Diffusion Prior-Based Amortized Variational Inference for Noisy Inverse Problems

Sojin Lee*, Dogyun Park*, Inho Kong, Hyunwoo J. Kim†

Department of Computer Science and Engineering, Korea University

Sojin Lee

Dogyun Park

Inho Kong

Hyunwoo J. Kim

Problem definition

✓ **Forward model**

$\mathbf{y} = \mathbf{H}\mathbf{x}_0$

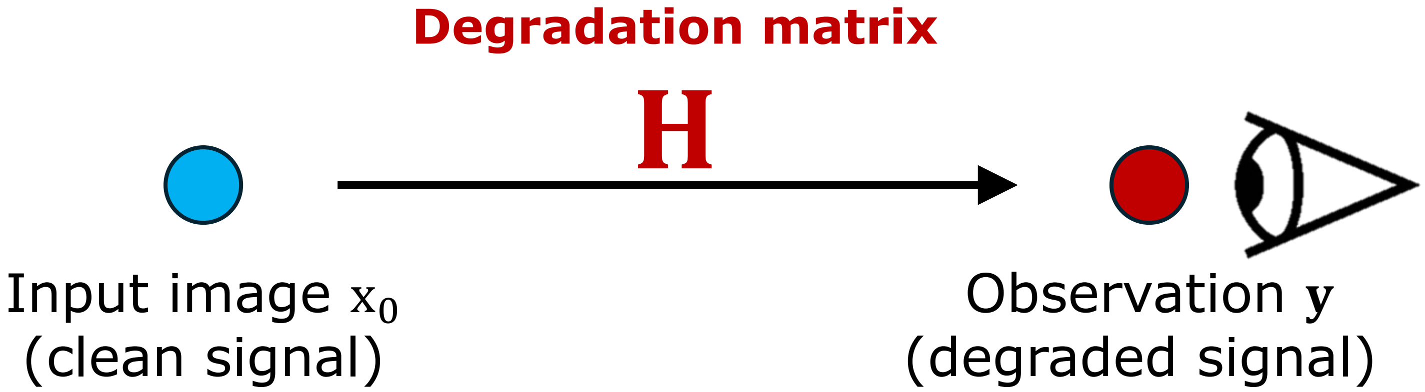

✓ **Inverse problem**

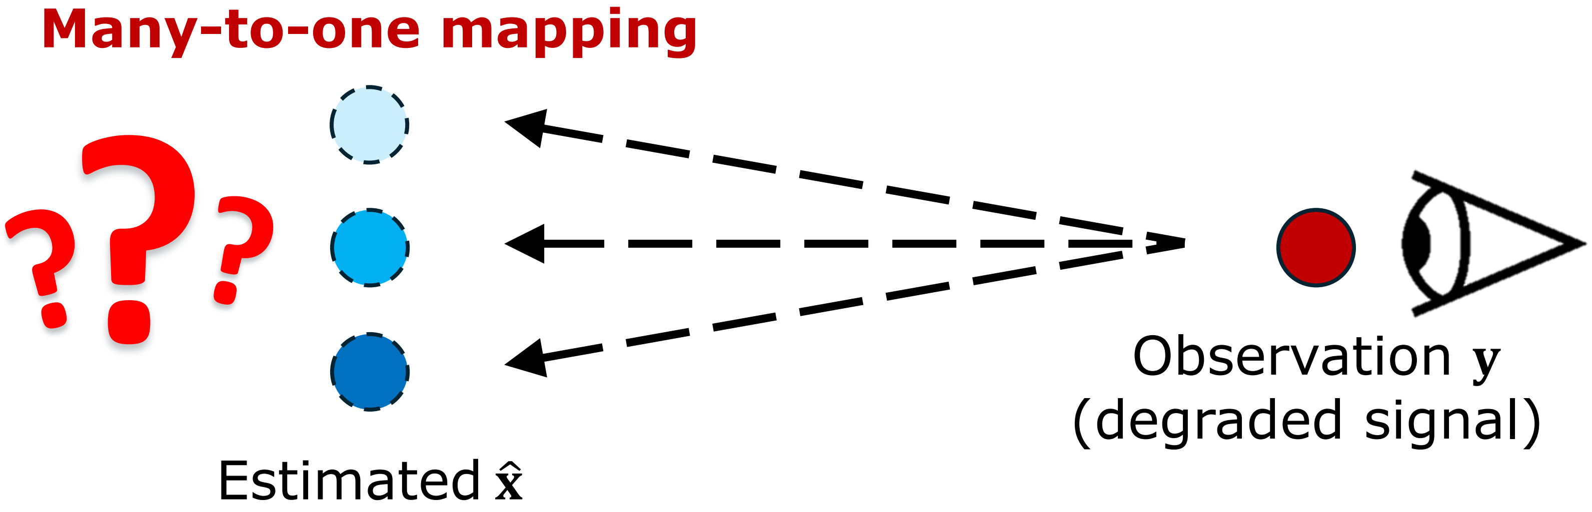

Problem definition

- ✓ **Forward model**

degradation matrix $\mathbf{H} \in \mathbb{R}^{d_\mathbf{y} \times d_{\mathbf{x}_0}}$

i.i.d white Gaussian noise $\boldsymbol{n} \in \mathbb{R}^{d_\mathbf{y}}$

$$\mathbf{y} = \mathbf{H}\mathbf{x}_0 + \mathbf{n}, \qquad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_\mathbf{y}^2 \mathbf{I})$$

Measurement $\mathrm{y} \in \mathbb{R}^{d_\mathbf{y}}$

clean image $\mathbf{x}_0 \in \mathbb{R}^{d_{\mathbf{x}_0}}$

$\hat{\mathbf{x}}_0$

$\mathbf{x}_0$

$\mathbf{H}$ →

$\mathbf{H}\mathbf{x}_0$

+

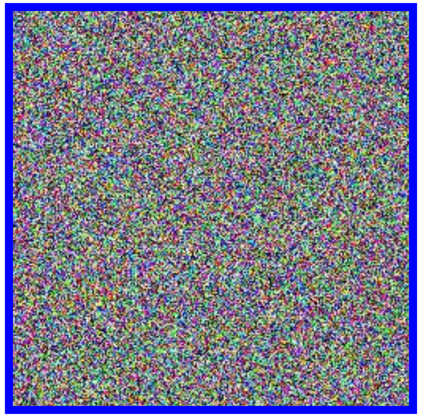

$\mathbf{n}$

=

$\mathbf{y}$

Inference speed

✓ **Previous methods**

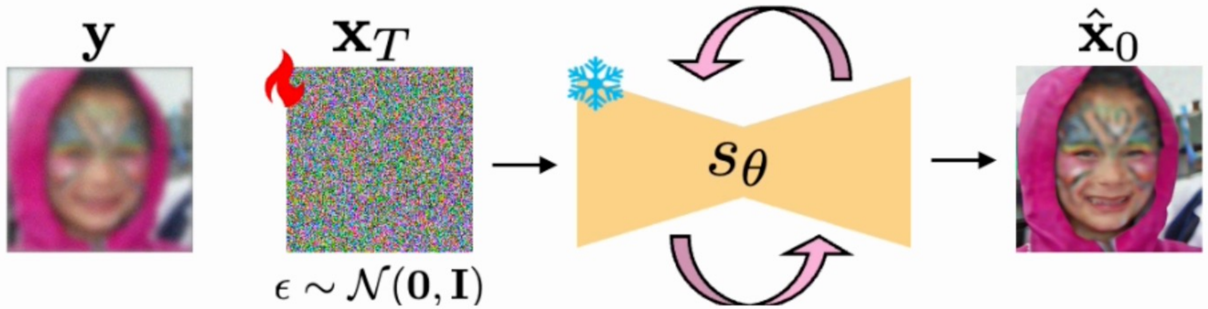

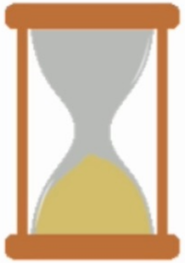

Generated 1 image

✓ **DAVI**

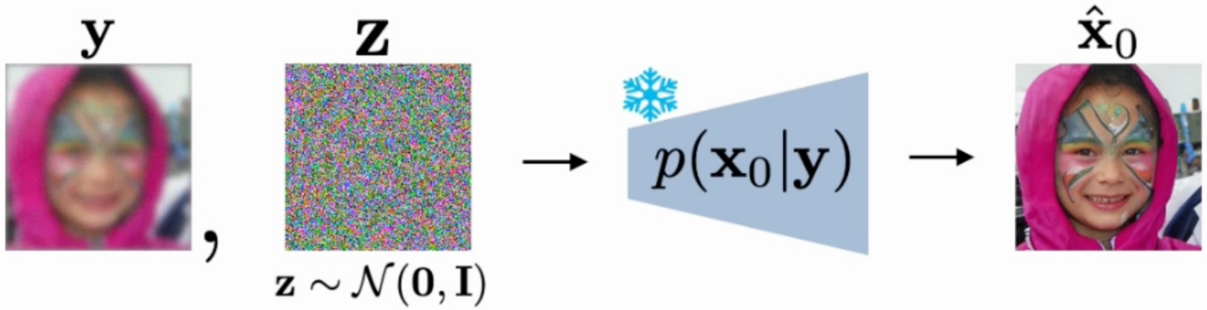

Generated 1,000 images

Inference speed

✓ **Previous methods**

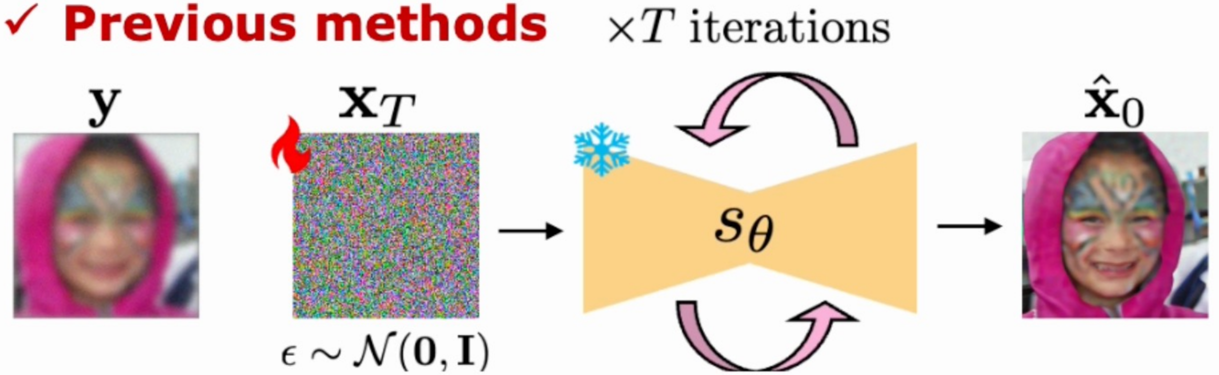

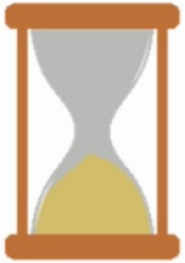

Generated 1 image

Generated 1,000 images

✓ **DAVI**

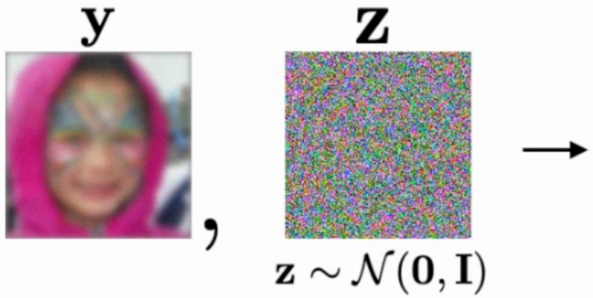

RED-diff 1 image ↔ DAVI 1,000 images

DAVI shows 1000x faster inference speed

Inference speed

DAVI requires 20~1,000 times fewer steps than baselines

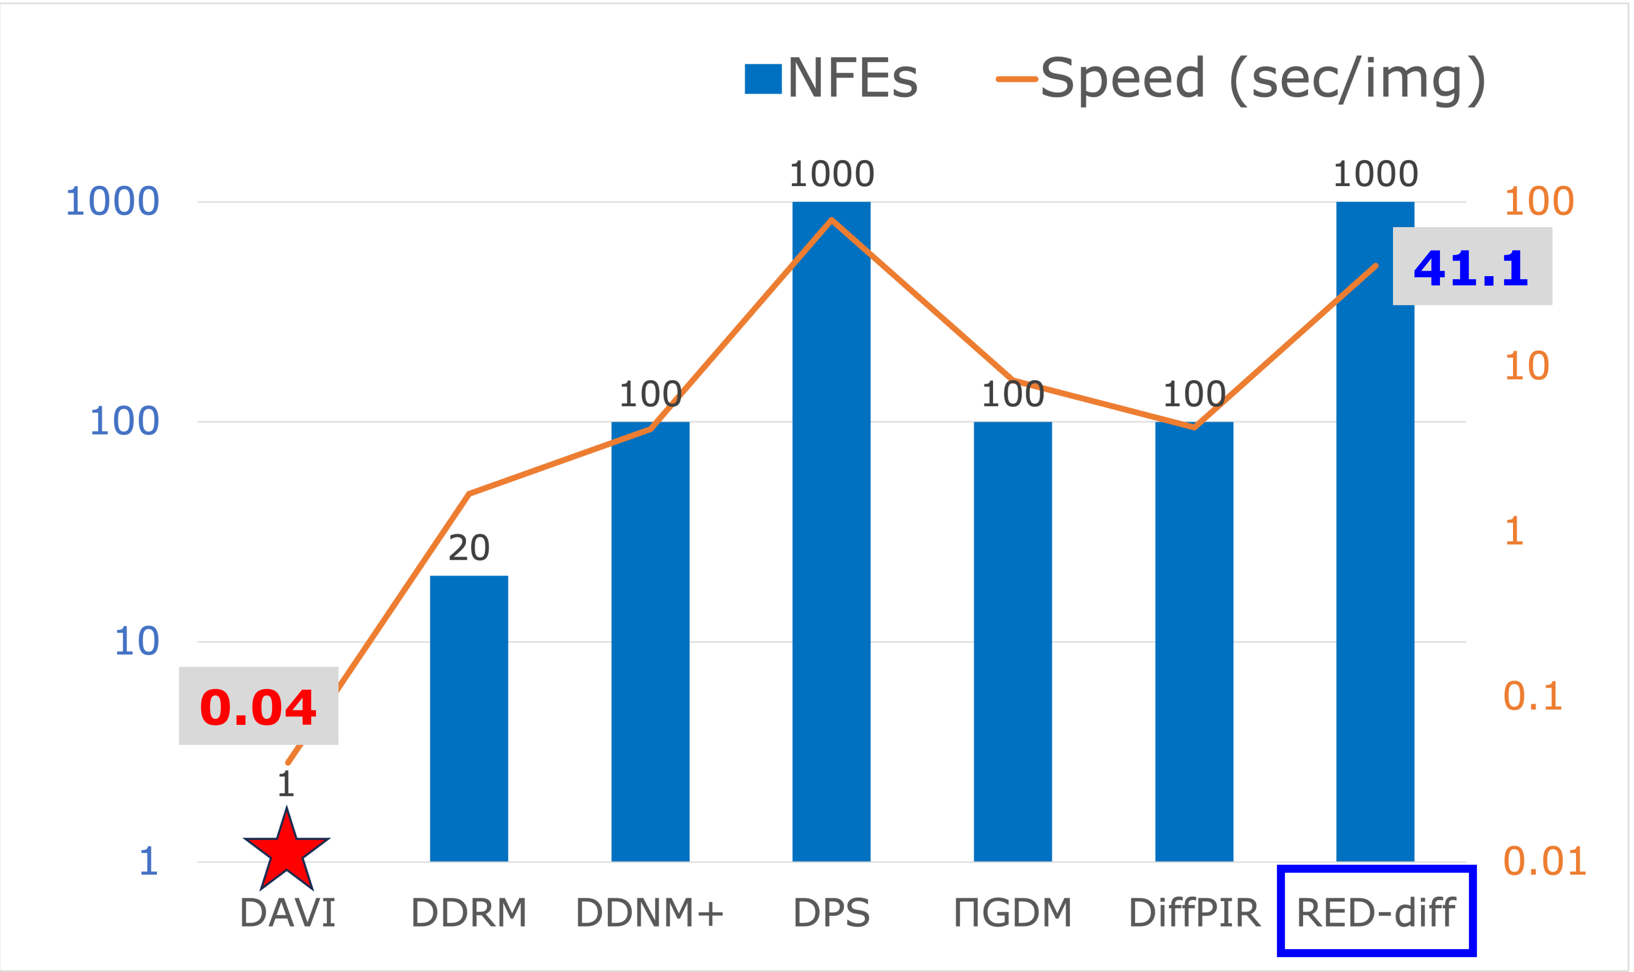

Qualitative results

Comparison with baselines

✓ Previous methods

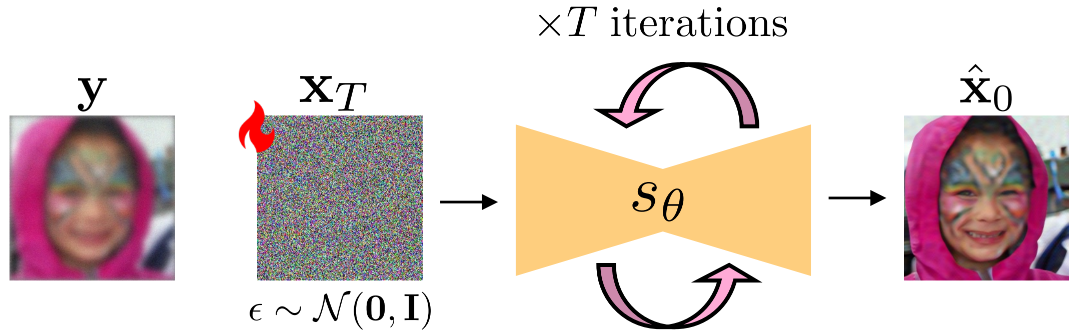

✓ DAVI

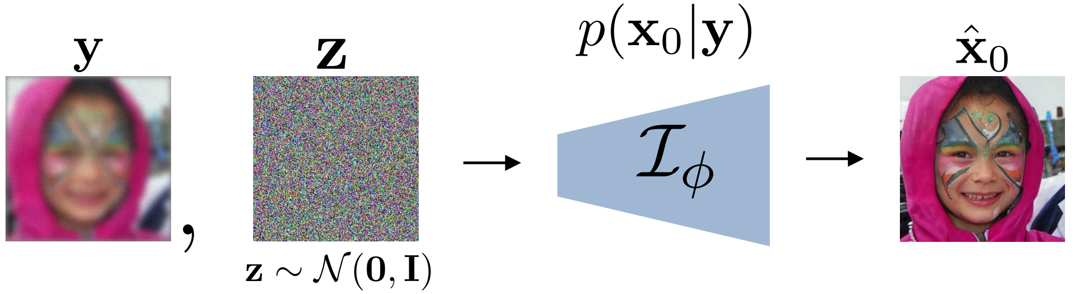

Inference

Previous methods:

- ✓ **Iterative sampling** along reverse diffusion process
- ✓ Pre-trained **diffusion model** s_θ

DAVI:

- ✓ **Single-step sampling**
- ✓ Parameterized **Implicit distribution** $\mathcal{I}_\phi$

Optimization

Previous methods:

- ✓ **Measurement-wise optimization** $\mathbf{x}_T$
- ✓ **Test-time optimization** for each sample

DAVI:

- ✓ **Amortized Optimization** $\mathcal{I}_\phi$
- ✓ **No additional optimization** at test-time

Proposed method

✓ Previous methods

✓ DAVI

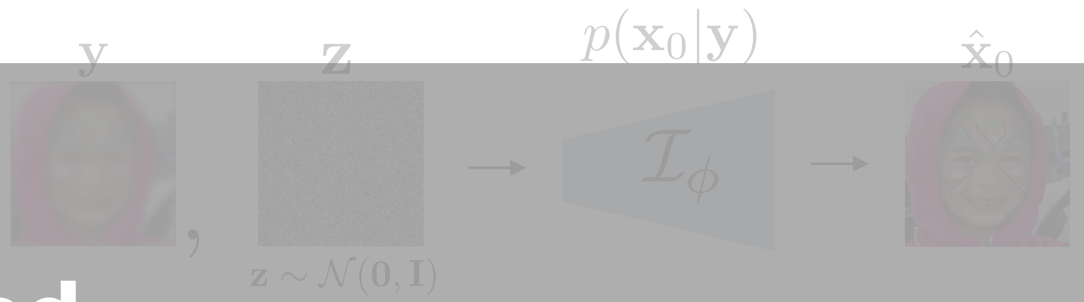

Diffusion Prior-Based Amortized Variational Inference (DAVI)

Inference

✓ **Iterative sampling** along reverse diffusion process

✓ Pre-trained diffusion model s_θ

✓ **Single-step sampling**

✓ Parameterized Implicit distribution $\mathcal{I}_\phi$

Optimization

✓ Measurement-wise optimization $\mathbf{x}_T$

✓ Test-time optimization for each sample

✓ Amortized Optimization $\mathcal{I}_\phi$

✓ No additional optimization at test-time

Method

- **[Goal] Variational Optimization between $q_\phi(\mathbf{x}_0|\mathbf{y})$ and $p(\mathbf{x}_0|\mathbf{y})$**

Implicit distribution

$$\phi^\star = \arg\min_\phi \left[D_{KL}(q_\phi(\mathbf{x}_0|\mathbf{y}) \,\|\, p(\mathbf{x}_0|\mathbf{y})) \right]$$

True posterior distribution

Method

- **Expand the KL divergence and drop the constant term** $\log p(\mathbf{y})$

Implicit distribution   True posterior

$$D_{KL}(q_\phi(\mathbf{x}_0|\mathbf{y}) \,||\, p(\mathbf{x}_0|\mathbf{y}))$$

Data consistency term   **KL divergence**

$$= -\mathbb{E}_{q_\phi(\mathbf{x}_0|\mathbf{y})}\left[\log p(\mathbf{y}|\mathbf{x}_0)\right] + D_{KL}(q_\phi(\mathbf{x}_0|\mathbf{y}) \,||\, p(\mathbf{x}_0))$$

Eq. Forward model

known   estimated   unknown

$$\mathbf{y} = \mathbf{H}\mathbf{x}_0 + \mathbf{n}, \qquad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_\mathbf{y}^2 \mathbf{I})$$

clean image prior

$\approx$ **diffusion prior**

Method

- **Optimize $q_\phi(\mathbf{x}_0|\mathbf{y})$ to align with the clean image distribution $p(\mathbf{x}_0)$**

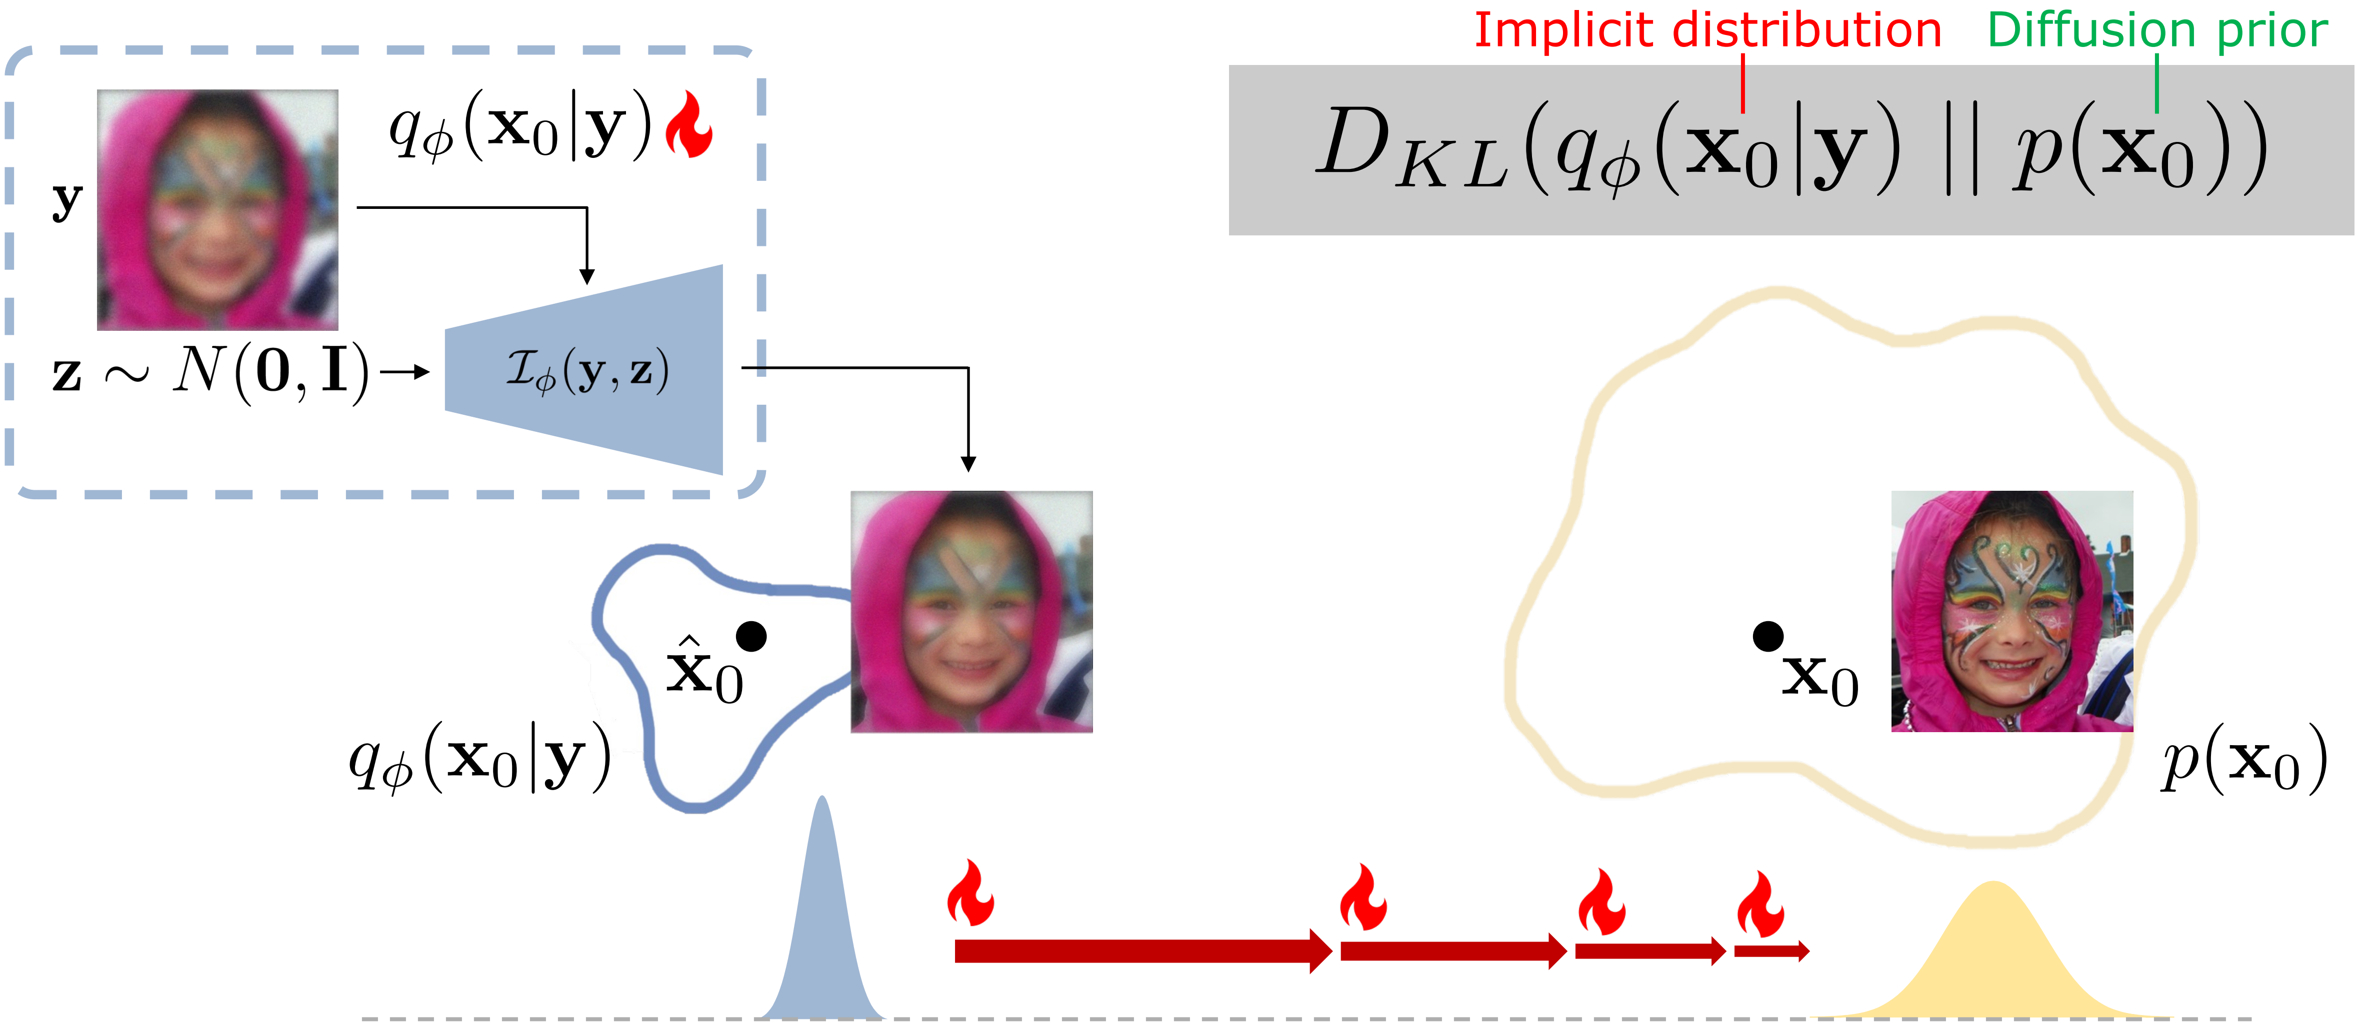

Method

- **Optimize $q_\phi(\mathbf{x}_0|\mathbf{y})$ to align with the clean image distribution $p(\mathbf{x}_0)$**

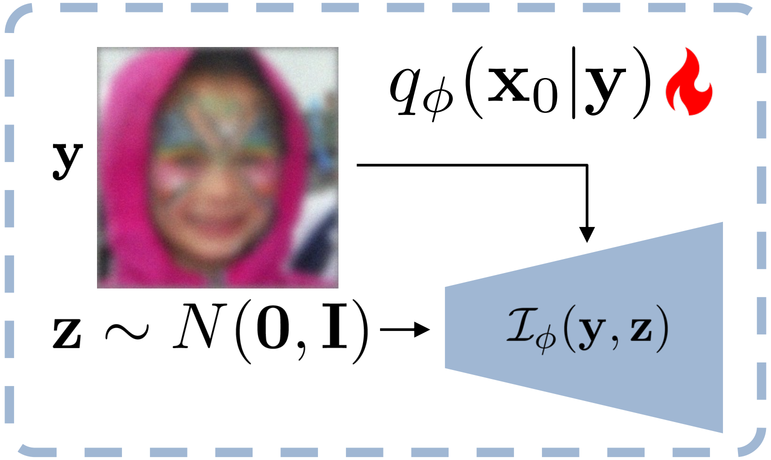

Implicit distribution Diffusion prior

$$D_{KL}(q_\phi(\mathbf{x}_0|\mathbf{y}) \,||\, p(\mathbf{x}_0))$$

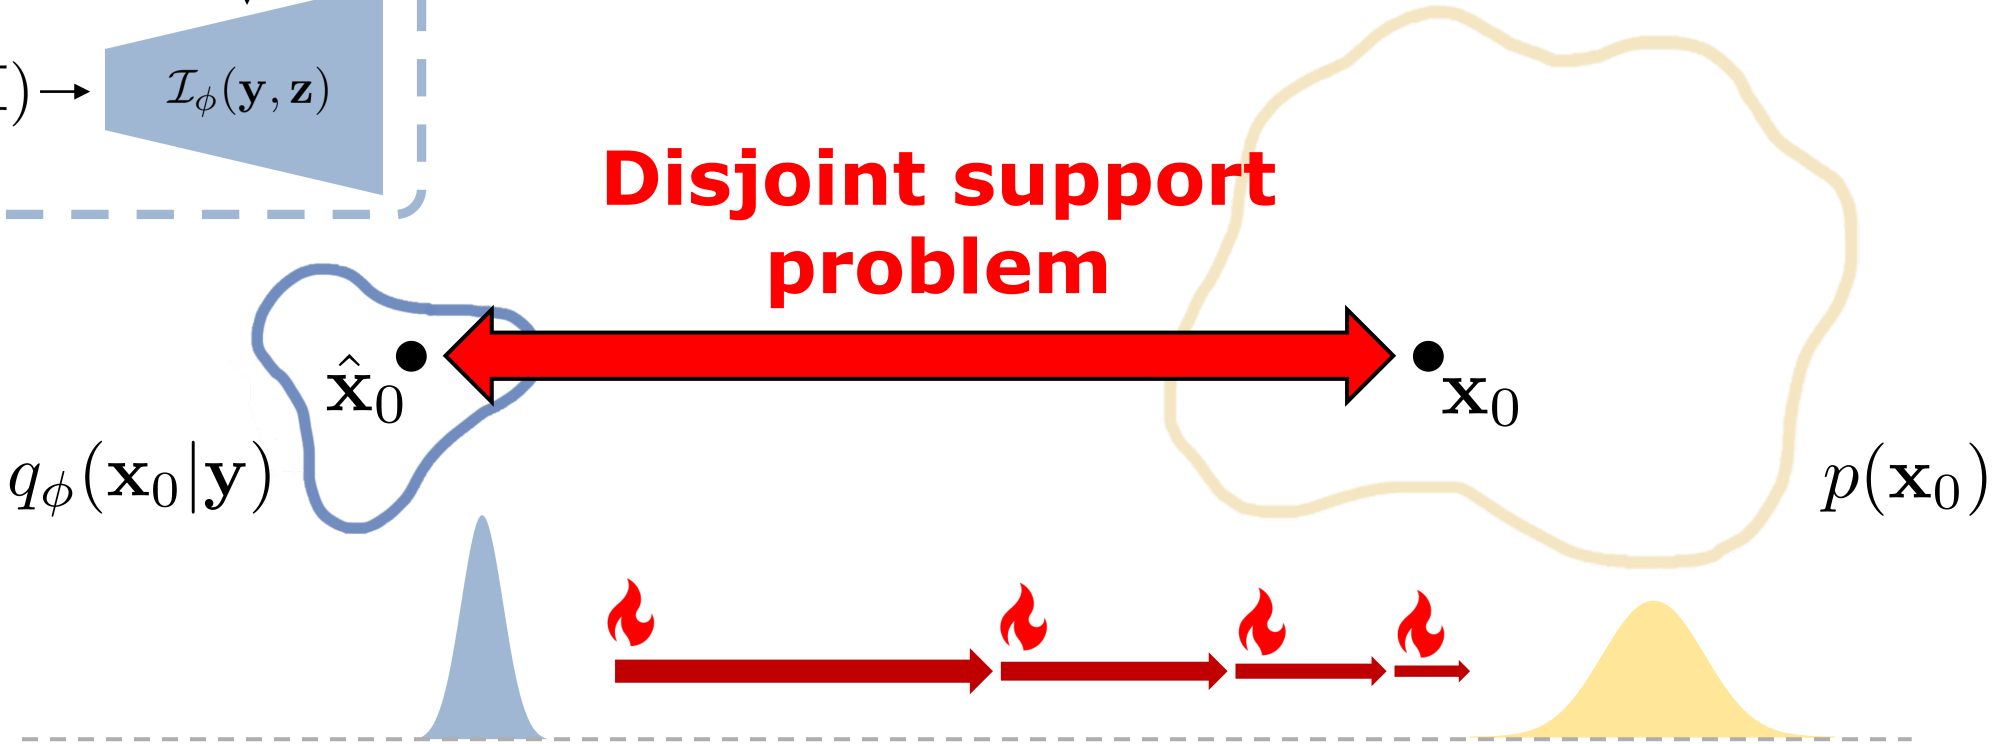

Integrated KL divergence (IKL)

- **Alleviate the disjoint support problem by smoothing the distributions**

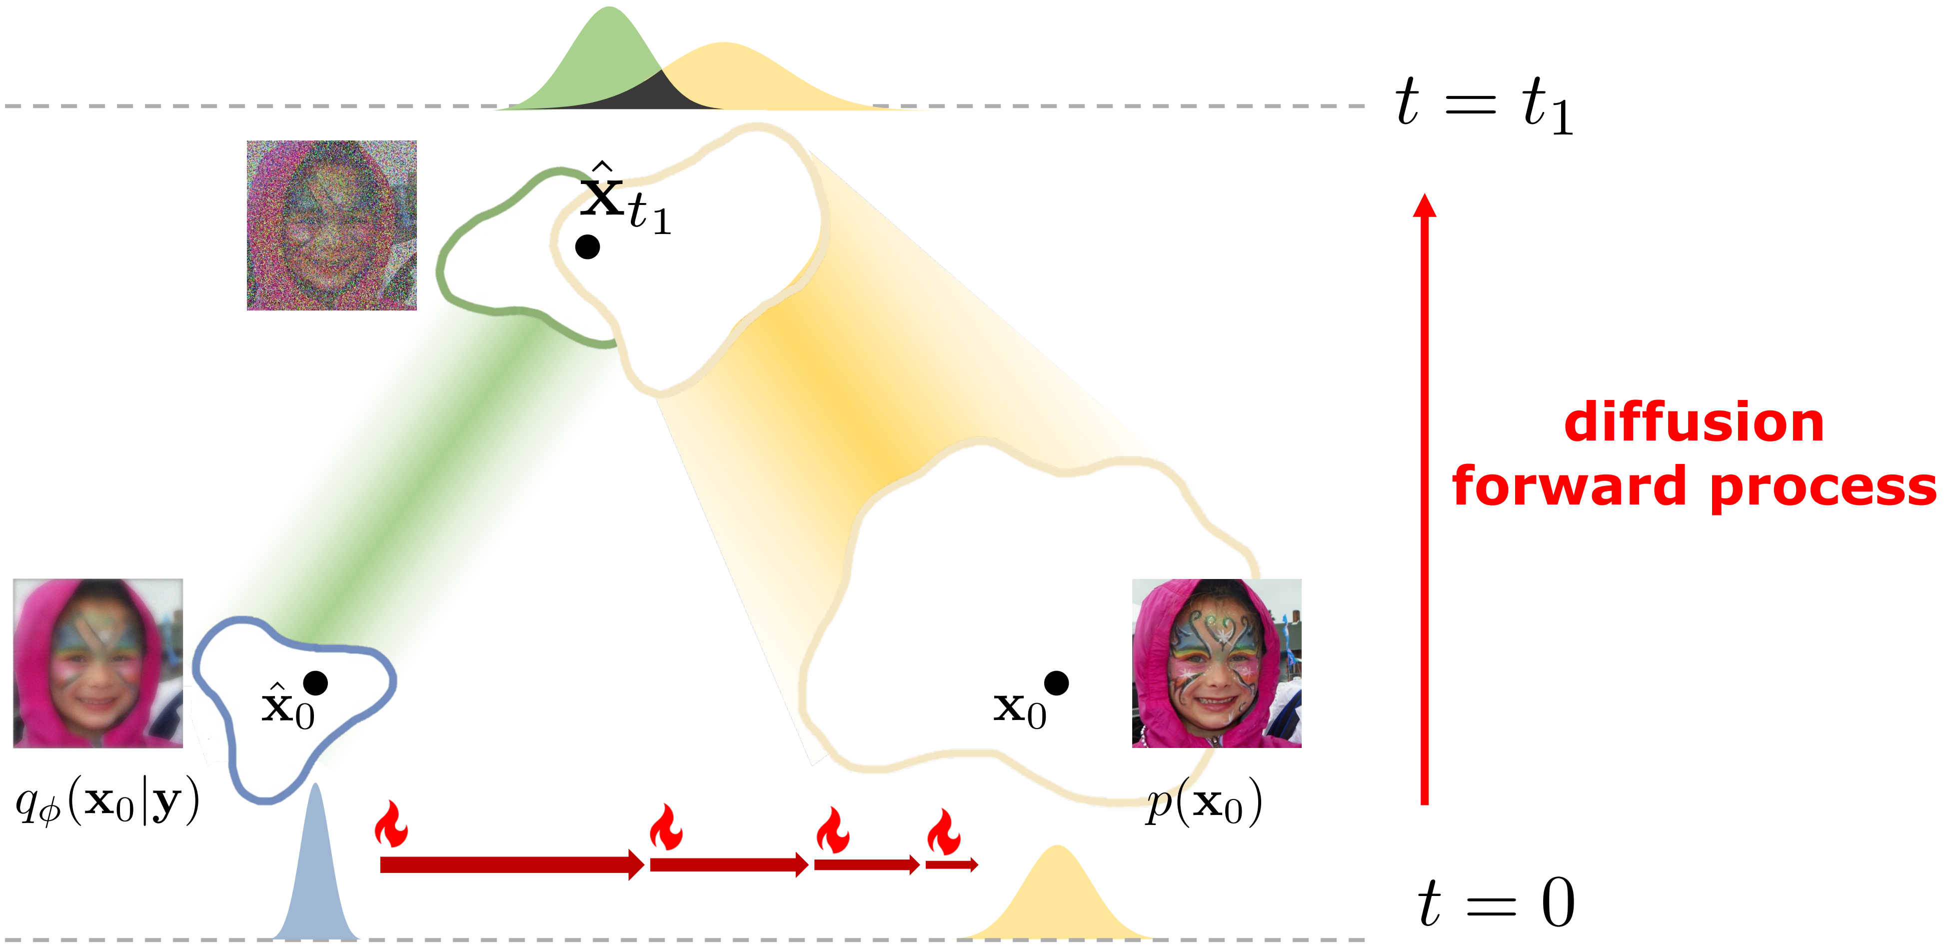

Score distillation gradient

- **Minimize the discrepancy between $q_\phi(\mathbf{x}_0|\mathbf{y})$ and $p(\mathbf{x}_0)$**

[Requirements]
Implicit score function s_ψ

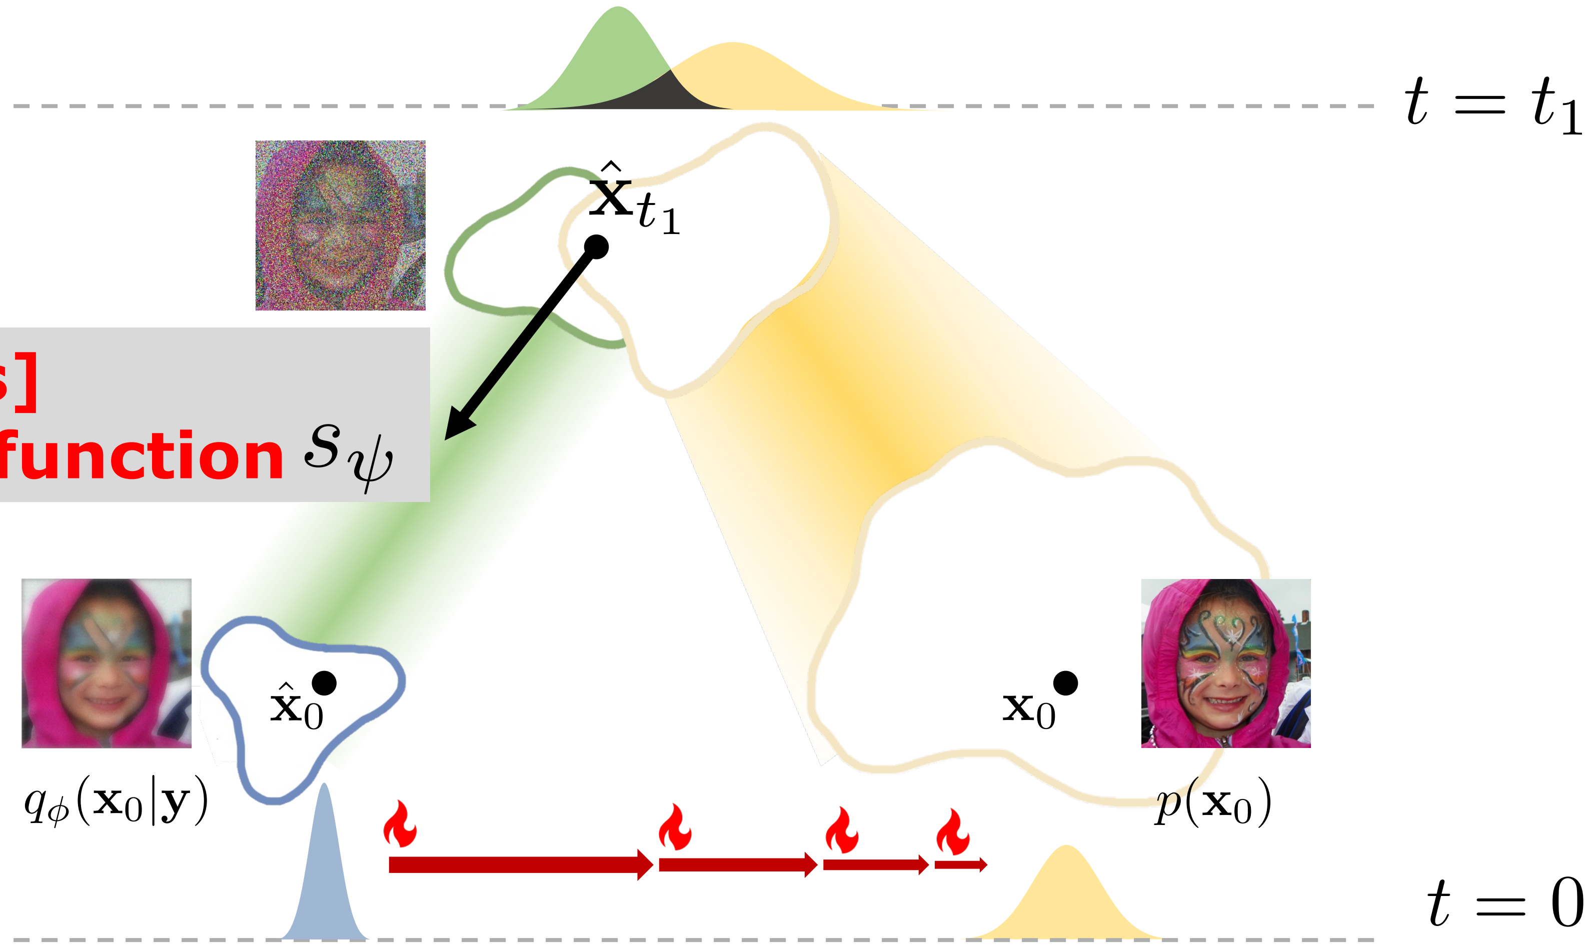

Score distillation gradient

- **Minimize the discrepancy between $q_\phi(\mathbf{x}_0|\mathbf{y})$ and $p(\mathbf{x}_0)$**

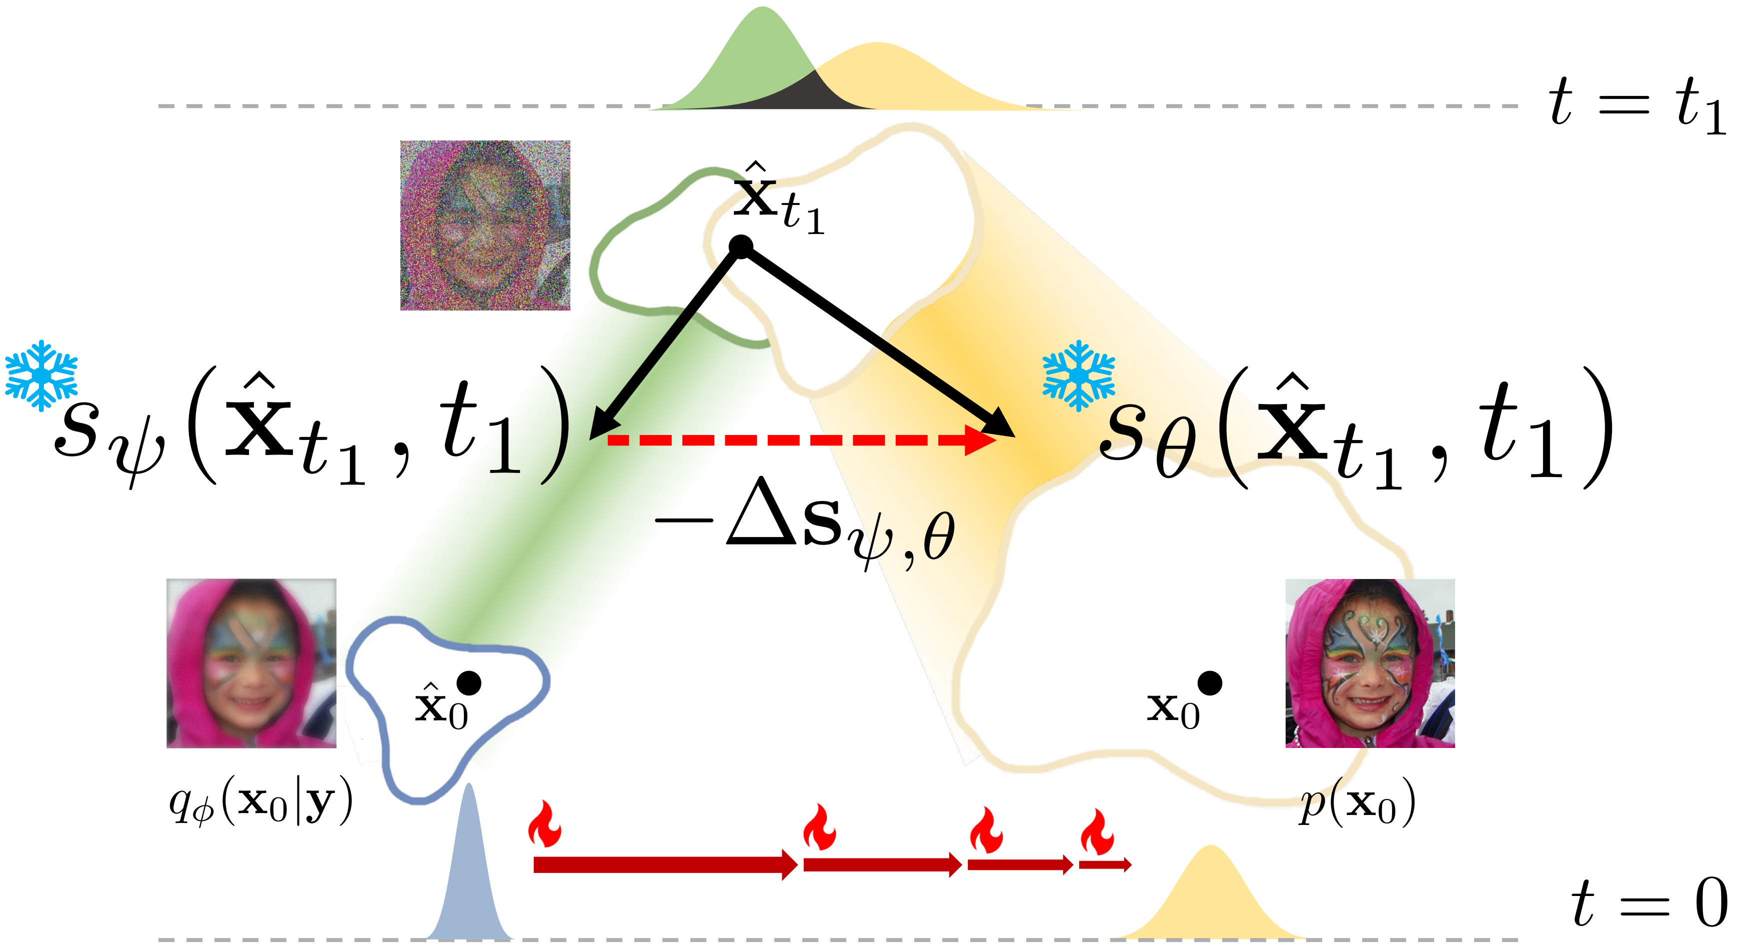

Denoising score matching loss

- **Optimize the implicit score function $\mathbf{s}_\psi$**

$$\mathbb{E}_{q_\phi(\mathbf{x}_t|\mathbf{y}),t}\left[\|\mathbf{s}_\psi(\mathbf{x}_t,t) - \nabla_{\mathbf{x}_t}\log q_\phi(\mathbf{x}_t|\mathbf{y})\|_2^2\right]$$

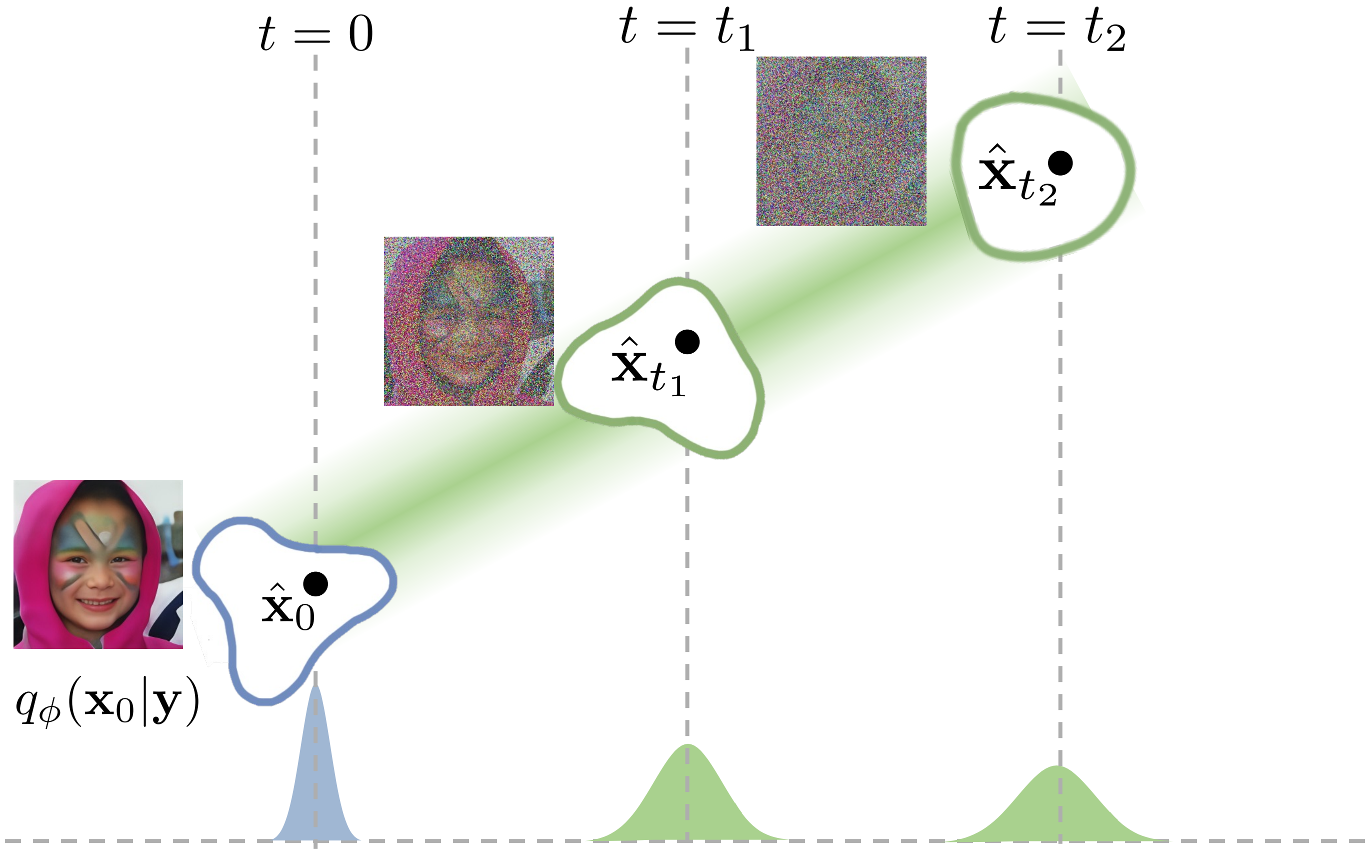

Perturbed Posterior Bridge (PPB)

- **Improve the robustness of the implicit neural network to diverse $\mathbf{y}$**

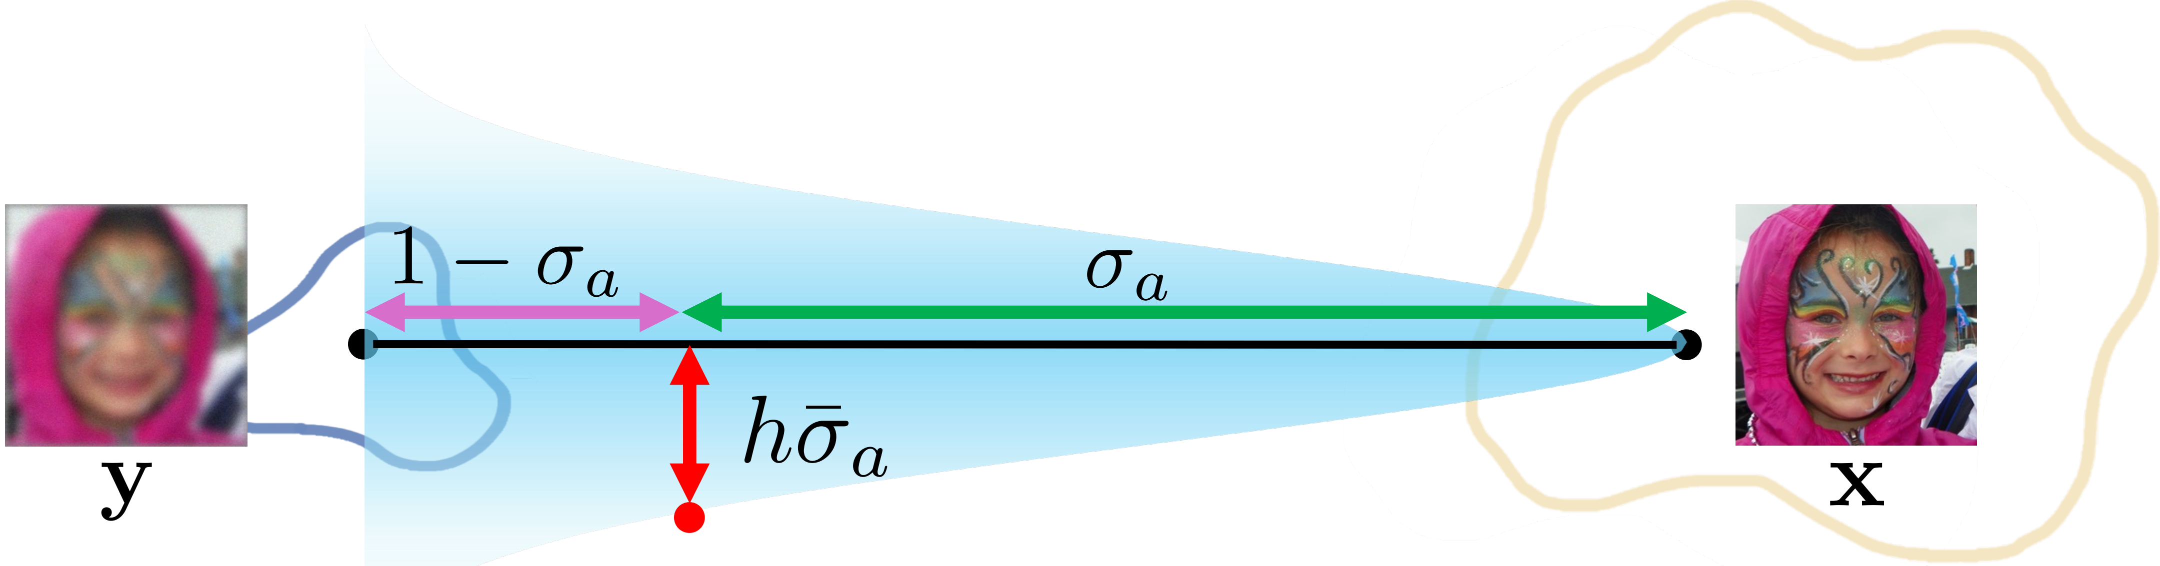

$$\mathbf{y}_a = (1-\sigma_a)\mathbf{y} + \sigma_a\mathbf{x} + h\bar{\sigma}_a\mathbf{z},\ \mathbf{z}\sim\mathcal{N}(\mathbf{0},\mathbf{I})$$

$$\sigma_a = \frac{\int_a^1 \beta_t dt}{\int_0^a \beta_t dt + \int_a^1 \beta_t dt}$$

Perturbation schedule $\bar{\sigma}_a = 1 - \Pi_{i=1}^{a}\beta_i$

hyperparameter a and h

Perturbed Posterior Bridge (PPB)

- **Intermediary set of trajectories between $\mathbf{y}$ and $\mathbf{x}$**

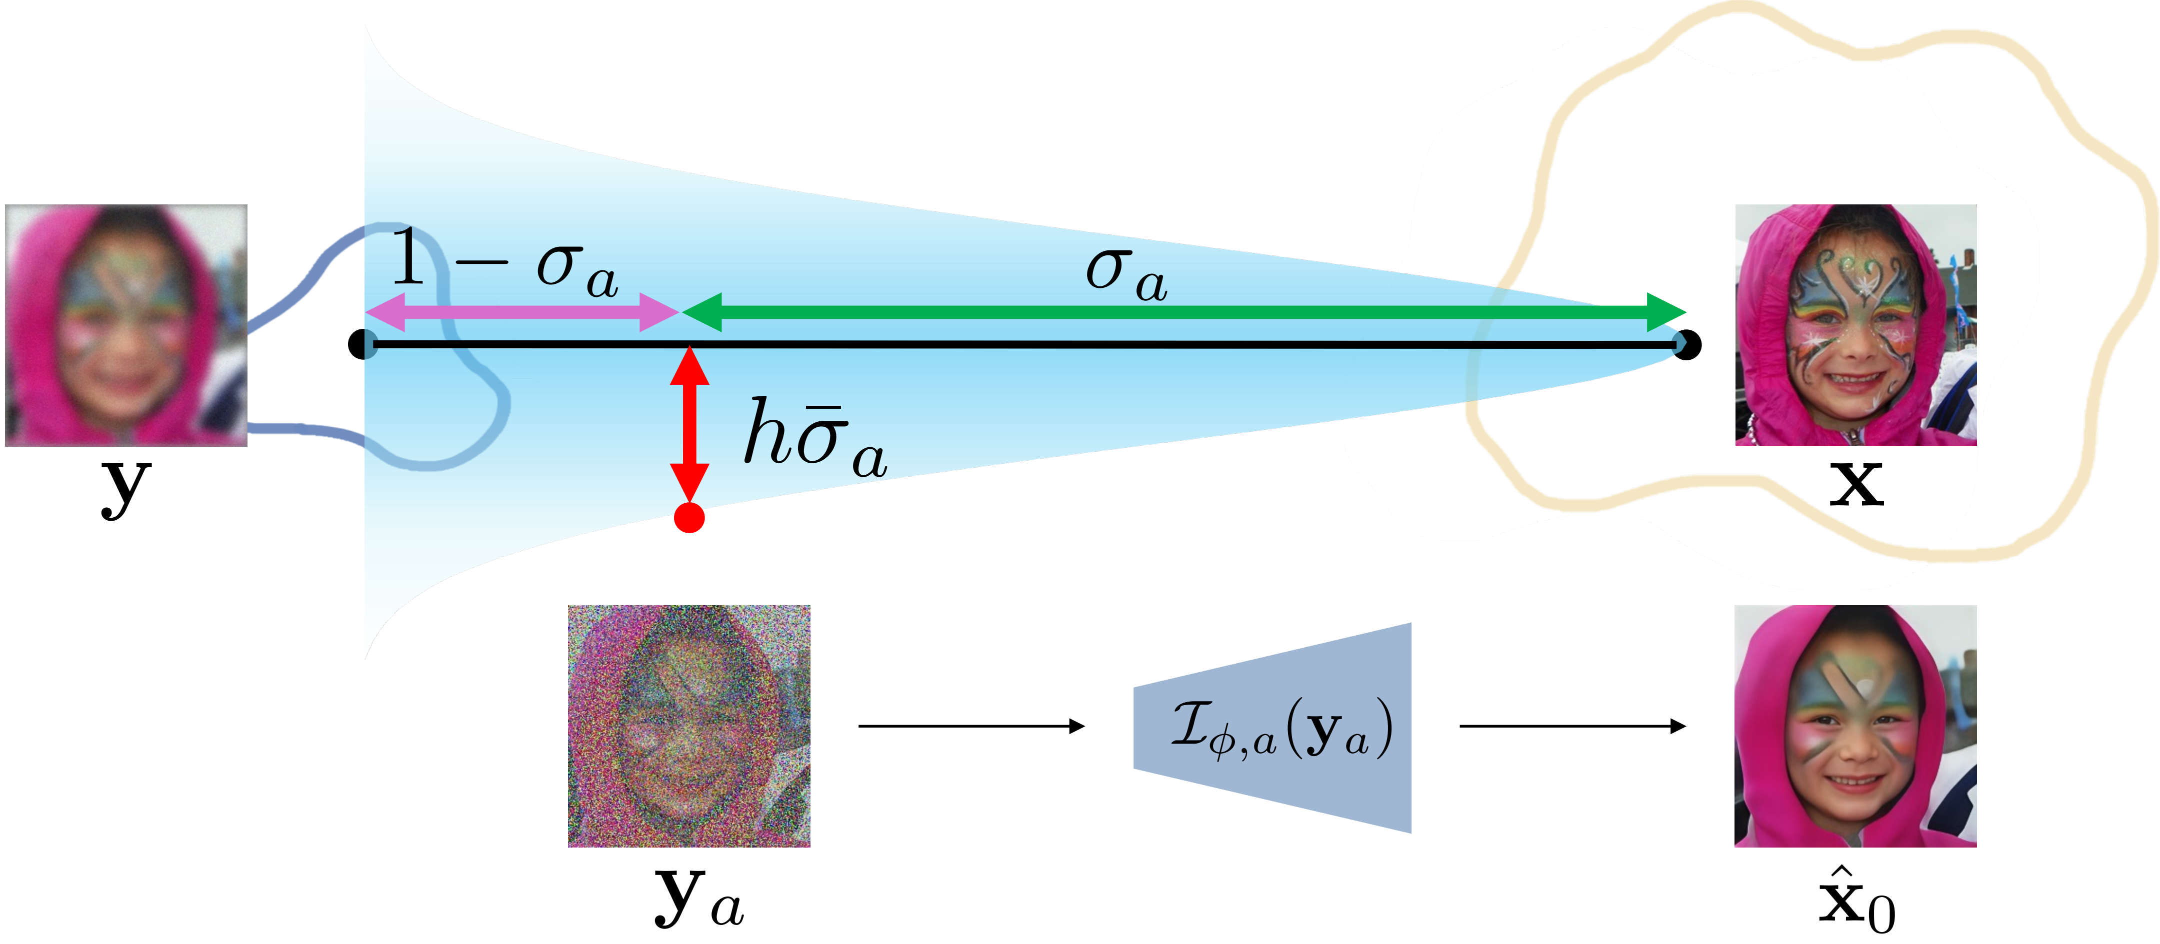

Out-of-distribution (OOD)

[Train] FFHQ / [Test] CelebA-HQ

[Train] ImageNet / [Test] GoPro

Experiments

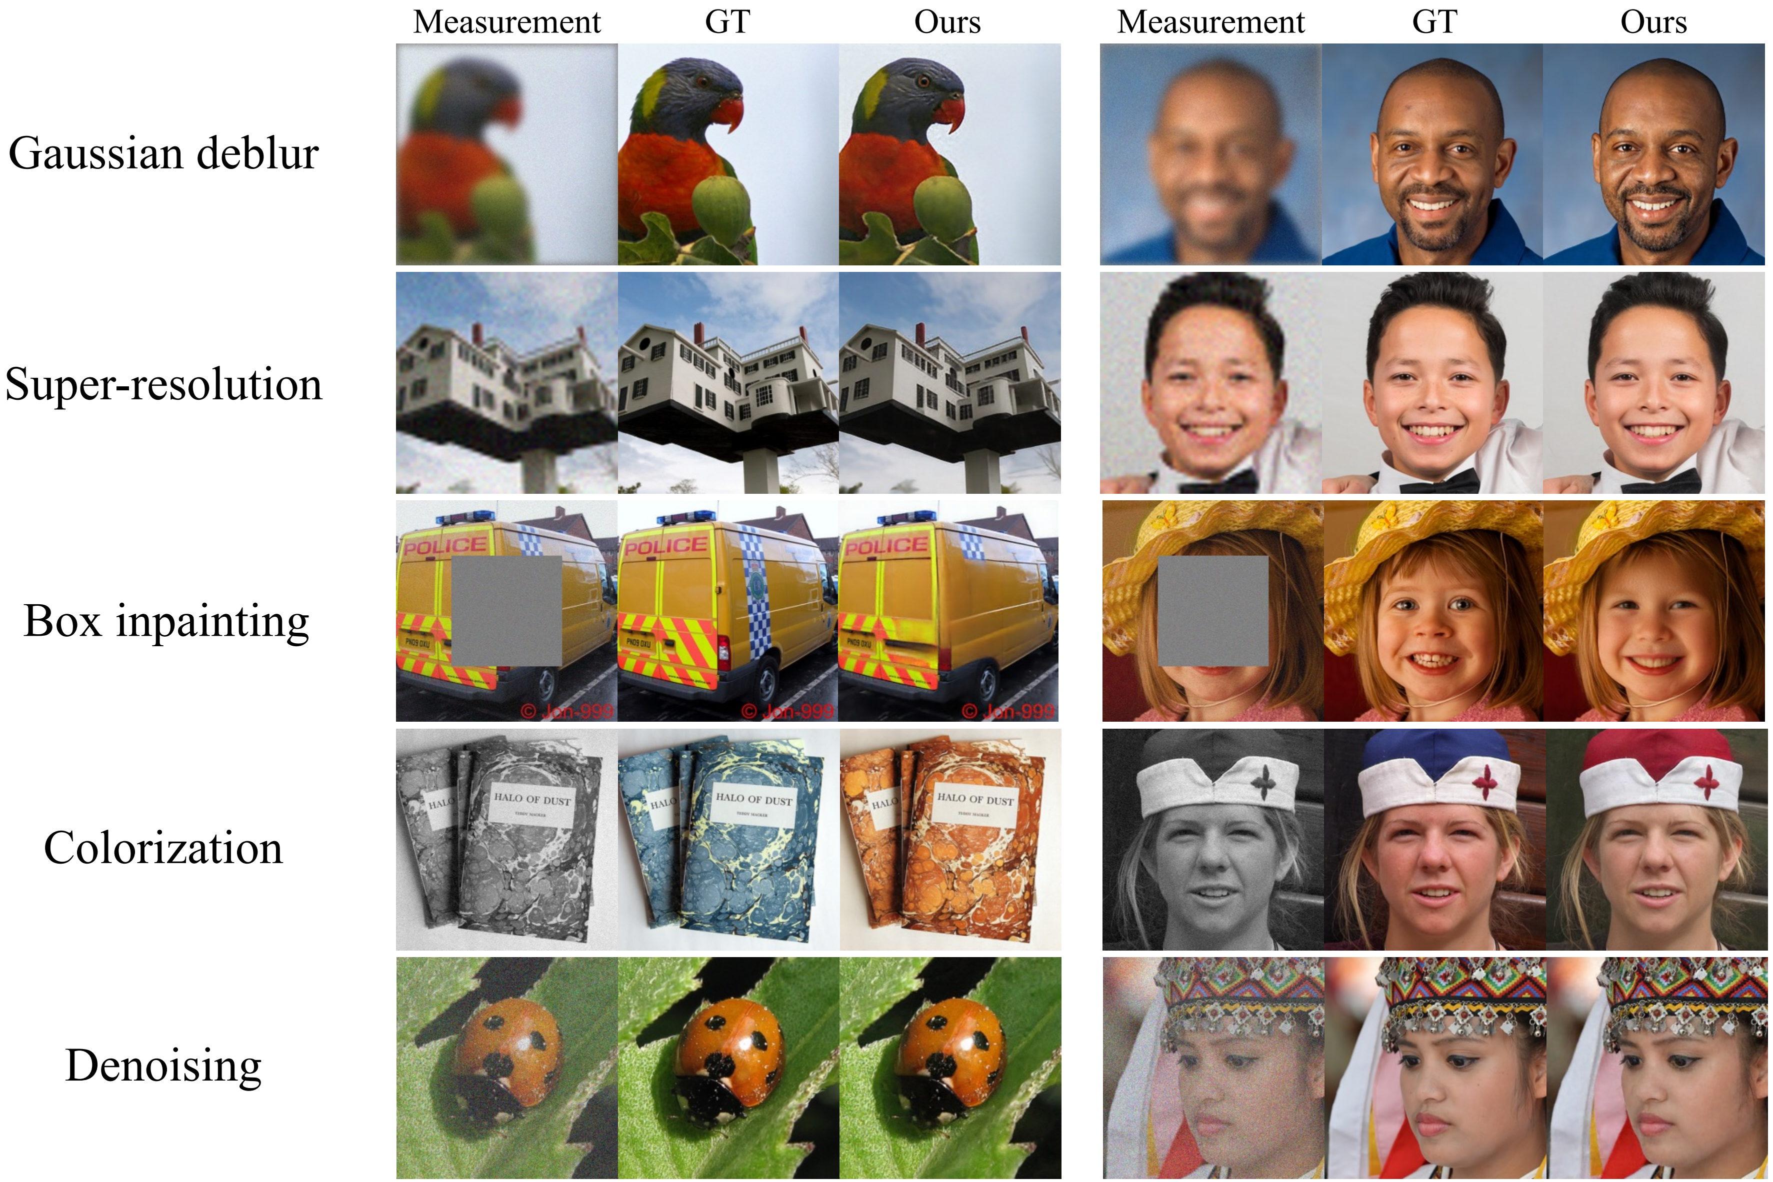

Conclusion

- **Diffusion prior-based Amortized Variational Inference (DAVI)**
- **Efficient posterior sampling by a single neural network evaluation**
- **Generalization** for both seen and unseen measurements **without any optimization at test-time**
- **Perturbed Posterior Bridge** further enhances the generalization capabilities

Thank you 😀

Paper

GitHub