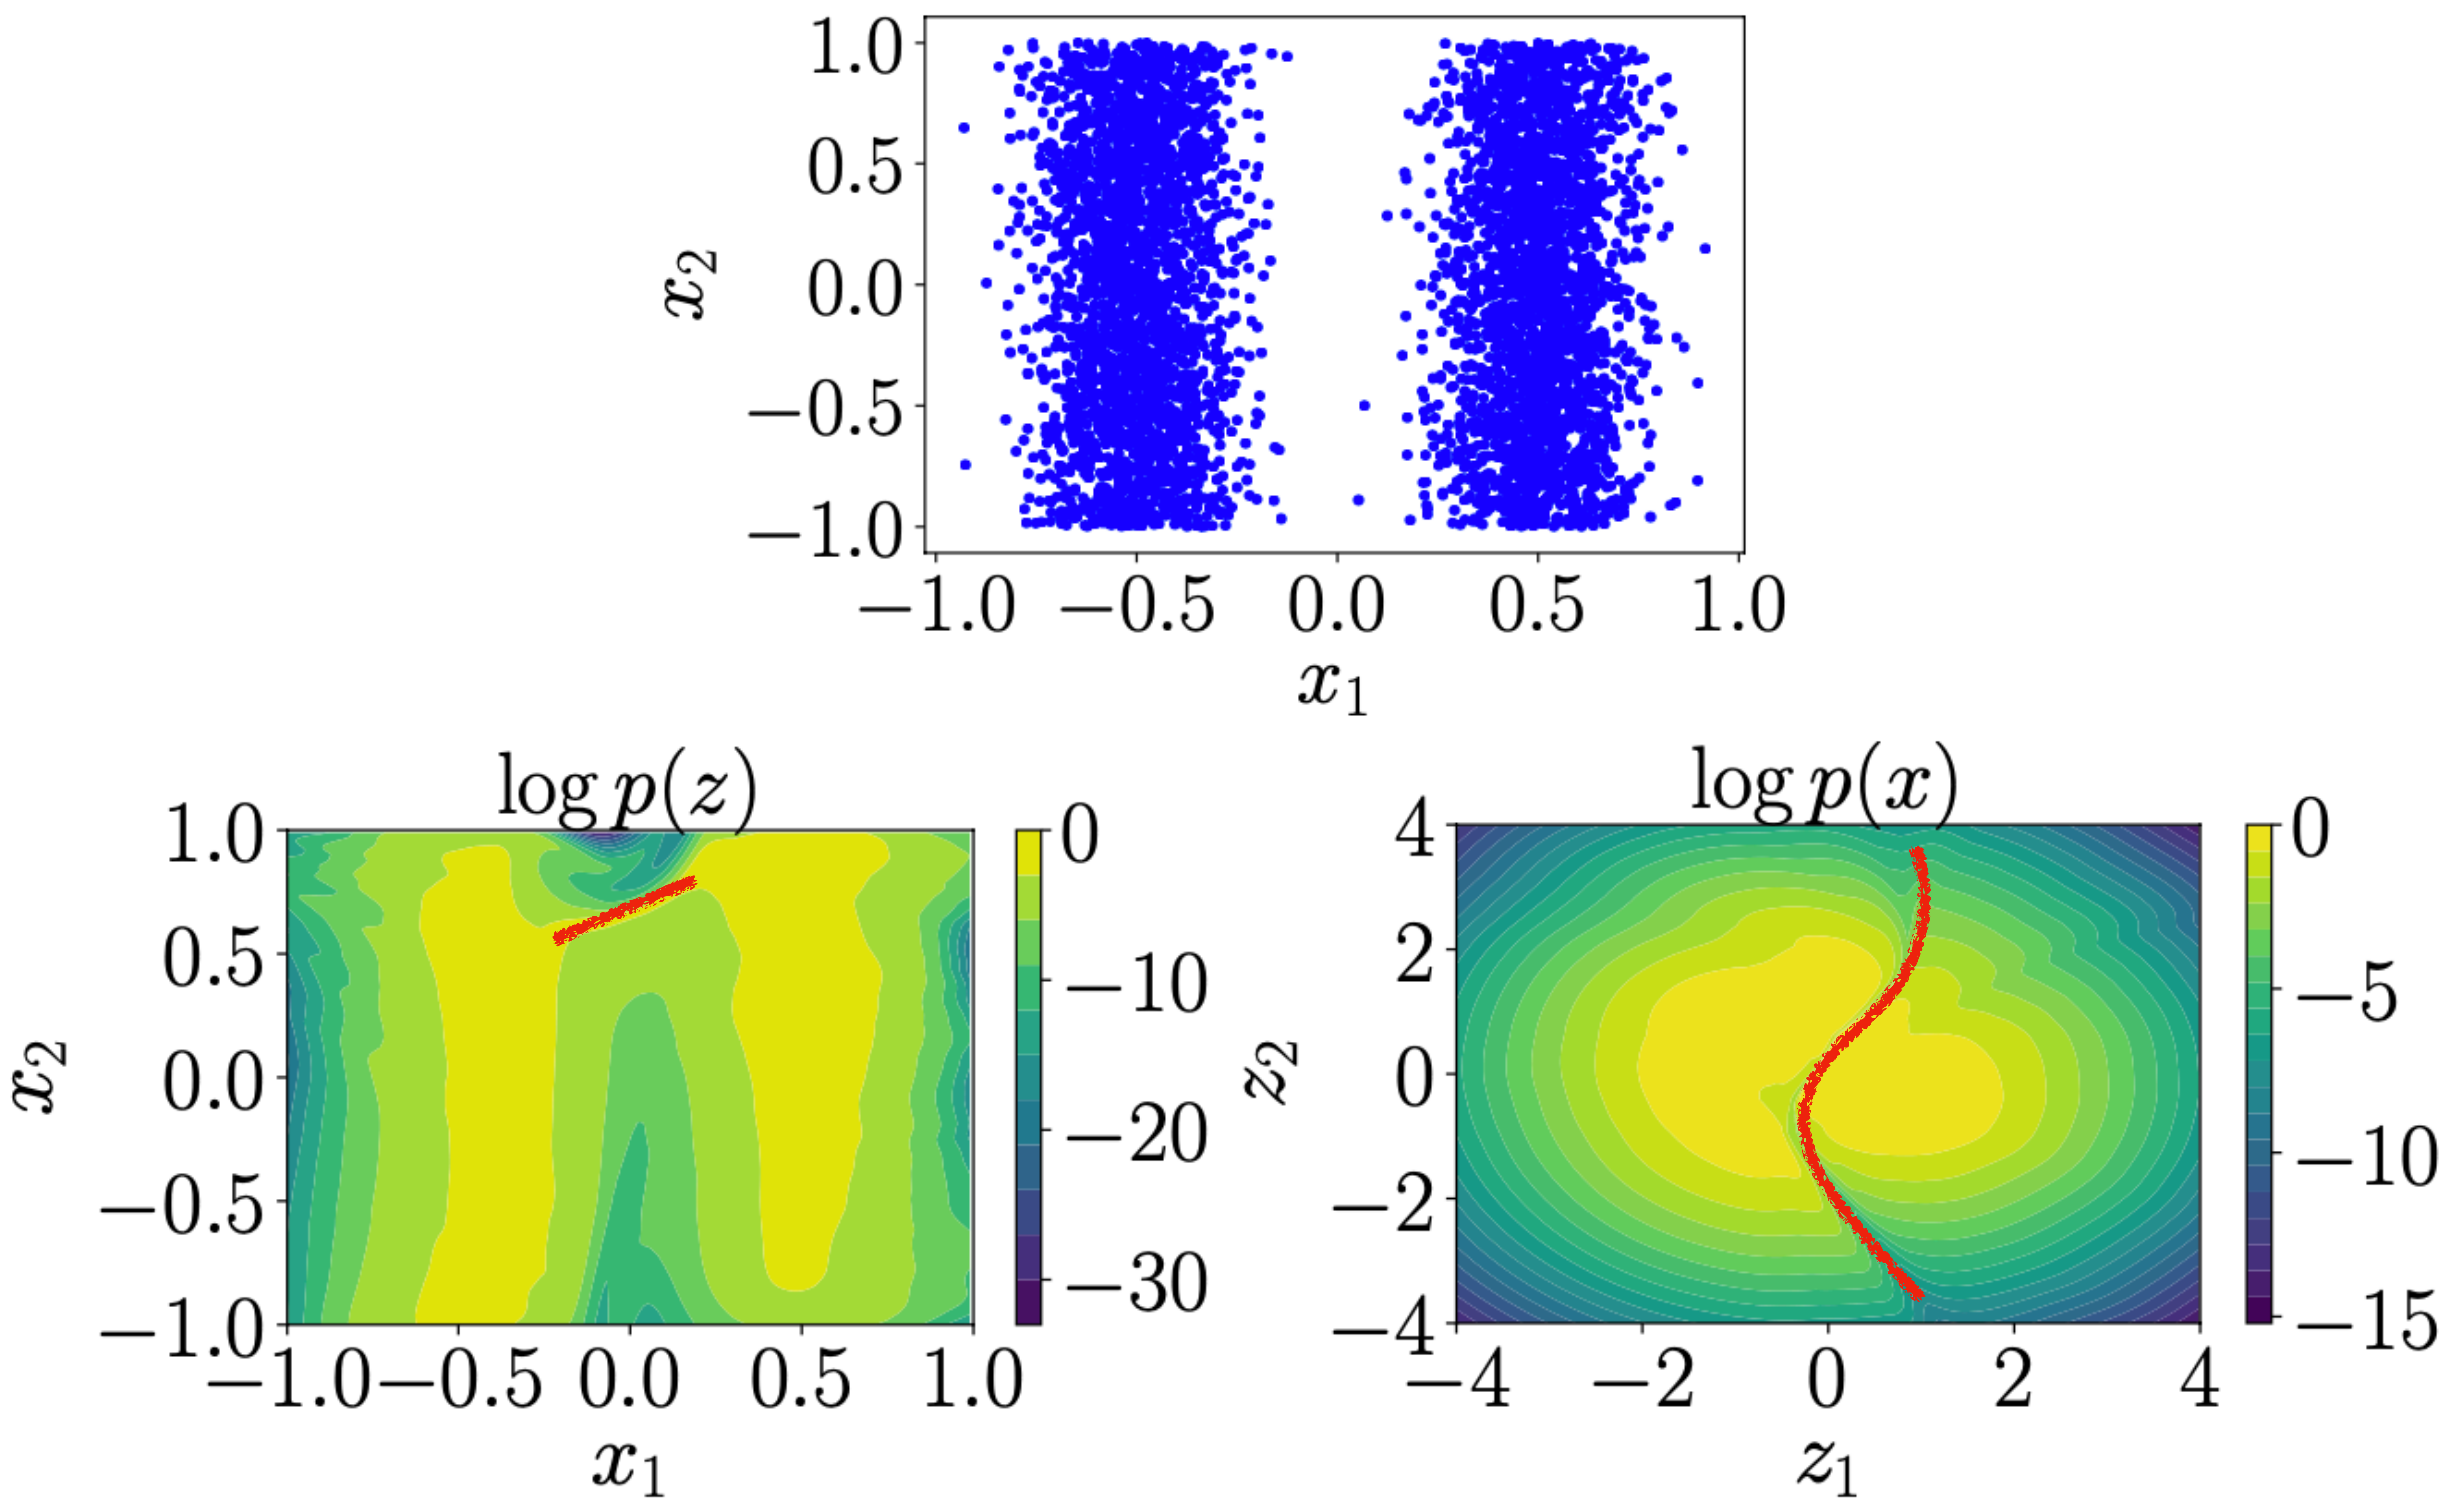

High likelihood regions in latent space generally correspond to high likelihood regions in image space

Contributions

1. Trained INN priors provide SOTA performance in a variety of inverse problems

2. Trained INN priors exhibit strong performance on out-of-distribution images

3. Theoretical guarantees in the case of linear invertible model