DiracDiffusion: Denoising and Incremental Reconstruction with Assured Data-Consistency

Zalan Fabian, Berk Tinaz, Mahdi Soltanolkotabi

Paper Code

Inverse problems

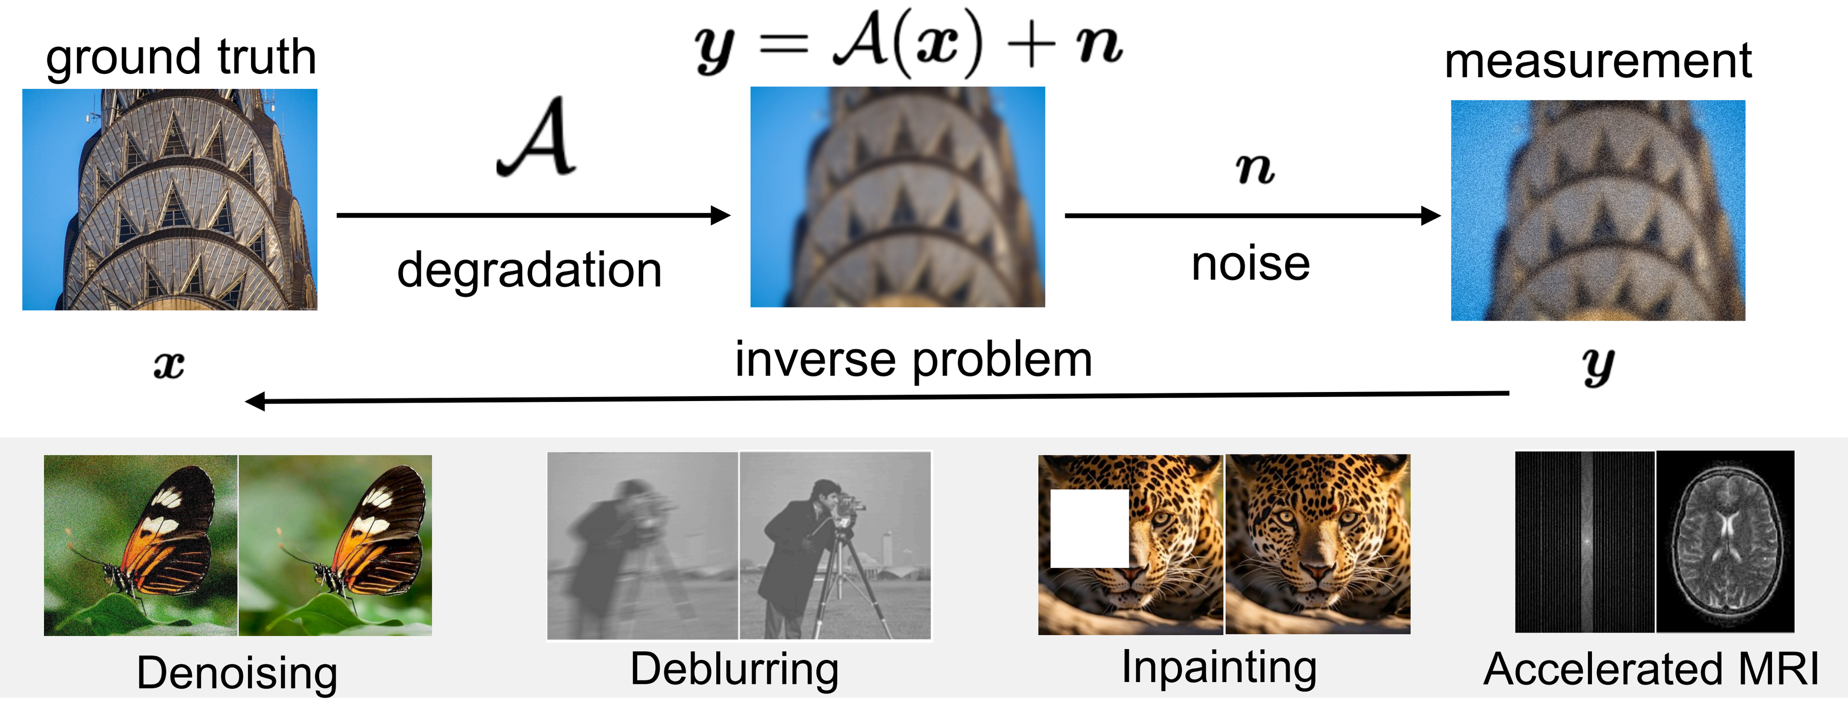

Consistency with the measurement is crucial!

Diffusion solvers

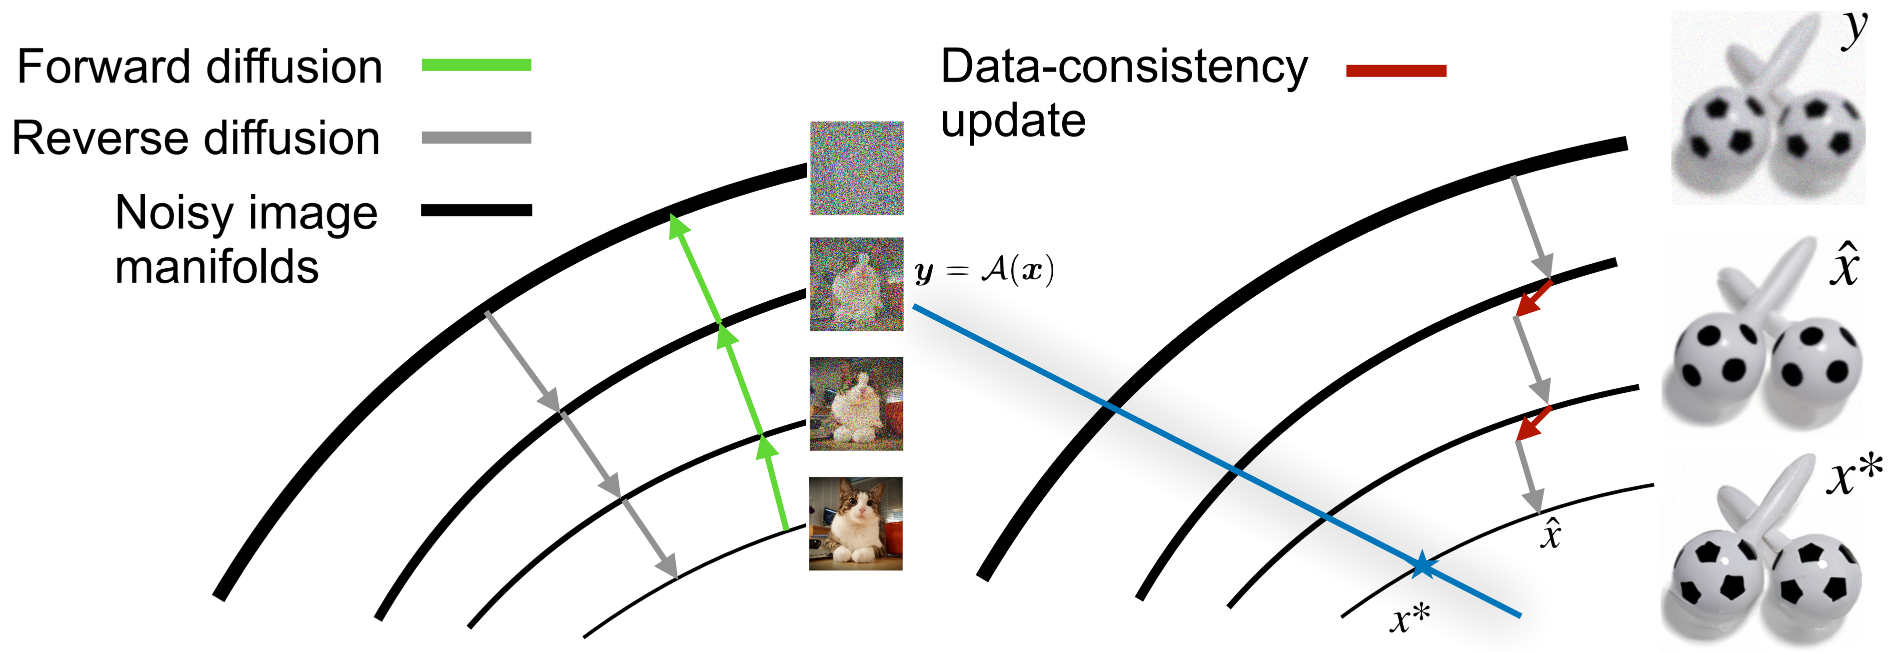

Diffusion models reverse a Gaussian noising process for image generation.

Data-consistency is encouraged by additional guidance.

Key idea: tailor diffusion process to the image degradation, such that **data-consistency is ensured!**

Stochastic Degradation Process (SDP)

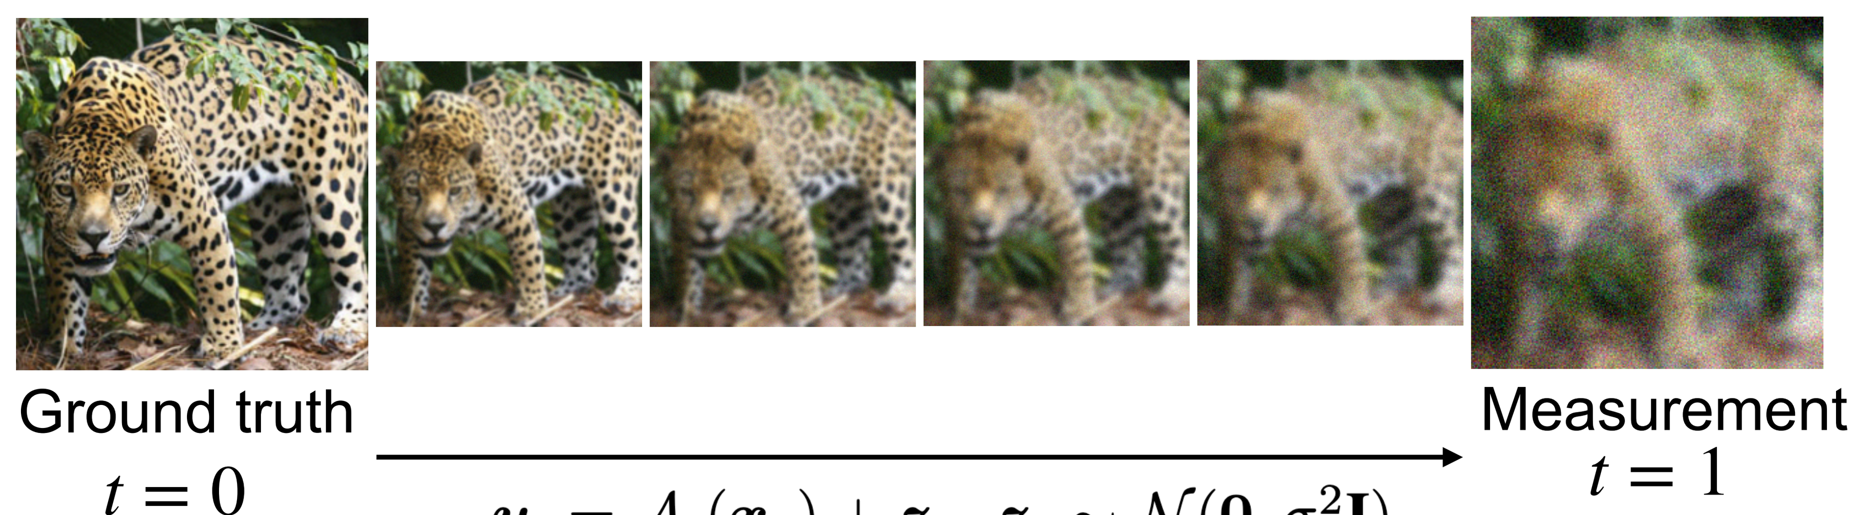

$$y_t = \mathcal{A}_t(x_0) + z_t, \; z_t \sim \mathcal{N}(0, \sigma_t^2 \mathbf{I})$$

Degradation severity

SDP is defined based on a notion of *degradation severity:*

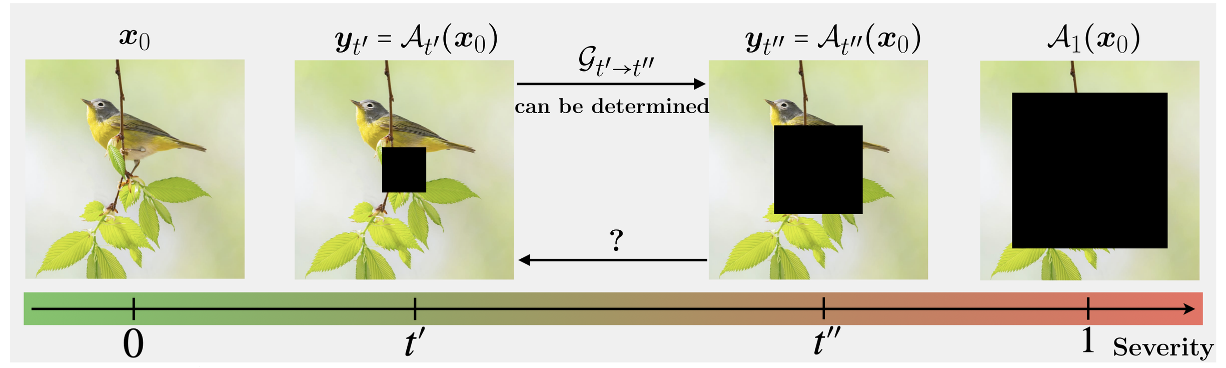

Property 1: $\mathcal{A}_t$ is more severe than $\mathcal{A}_{t'}$, $\forall t > t'$

Property 2: $\mathcal{A}_0(\mathbf{x_0}) = \mathbf{x_0}$

DiracDiffusion

Reversing the SDP

$$\mathrm{d}y_t = \dot{\mathcal{A}}_t(x_0)\mathrm{d}t + \sqrt{\frac{\mathrm{d}}{\mathrm{d}t}\sigma_t^2}\,\mathrm{d}w \xrightarrow[\text{discretize}]{\text{reverse}}$$

[UPDATE]:

$$y_{t-\Delta t} = y_t + \underbrace{\mathcal{A}_{t-\Delta t}(x_0) - \mathcal{A}_t(x_0)}_{\text{incremental reconstruction}} - \underbrace{(\sigma_{t-\Delta t}^2 - \sigma_t^2)\nabla_{y_t}\log q_t(y_t)}_{\text{denoising}} + \sqrt{\sigma_t^2 - \sigma_{t-\Delta t}^2}\, z$$

We learn to iteratively reverse small steps of degradation, which we call **incremental reconstruction** (IR).

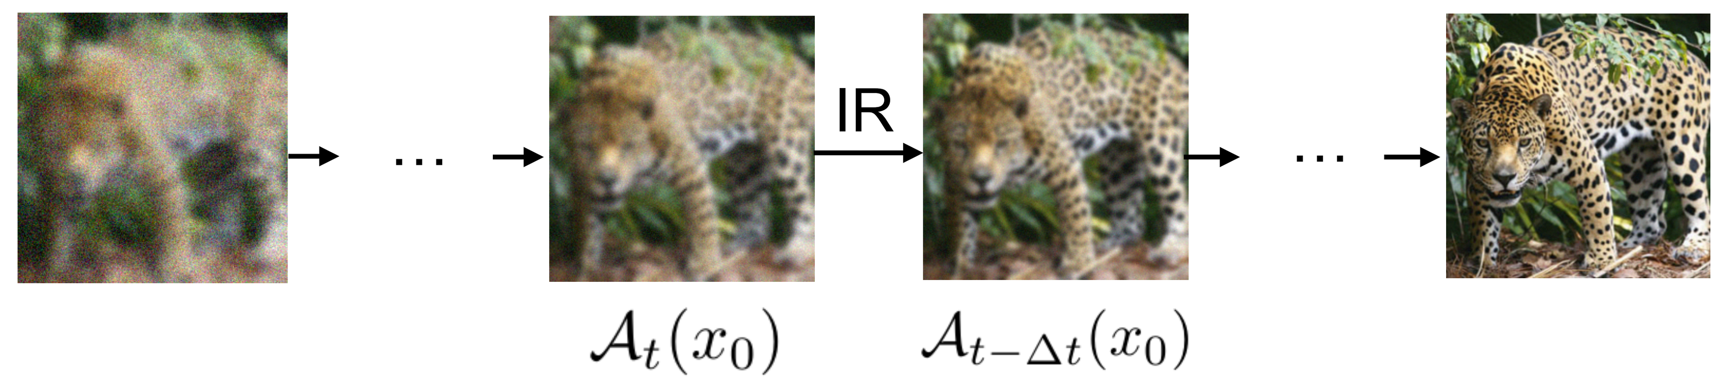

Incremental Reconstruction Loss (IRL)

$$\mathcal{L}_{IR}(\theta) = \mathbb{E}_{t,(x_0,y_t)}\left[\left\|\mathcal{A}_{t-\Delta t}(\Phi_\theta(y_t,t)) - \mathcal{A}_{t-\Delta t}(x_0)\right\|^2\right]$$

Given a degraded image with severity t, we predict a slightly less severe $(t - \Delta t)$ degradation of the clean image.

Theoretical insights

1) Upper-bound on IR error depends on **degradation smoothness** (Theorem 3.4).

2) Minimizing IRL enables learning **both** incremental reconstruction and denoising (Proposition A.6).

3) Running **[UPDATE] ensures data-consistency** in every reverse diffusion step (Theorem 3.6).

Perception-distortion trade-off

Perceptual image quality (LPIPS, FID) is fundamentally at odds with distortion (PSNR, SSIM). We control the trade-off via **early-stopping** the reverse process.

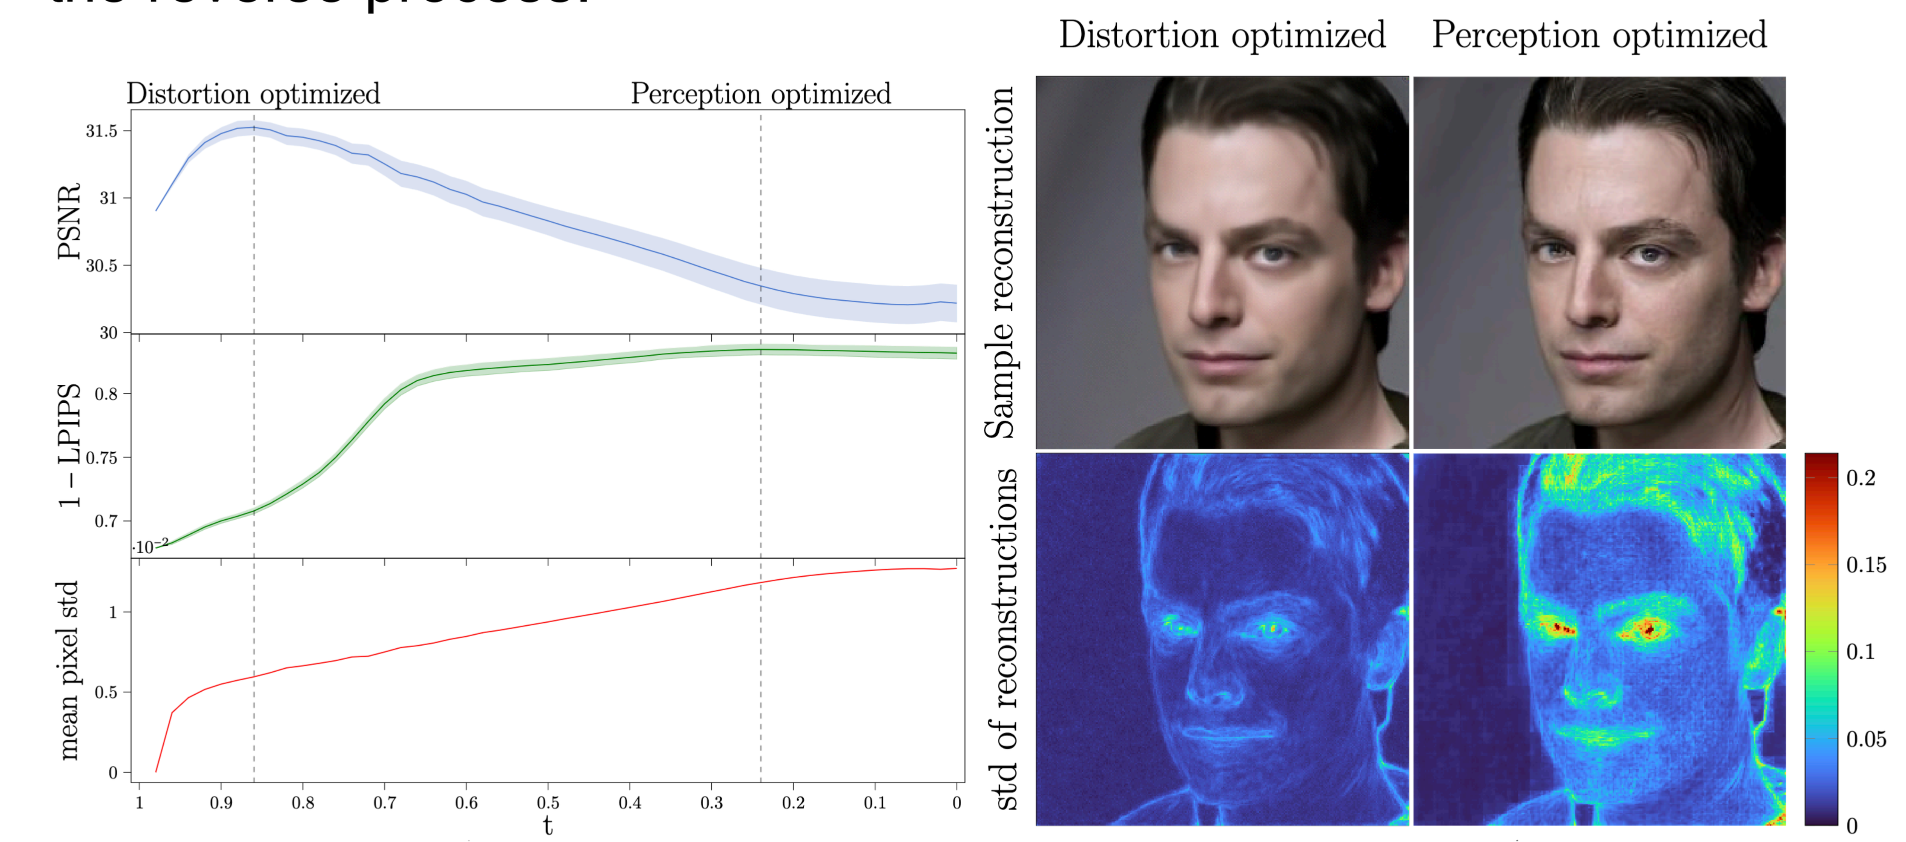

Results

Data-consistency

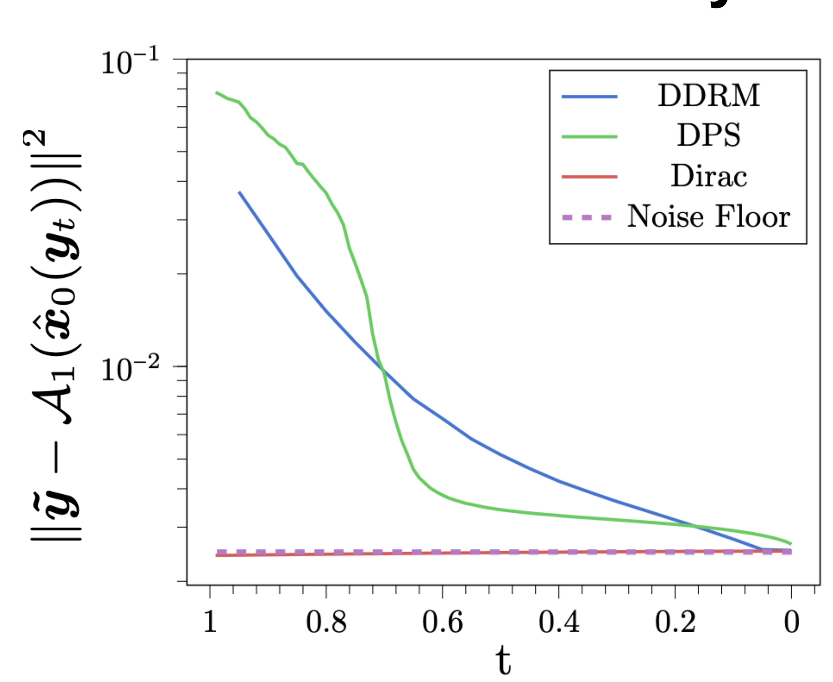

Excellent reconstruction quality

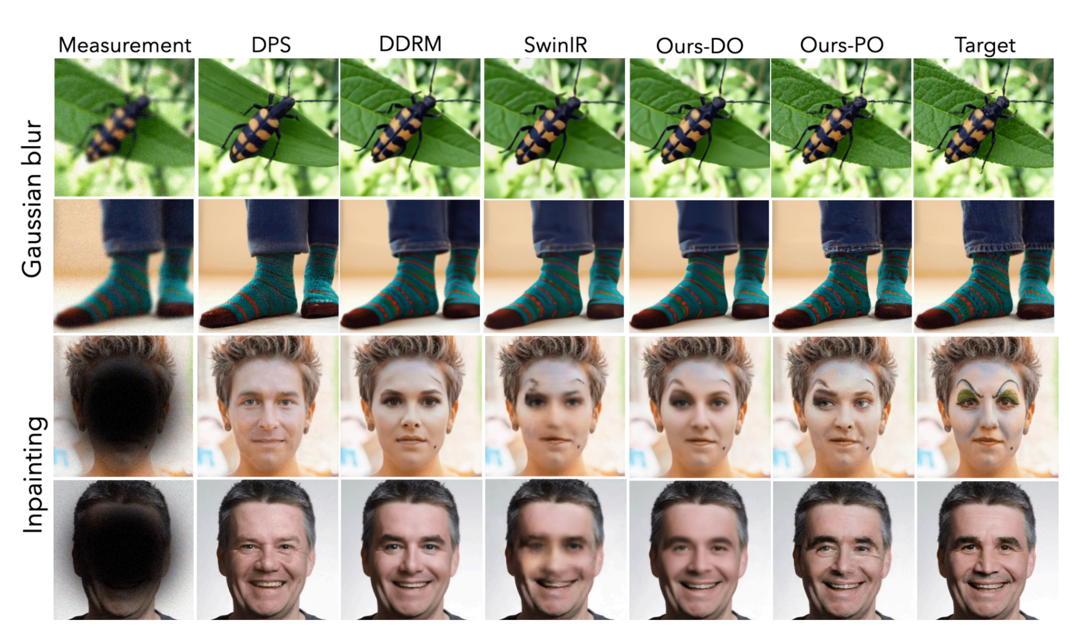

Fast sampling

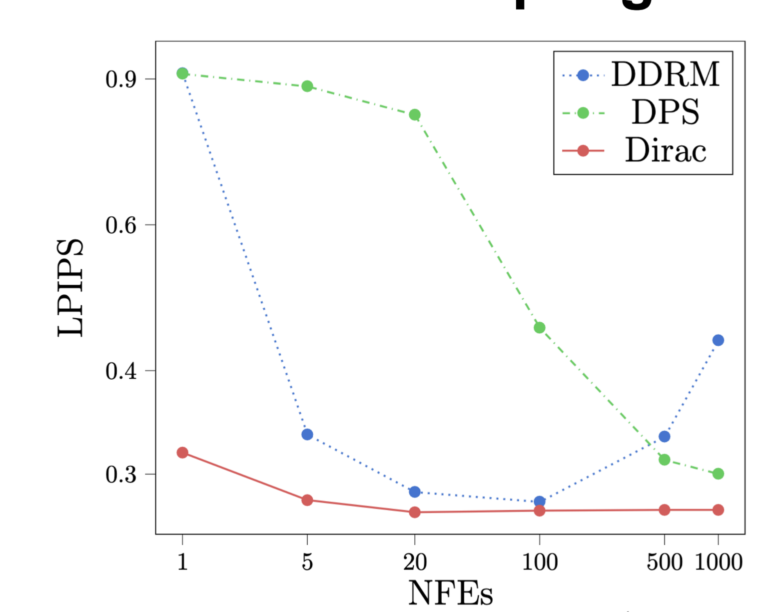

Method	Deblurring PSNR(↑)	Deblurring SSIM(↑)	Deblurring LPIPS(↓)	Deblurring FID(↓)	Inpainting PSNR(↑)	Inpainting SSIM(↑)	Inpainting LPIPS(↓)	Inpainting FID(↓)
Dirac-PO (ours)	26.67	0.7418	**0.2716**	**53.36**	25.41	0.7595	0.2611	**39.43**
Dirac-DO (ours)	28.47	0.8054	0.2972	69.15	**26.98**	**0.8435**	**0.2234**	51.87
DPS (Chung et al., 2022a)	25.56	0.6878	0.3008	65.68	21.06	0.7238	0.2899	57.92
DDRM (Kawar et al., 2022a)	27.21	0.7671	0.2849	65.84	25.62	0.8132	0.2313	54.37
SwinIR (Liang et al., 2021)	**28.53**	**0.8070**	0.3048	72.93	24.46	0.8134	0.2660	59.94
PnP-ADMM (Chan et al., 2016)	27.02	0.7596	0.3973	74.17	12.27	0.6205	0.4471	192.36
ADMM-TV	26.03	0.7323	0.4126	89.93	11.73	0.5618	0.5042	264.62

Method	Deblurring PSNR(↑)	Deblurring SSIM(↑)	Deblurring LPIPS(↓)	Deblurring FID(↓)	Inpainting PSNR(↑)	Inpainting SSIM(↑)	Inpainting LPIPS(↓)	Inpainting FID(↓)
Dirac-PO (ours)	24.68	0.6582	**0.3302**	53.91	26.36	0.8087	0.2079	34.33
Dirac-DO (ours)	**25.76**	**0.7085**	0.3705	83.23	**28.92**	**0.8958**	**0.1676**	38.25
DPS (Chung et al., 2022a)	21.51	0.5163	0.4235	**52.60**	22.71	0.8026	0.1986	34.55
DDRM (Kawar et al., 2022a)	24.53	0.6676	0.3917	61.06	25.92	0.8347	0.2138	**33.71**
SwinIR (Liang et al., 2021)	25.07	0.6801	0.4159	84.80	26.87	0.8490	0.2161	45.69
PnP-ADMM (Chan et al., 2016)	25.02	0.6722	0.4565	98.72	18.14	0.7901	0.2709	101.25
ADMM-TV	24.31	0.6441	0.4578	88.26	17.60	0.7229	0.3157	120.22

Table 1. Experimental results on the FFHQ (top) and ImageNet (bottom) test splits.