UNIVERSITÄT BERN

ICCV HONOLULU HAWAII OCT 19-23, 2025

Diffusion Image Prior

Hamadi Chihaoui, Paolo Favaro

Computer Vision Group, University of Bern, Switzerland

Image Restoration

- Structured Degradations
  - Fully Known or Modeled Operators
  - Examples: Gaussian deblurring, super-resolution..

Gaussian blur

Low resolution

Heavily explored

- Unstructured Degradations
  - Complex or Unknown operators
  - Examples: Water drop, non-uniform deformation..

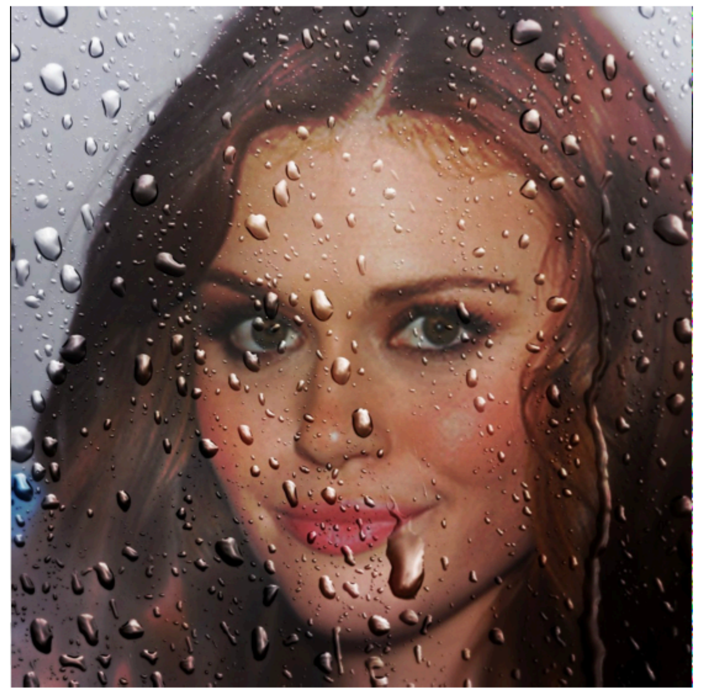

Water drop

Non-uniform deformation

Rarely explored

Tackling Unstructured Degradation

- Supervised methods
  - need an annotated dataset which is hard to collect
- Training-based methods
  - Rely on a dataset and usually fail to generalize on new degradation patterns

- Unsupervised methods
  - Degradation operator fully known
- Training-free methods
  - Applied directly to the input, no generalization problem.

Approach

We introduce a new image restoration method that:

1. does not assume any prior knowledge of the degradation model.
2. is fully unsupervised.
3. is training-free (leverages a pre-trained model).
4. shows higher fidelity to the input compared to existing methods.

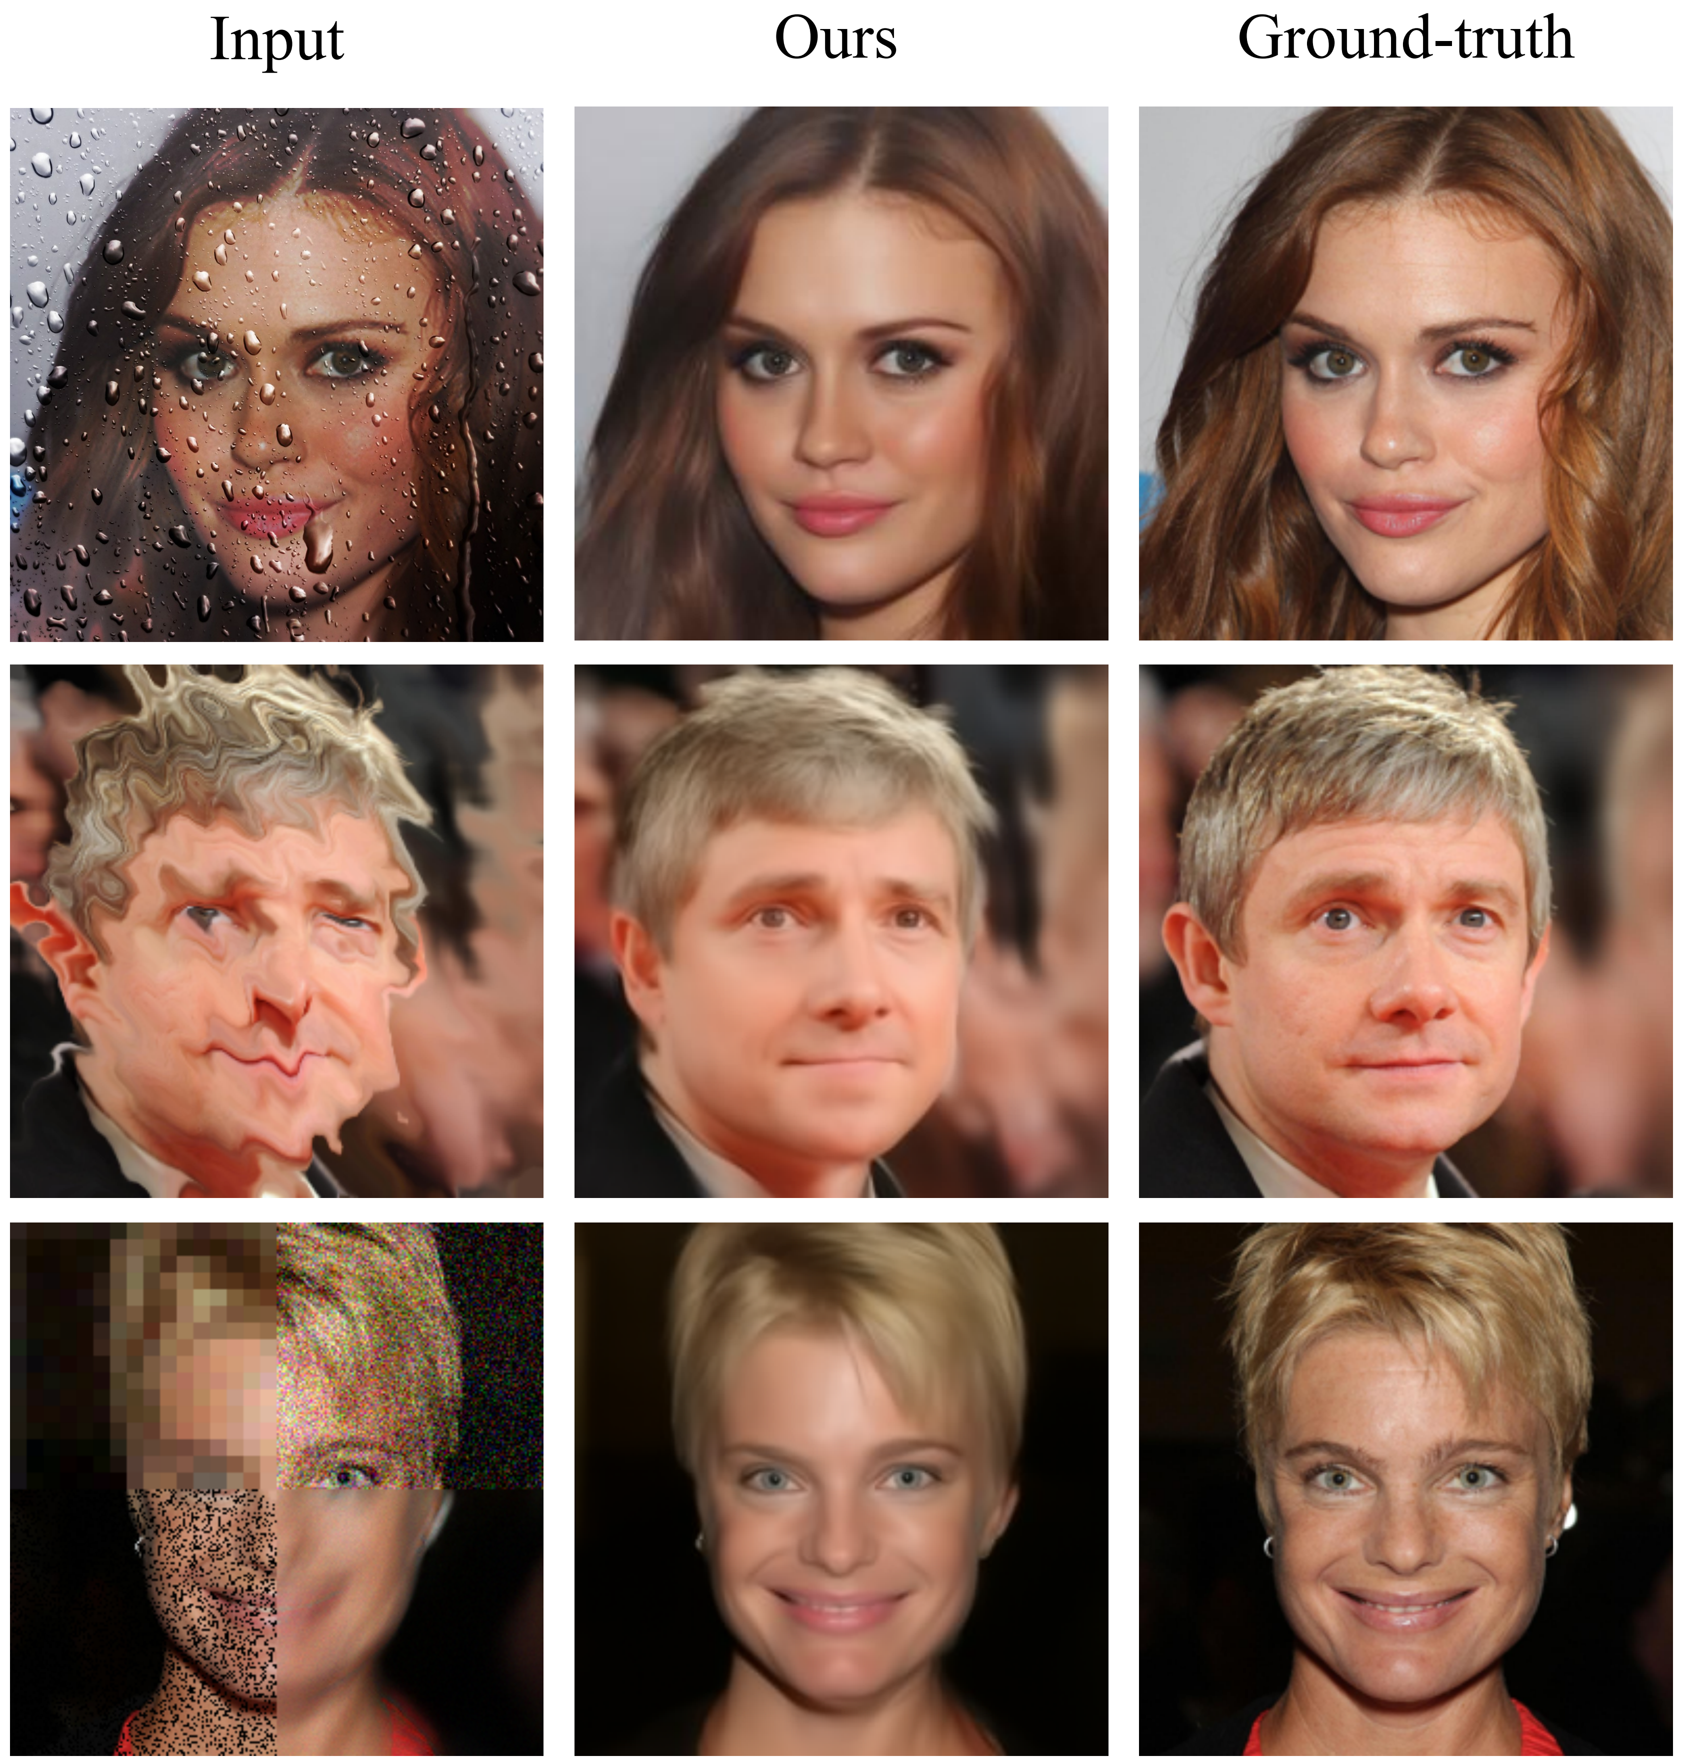

Inspiration and Motivation

Deep Image Prior (DIP) [Ulyanov et al. CVPR 2018]

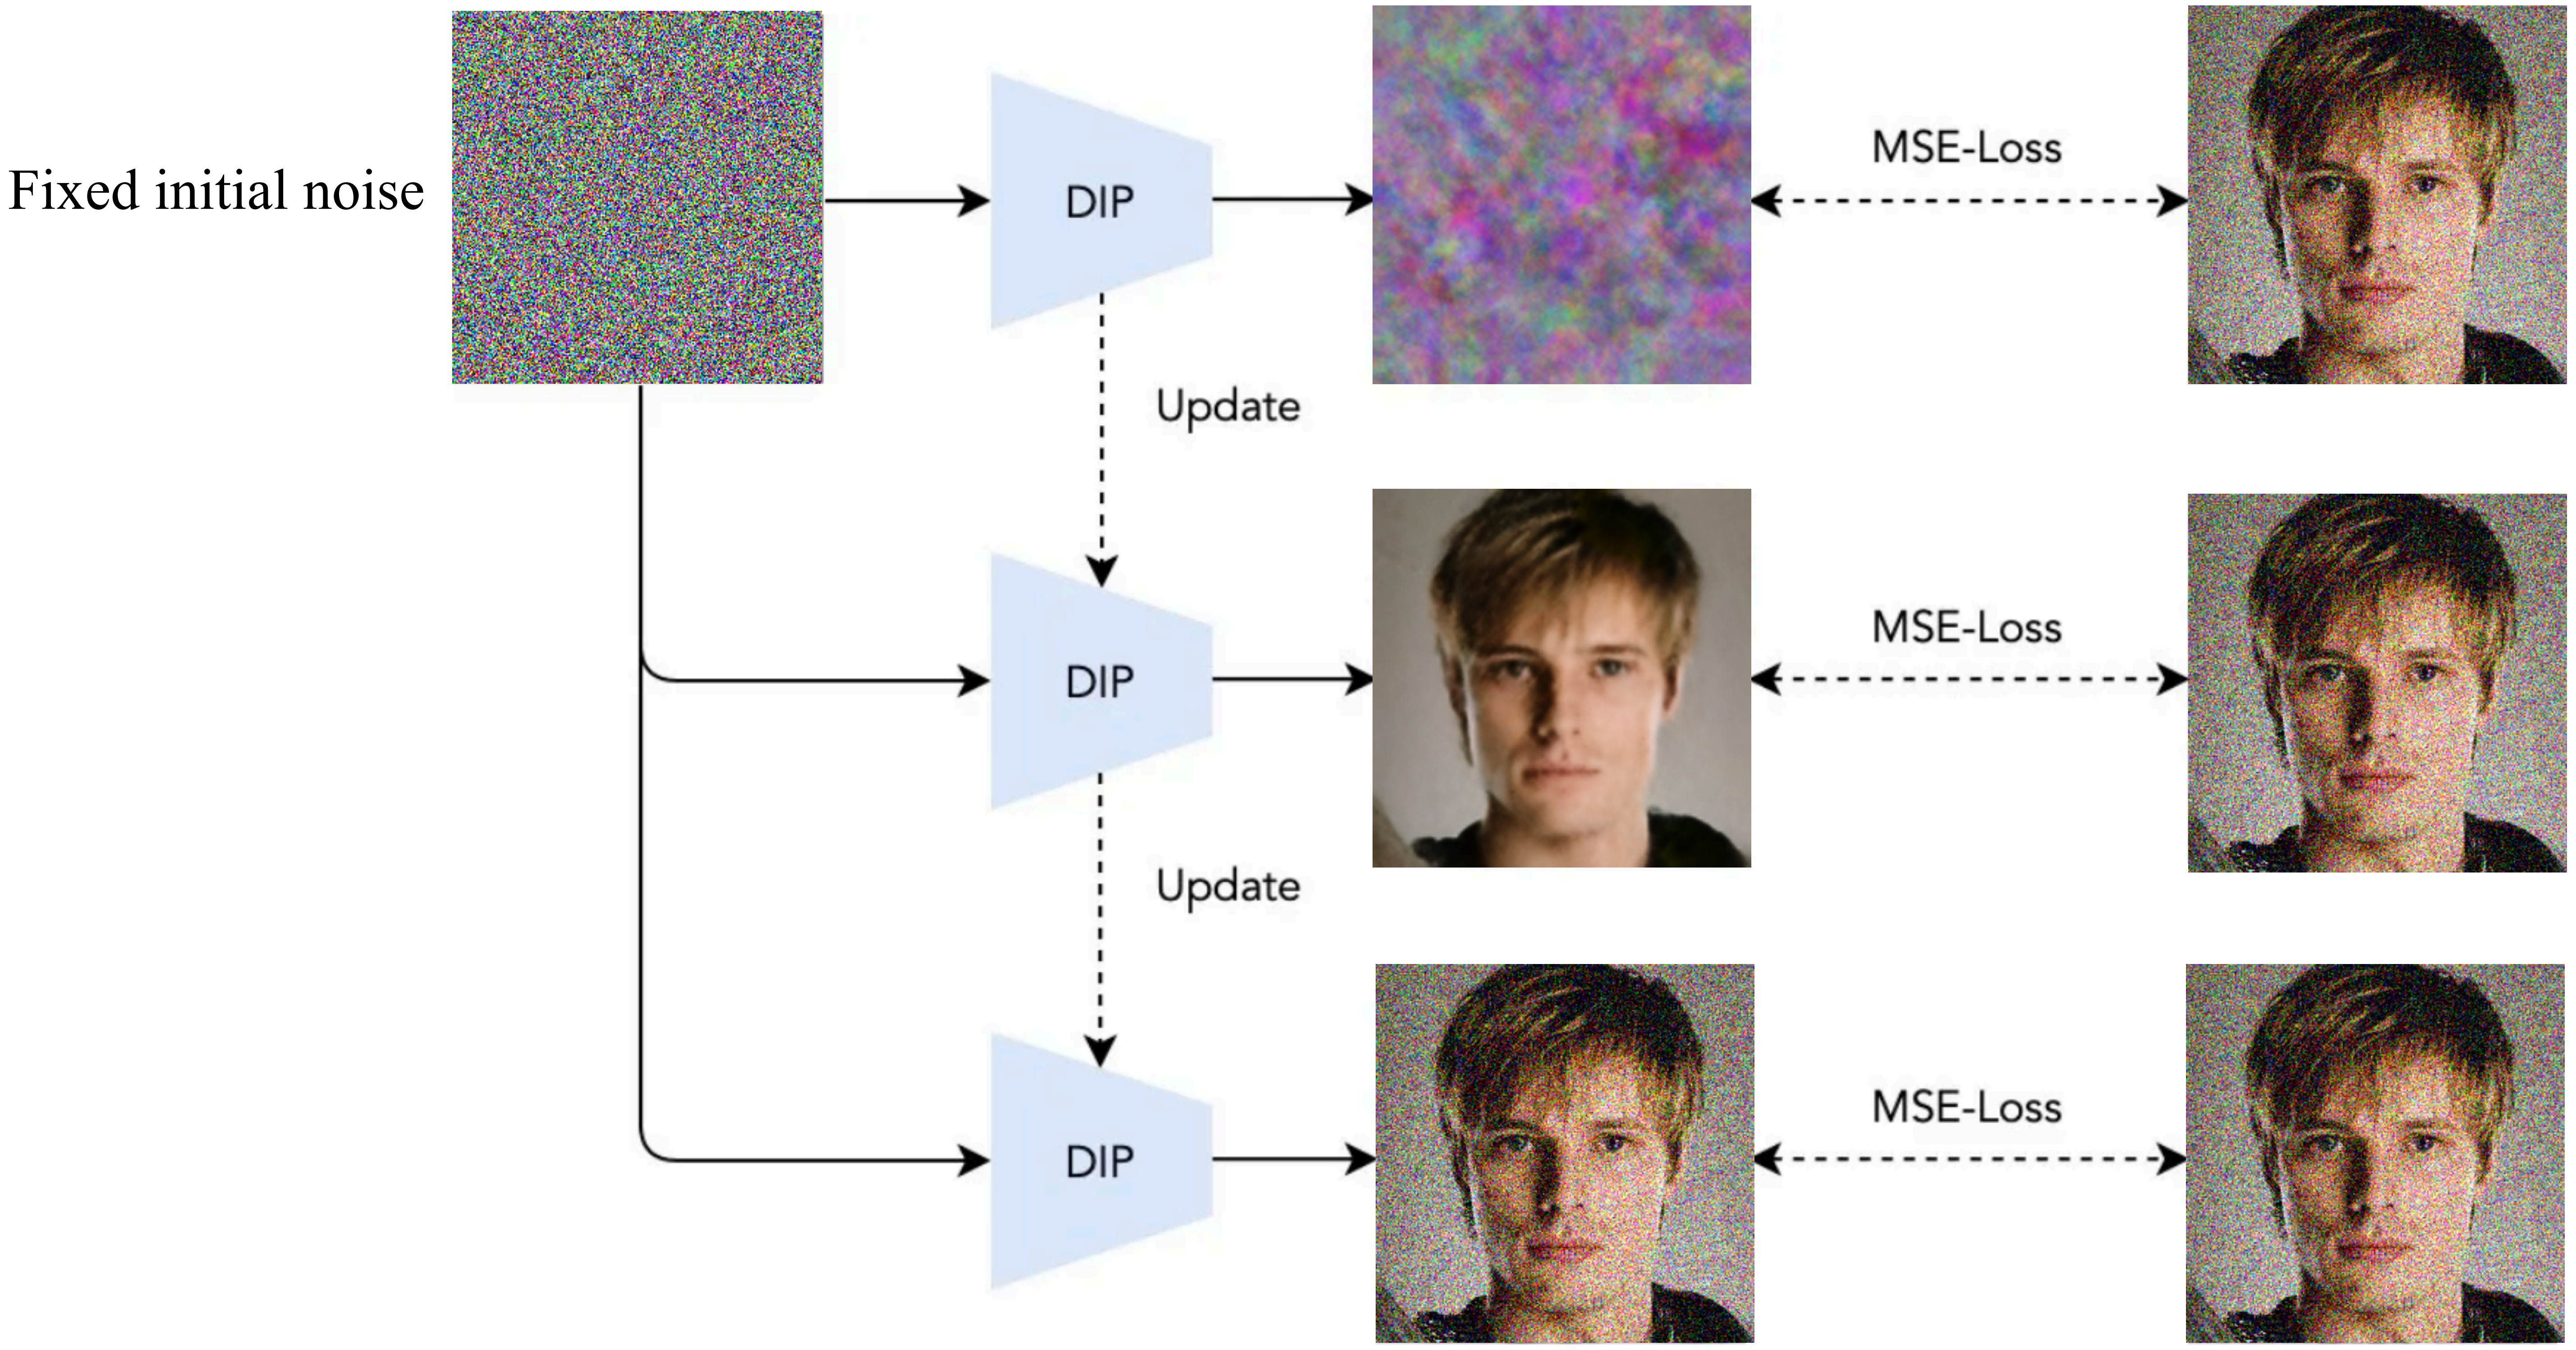

Visual illustration of Deep Image Prior (DIP)

Deep Image Prior (DIP)

- DIP optimization:

$$\theta^* = \arg\min_{\theta} ||f_{\theta}(z) - y||^2 \qquad (1)$$

$$\hat{x} = f_{\theta^*}(z)$$

where y is the input image, f_θ is an untrained CNN, and z is an initial noise.

- However, an untrained deep CNN is a too weak prior.

- What about using a stronger prior?

DIffusion Image Prior (DIIP)

Proposed optimization:

$$z^* = \arg\min_{z} ||g(z) - y||^2 \qquad (2)$$

$$\hat{x} = g(z^*)$$

where y is the input image, g is a pre-trained diffusion model, and z is noise.

Algorithm 1 Iterative Solver of Eq (2)

Require: Degraded image y, pre-trained diffusion g, learning rate α.
Ensure: Return $\hat{x}_0$

Initialize $z_T^0 \sim \mathcal{N}(0, \mathbf{I})$
for $k : 1 \rightarrow N-1$ **do**
  $x_0^k = g(z_T^k)$
  $z_T^{k+1} = z_T^k - \alpha \nabla_{x_T} ||x_0^k - y||^2$
end for
return $\hat{x}_0 = g(z_T^N)$

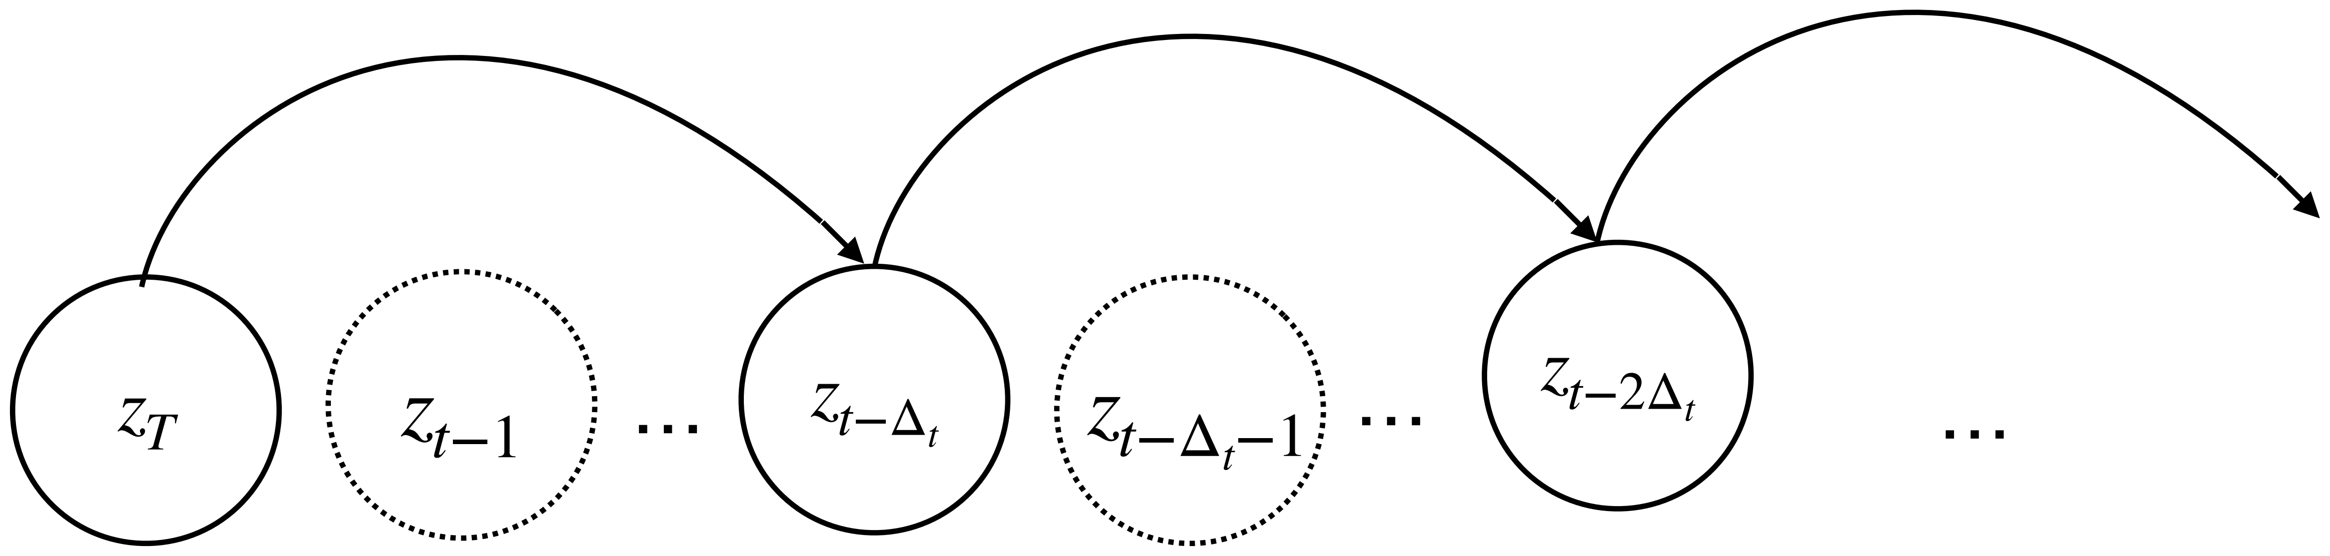

Fast computation of $g(z_T)$

Empirical Study

- Empirical study on synthetic data.
- *Assumption: any degradation either:*
  - *removes some high frequency details.*
  - *adds some high frequency artifacts.*

- Test DIP and our proposed optimization on blurry images and noisy ones.

Case of Noise

<u>DIP</u>

$$\theta^* = \arg\min_{\theta} \|f_\theta(z) - y\|^2 \qquad (1)$$

$$\hat{x} = f_{\theta^*}(z)$$

Ground-truth

Input

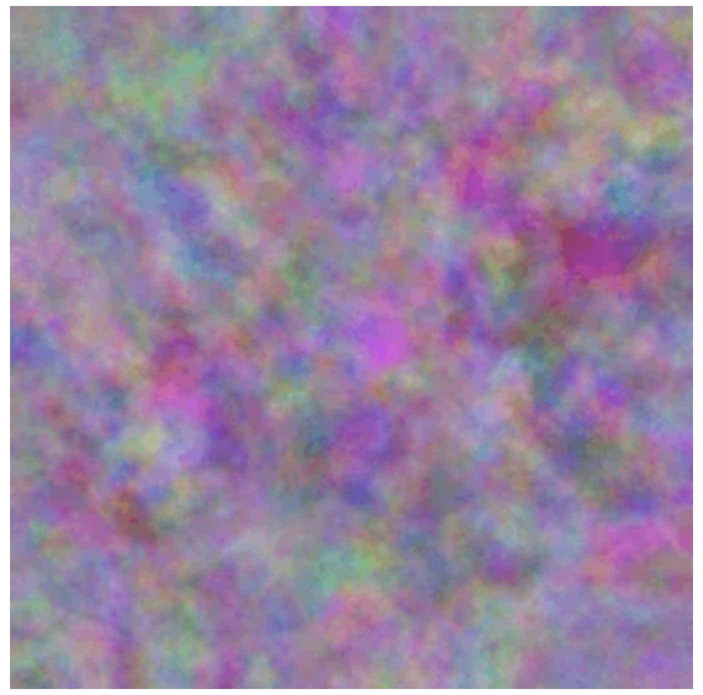

$\hat{x}$

<u>DIIP</u>

$$z^* = \arg\min_{z} \|g(z) - y\|^2 \qquad (2)$$

$$\hat{x} = g(z^*)$$

Ground-truth

Input

$\hat{x}$

Case of Blur

DIP

$$\theta^* = \arg\min_{\theta} \|f_\theta(z) - y\|^2 \qquad (1)$$

$$\hat{x} = f_{\theta^*}(z)$$

Ground-truth

Input

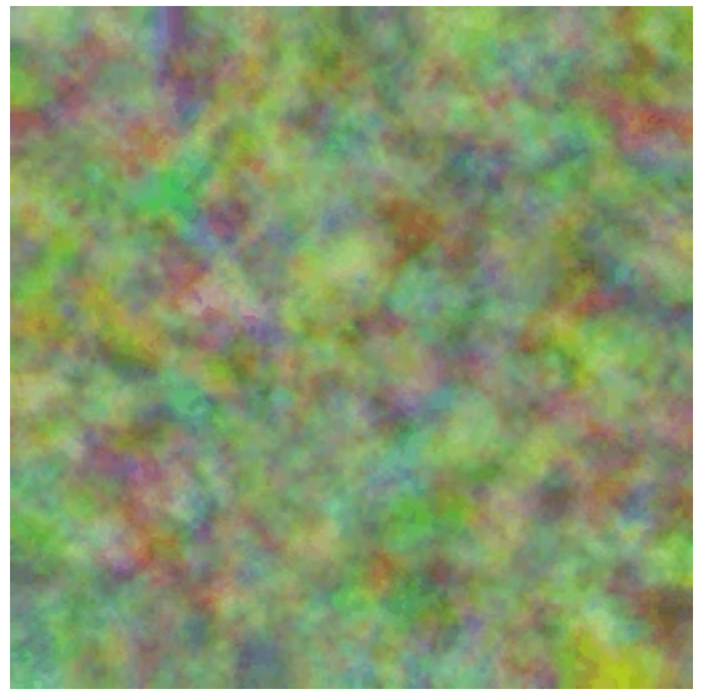

$\hat{x}$

DIIP

$$z^* = \arg\min_{z} \|g(z) - y\|^2 \qquad (2)$$

$$\hat{x} = g(z^*)$$

Ground-truth

Input

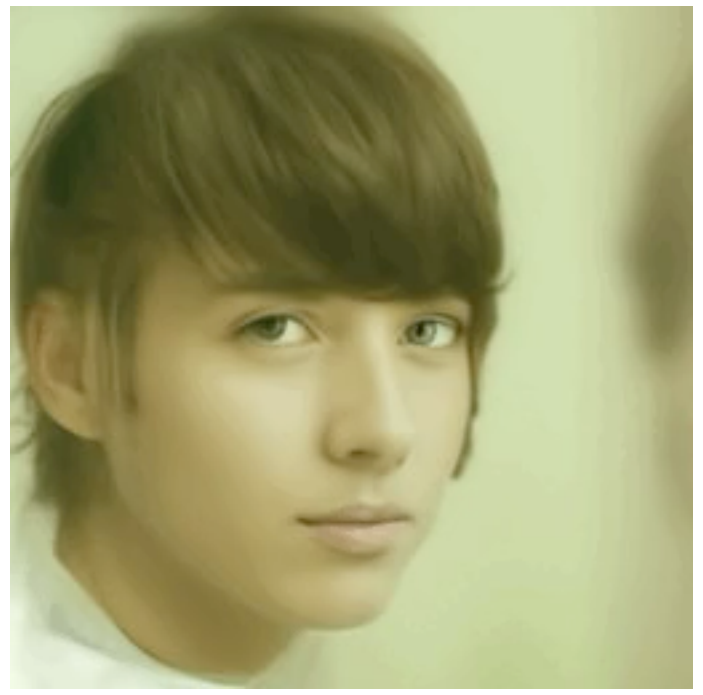

$\hat{x}$

Two Key Findings

1. Two regimes regardless of the degradation type.

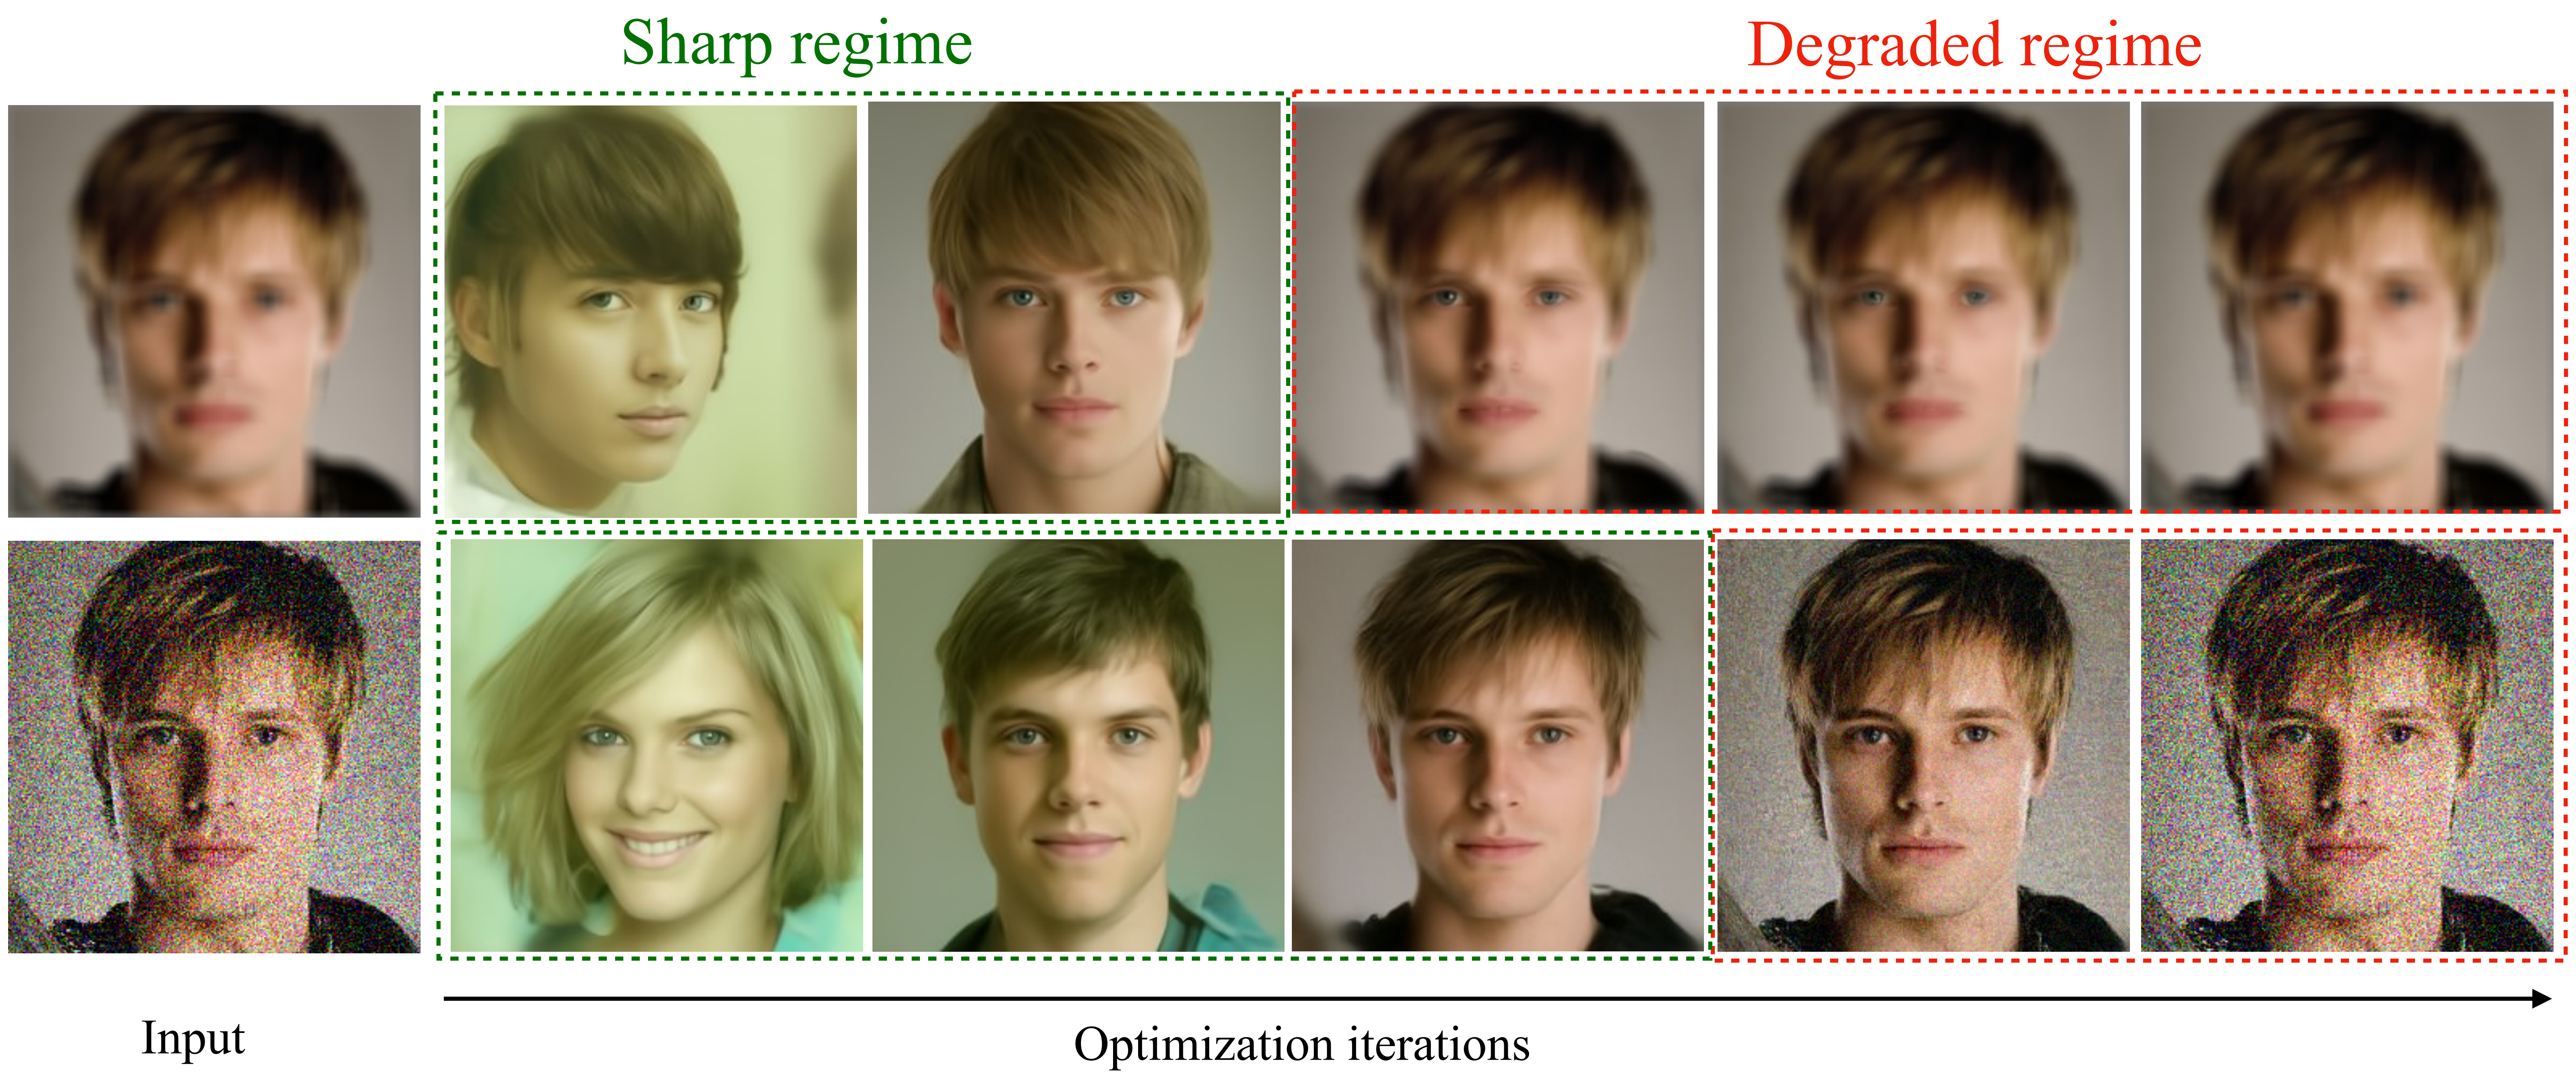

2. Higher resistance to high-frequency artifacts (noise) than to low-frequency distortions (blur).

When to Stop the Optimization?

Can we stop our proposed optimization when reaching the optimal reconstruction?

- Early stopping (similar to DIP)
- We propose a self-supervised stopping criterion based on slope of the loss function and the Laplacian of the reconstructed image

When to Stop the Optimization?

Degradations adding high-frequency artifacts

Normalized loss slope: $\Delta_k = \frac{\|g(z_k) - y\|^2 - \|g(z_{k+1}) - y\|^2}{\|g(z_k) - y\|^2}$

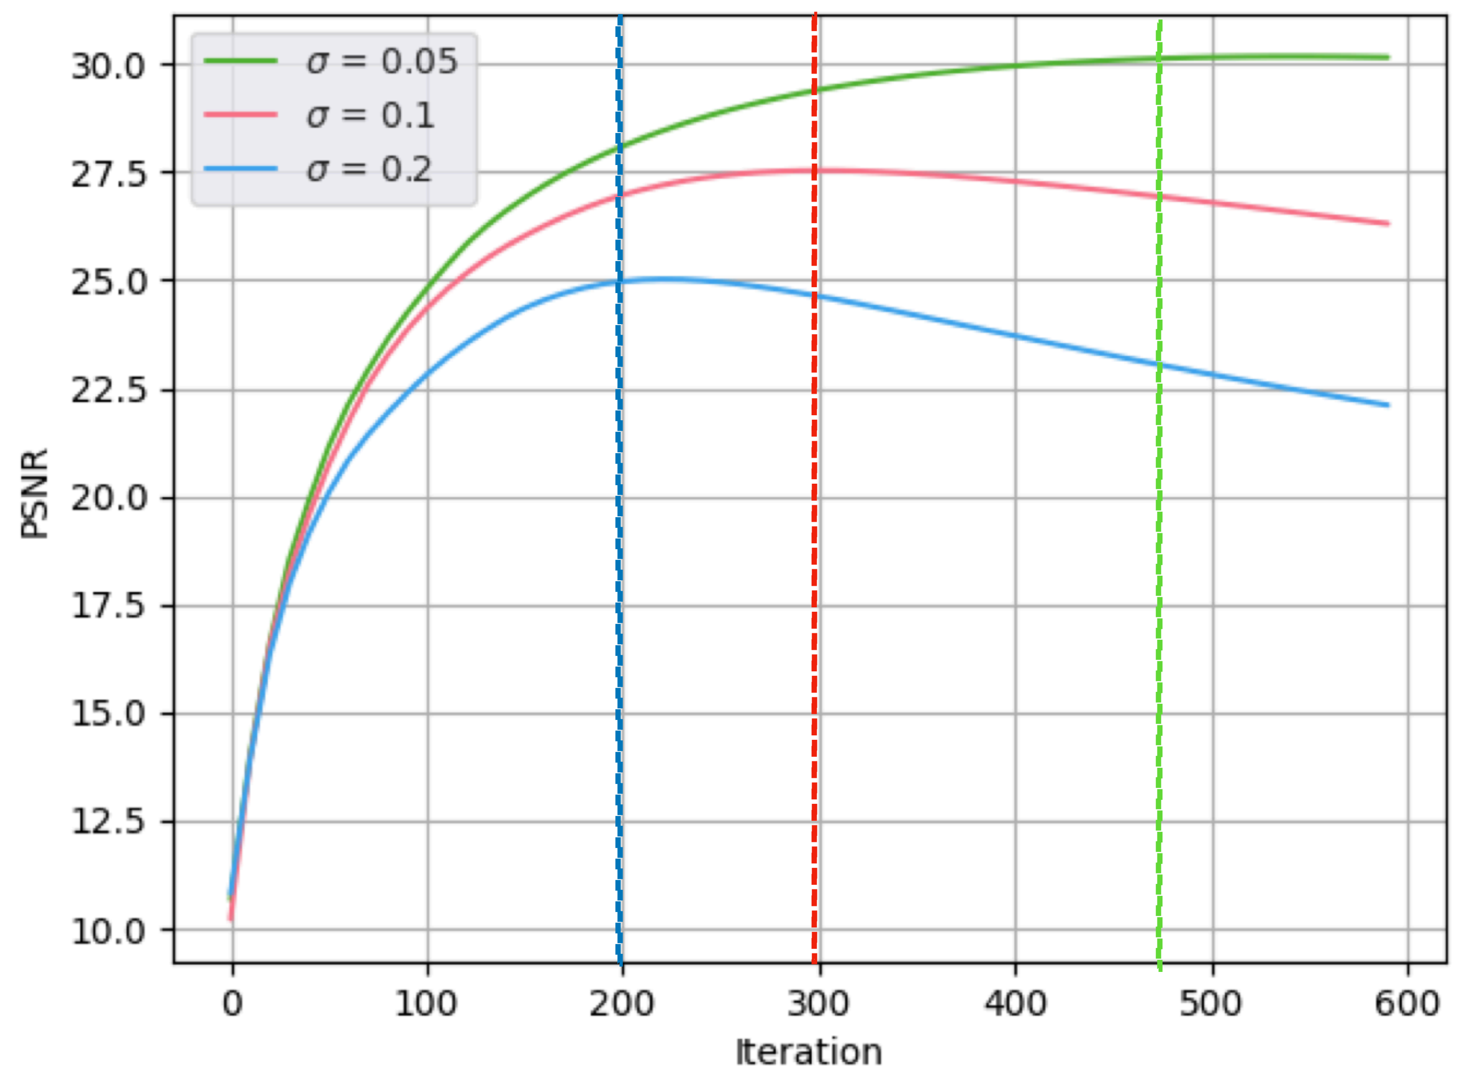

PSNR (w.r.t. ground truth) at different noise levels vs. Iterations

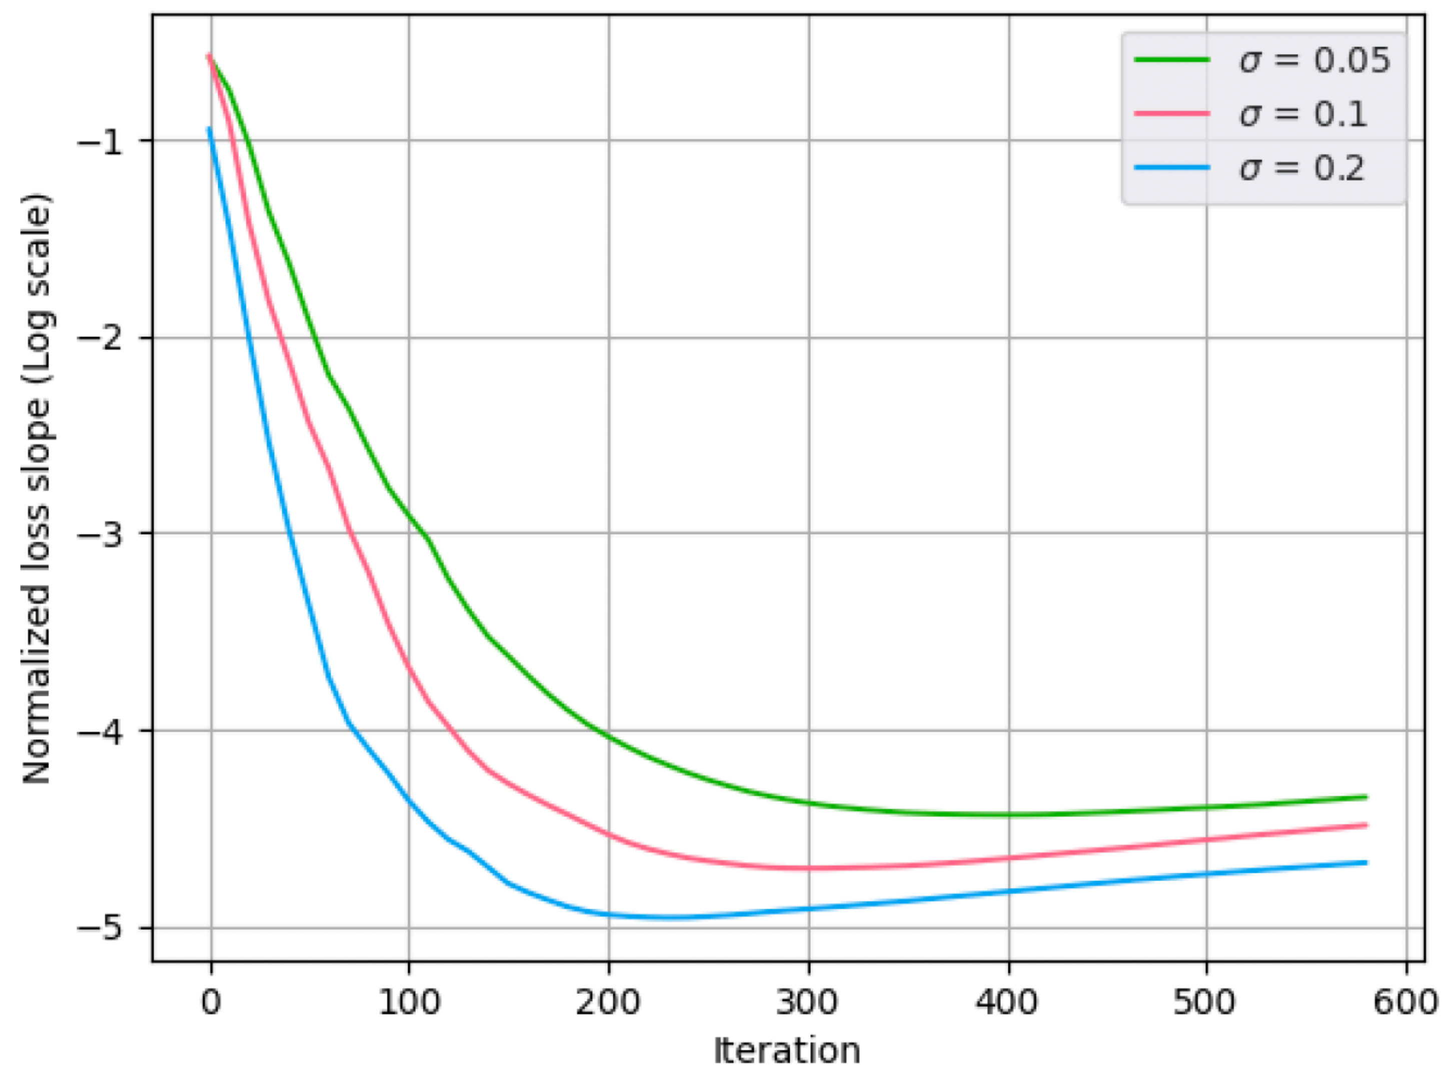

Normalized loss slope Δ_k vs. iterations

(a) The **noisier** an image is, the faster it reaches the optimal reconstruction.

(b) The **noisier** the image, the faster its normalized loss slope drops.

⇒ Stop the optimzation at iteration k if : $\Delta_k < \epsilon$

When to Stop the Optimization?

Degradations removing high-frequency details

- The variance of the Laplacian σ_k^2 is a robust measure of image sharpness. MAPS [Chao et al. BMC bioinformatics 2021]

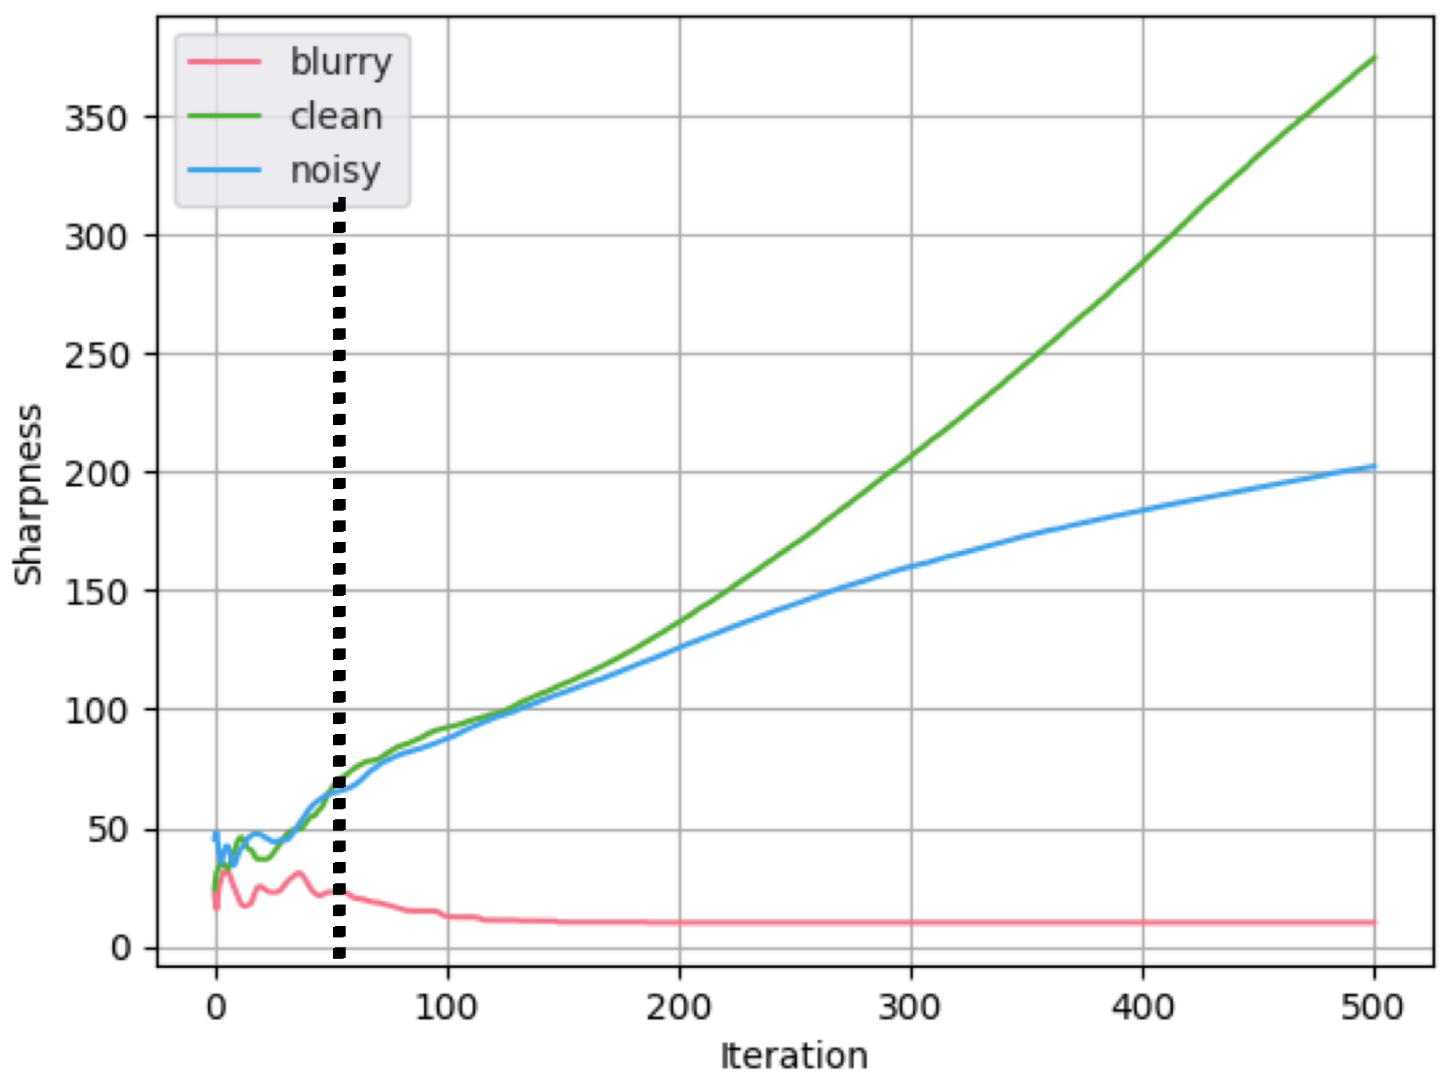

Variance of the Laplacian of reconstructions vs. Iterations

Stop the optimization at iteration k if: $k > k_{min}$ and $\sigma_k^2 > \sigma_{k+1}^2$

DIffusion Image Prior (DIIP)

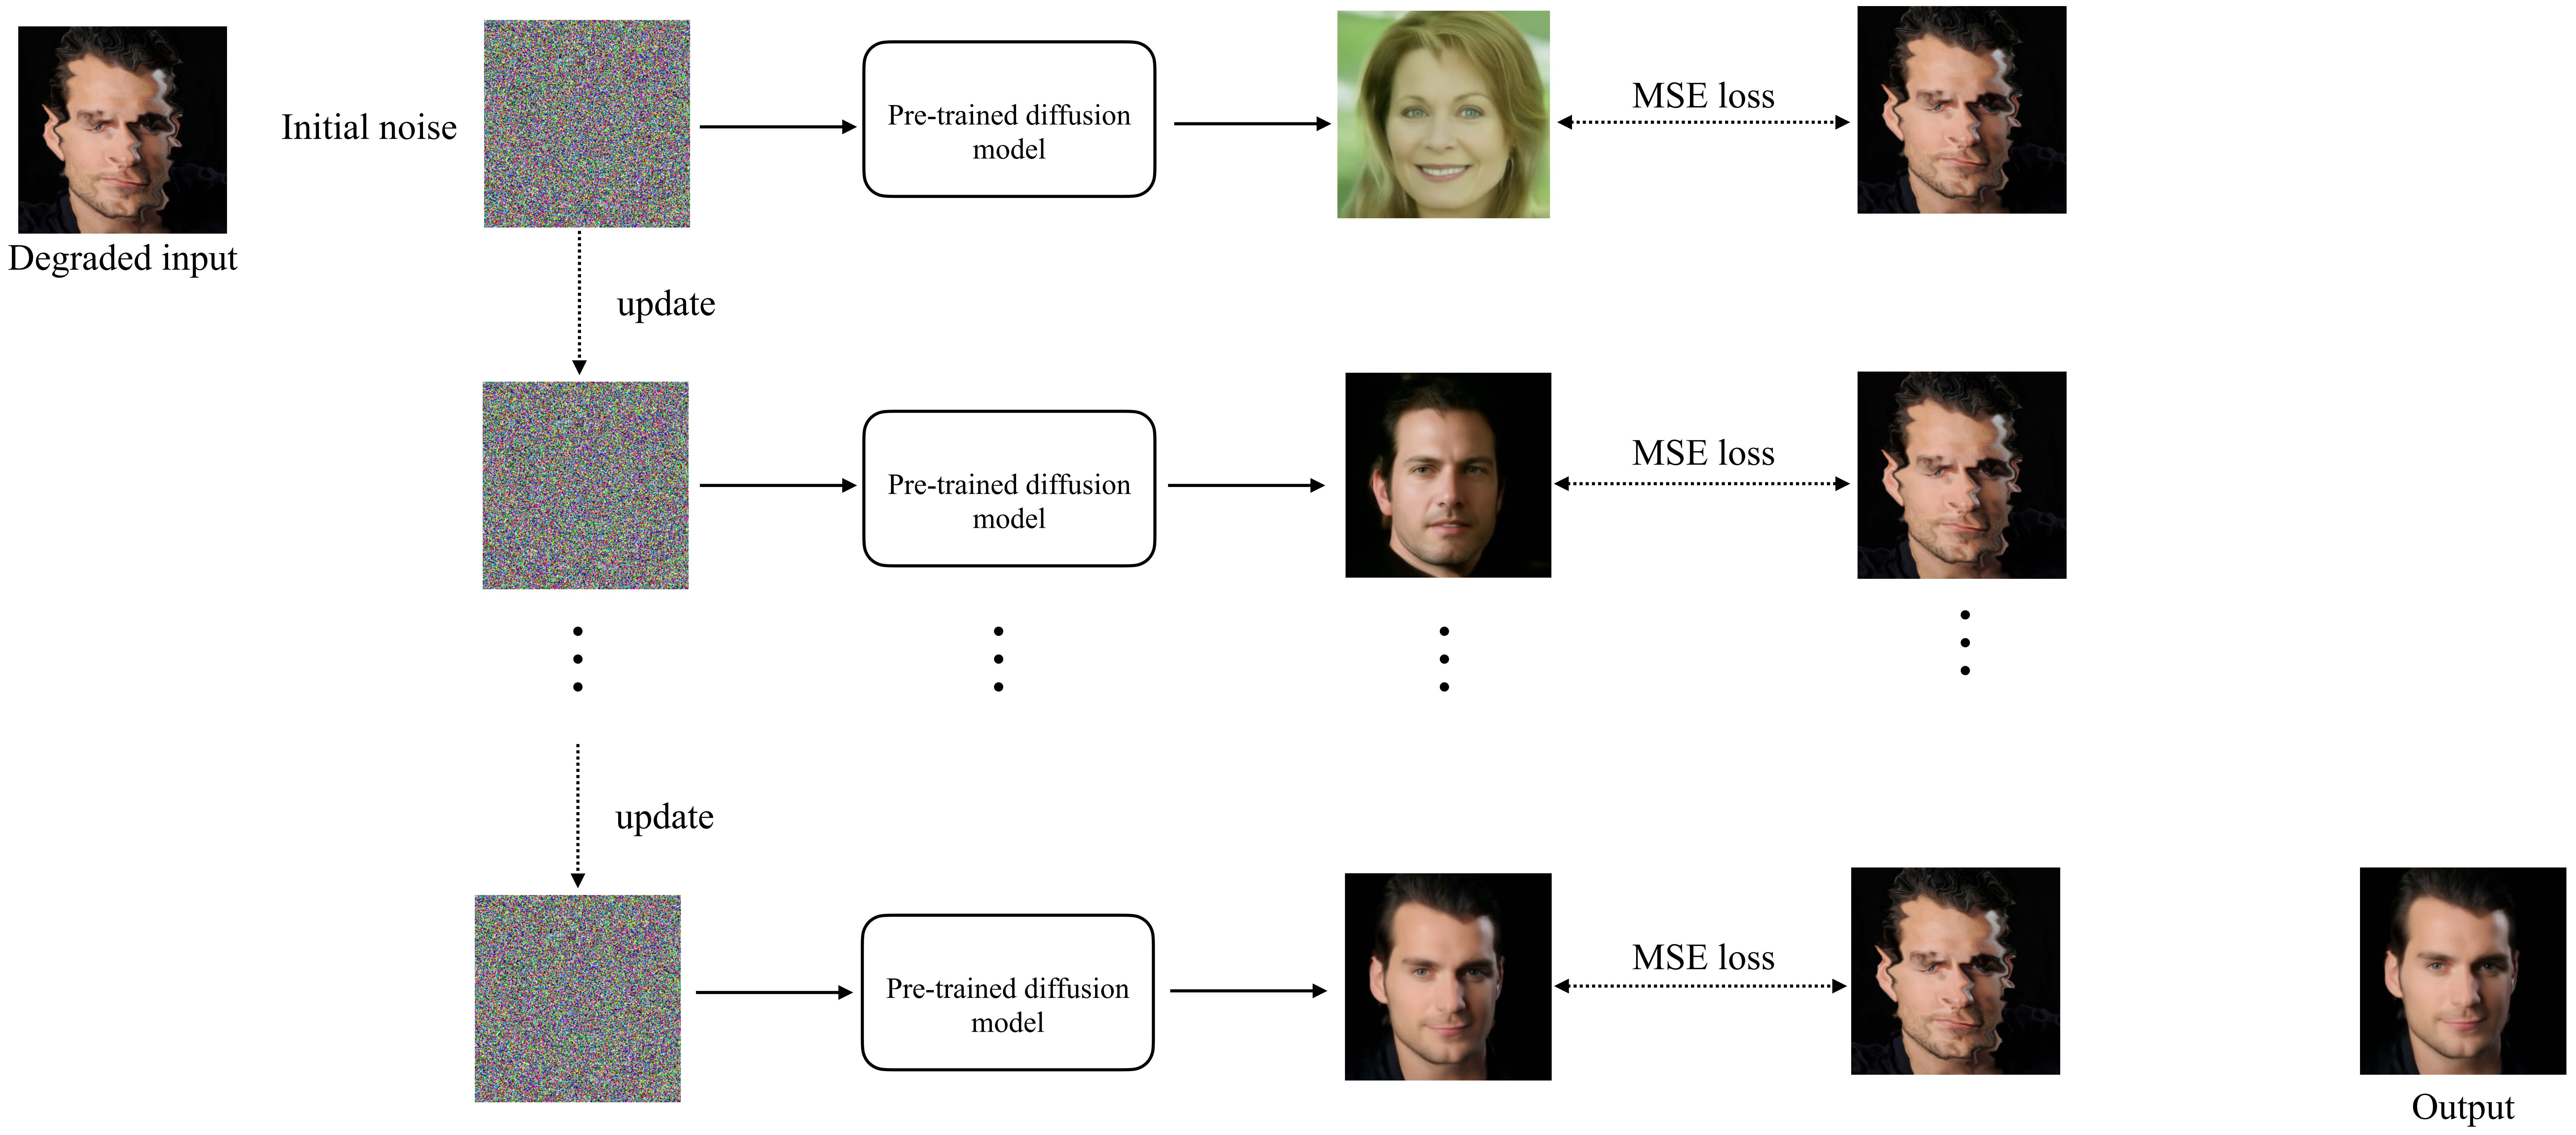

Experiments

- Unstructured degradations
  - water drop
  - non-uniform deformation
  - JPEG artifacts
- Structured degradations
  - Denoising
  - Super resolution

Visual Results - Unstructured Degradations

ICCV HONOLULU HAWAII OCT 19-23, 2025

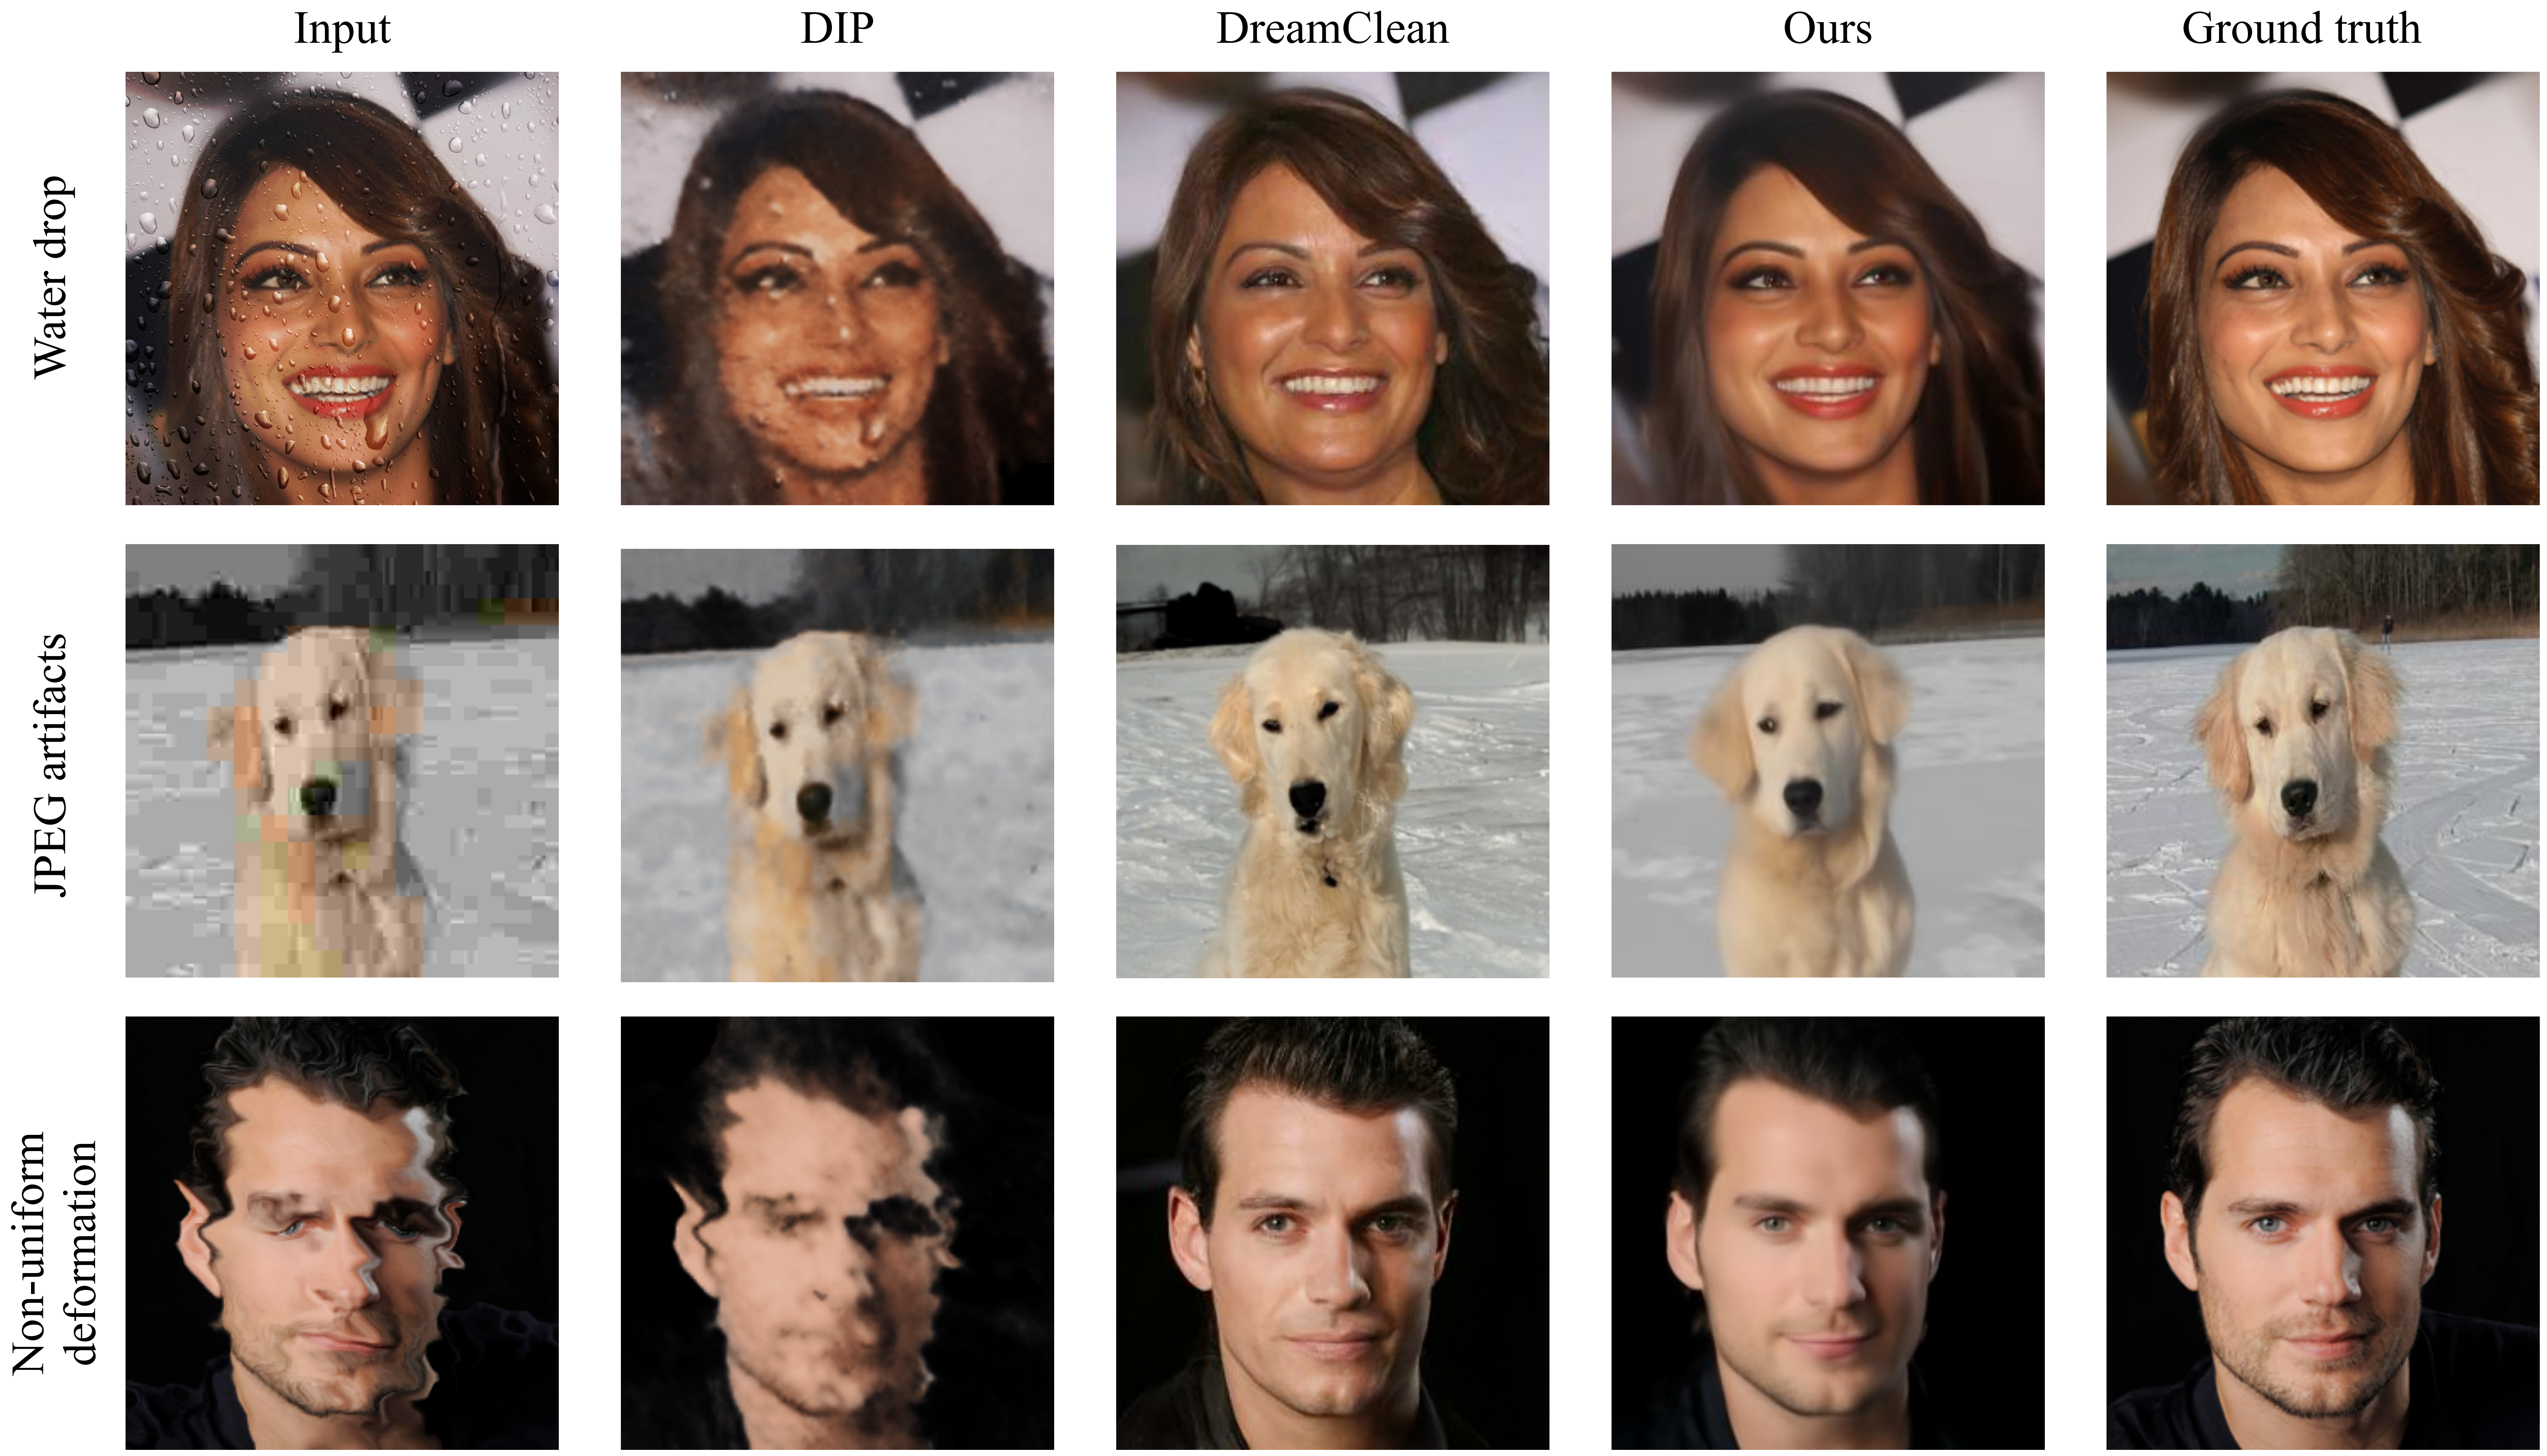

Visual Results - Structured Degradations

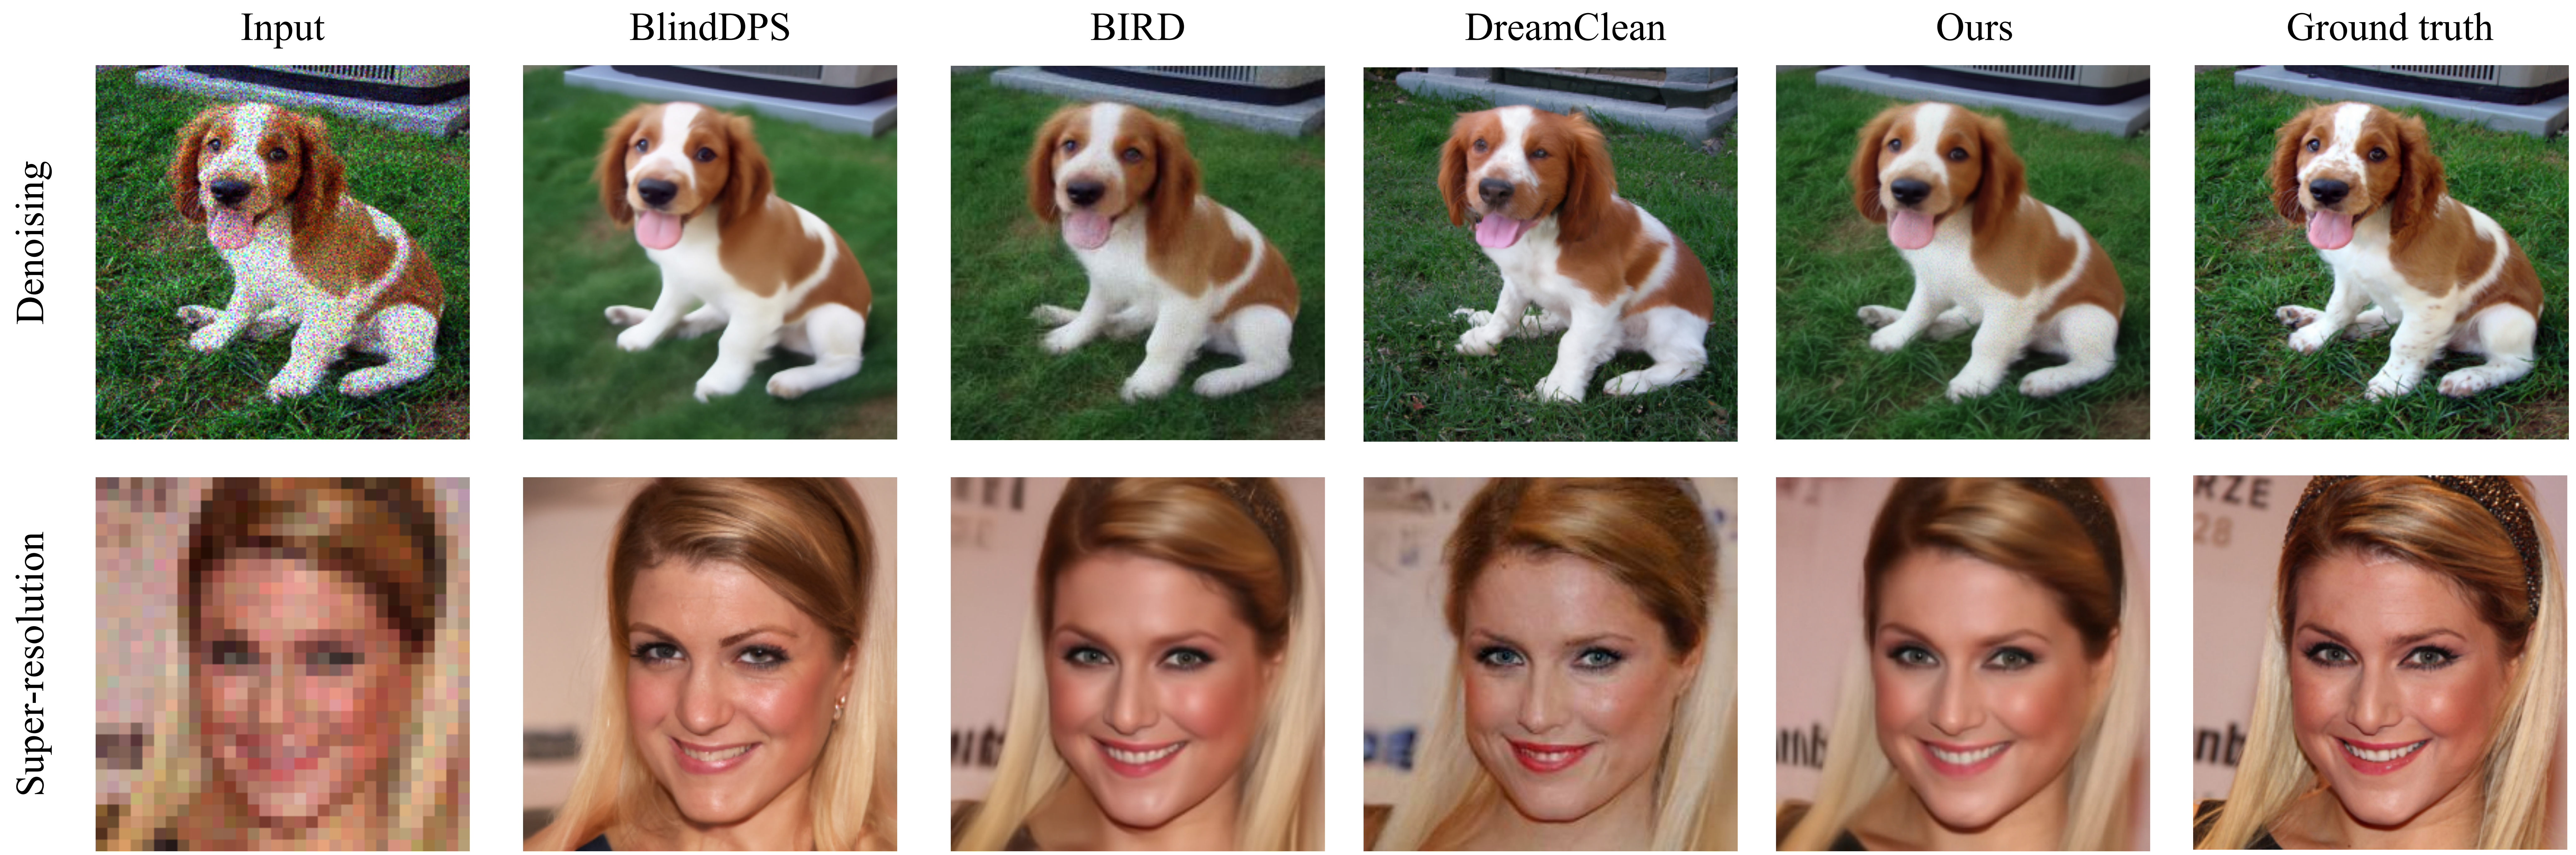

Quantitative Results

- Unstructured degradations

Method	JPEG De-artifacting			Non-uniform Deformation			Water-drop Removal		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
DIP [16]	20.43	0.593	0.622	18.83	0.437	0.643	20.37	0.517	0.642
DreamClean [19]	23.92	0.691	0.342	22.16	0.612	0.398	22.94	0.643	**0.361**
Ours	**25.29**	**0.783**	**0.325**	**23.45**	**0.689**	**0.392**	**23.78**	**0.702**	0.377

- Structured degradations

Method	Denoising			Superresolution (×4)			Superresolution (×8)		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
GDP [8]	27.73	0.817	0.232	24.21	0.708	0.337	21.66	0.618	0.374
Gibbsddrm [13]	27.38	0.809	0.255	24.38	0.689	0.330	21.45	0.605	0.364
BIRD [2]	27.92	0.821	0.238	**25.26**	0.751	**0.294**	22.63	0.626	0.352
BlindDPS [5]	27.56	0.813	0.246	24.51	0.722	0.324	21.73	0.620	0.360
DIP [16]	25.81	0.606	0.345	21.33	0.566	0.426	20.34	0.488	0.471
DreamClean [16]	27.05	0.771	0.236	23.44	0.663	0.322	21.33	0.586	0.344
Ours	**28.37**	**0.842**	**0.224**	25.14	**0.764**	0.301	**22.86**	**0.651**	**0.336**

Limitations

- Our stopping criteria are still non-optimal.

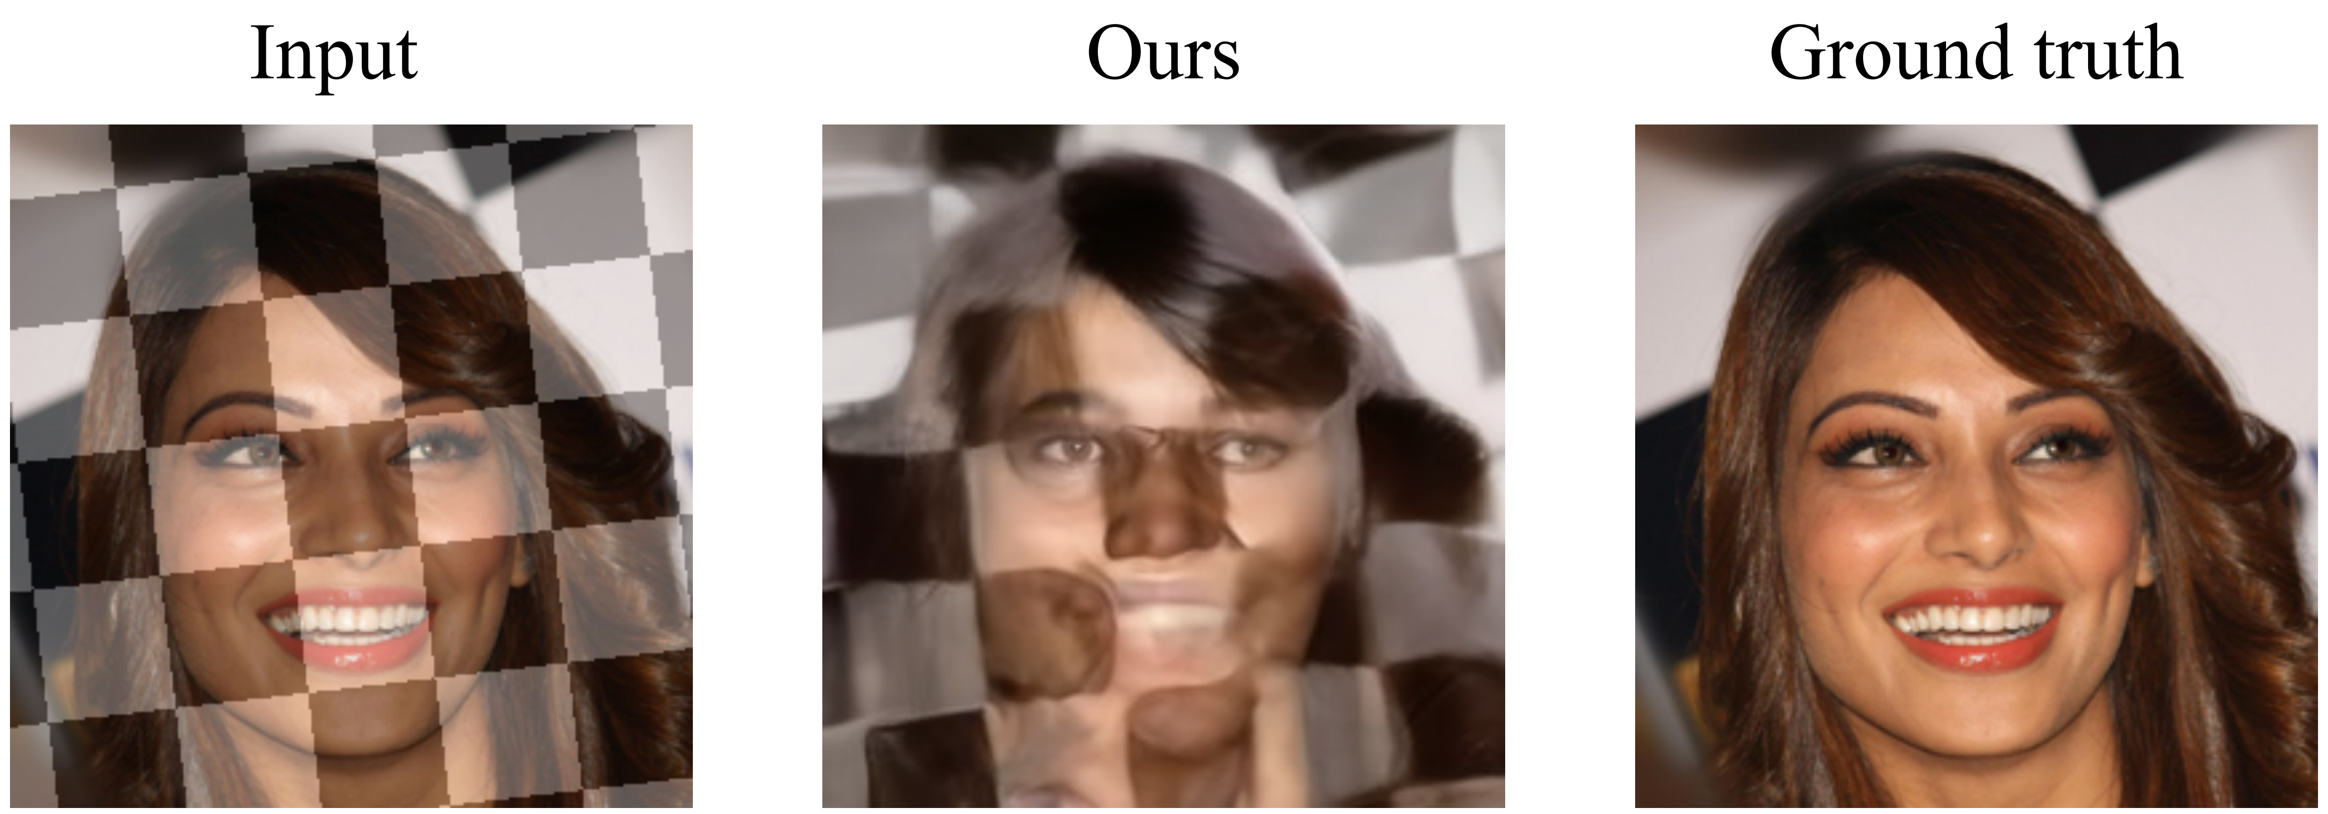

Failure case

- Over-smoothed outputs especially in some case of blurry input images.

- Investigating further self-supervised criteria is a promising research direction.

Summary

- We show the implicit prior of a frozen pre-trained diffusion model when used for degraded image reconstruction.
- We propose a new training-free and fully blind image restoration method, DIIP, which does not assume any prior knowledge of the degradation model.
- We demonstrate state-of-the-art performance on several unstructured degradation tasks.

Diffusion Image Prior

Hamadi Chihaoui, Paolo Favaro

Computer Vision Group, University of Bern, Switzerland