

A personal journey toward

Diffusion Models for Inverse Problems in **Medical Imaging**

Jong Chul Ye

Professor
Graduate School of Artificial Intelligence
KAIST, Korea

Future of AI? Yann LeCun's Cake Analogy

■ "Pure" Reinforcement Learning (cherry)

- ▶ The machine predicts a scalar reward given once in a while.

- ▶ **A few bits for some samples**

■ Supervised Learning (icing)

- ▶ The machine predicts a category or a few numbers for each input
- ▶ Predicting human-supplied data
- ▶ **10→10,000 bits per sample**

■ Unsupervised/Predictive Learning (cake)

- ▶ The machine predicts any part of its input for any observed part.
- ▶ Predicts future frames in videos
- ▶ **Millions of bits per sample**




■ (Yes, I know, this picture is slightly offensive to RL folks. But I'll make it up)

Self-Supervised Denoising: Noise2X

Input




Noise2Noise



Output



Target




Self-Supervised Denoising: Noise2X

Input




Noise2Self



Output



Target



Self-Supervised Denoising: SURE

Soltanayev et al, NeurIPS, 2018

Self-supervised denoising using Stein Unsupervised Risk Estimator (SURE)

$$\ell_{SURE}(\Theta) = \mathbb{E}_{y \sim P_Y} d(y, F_{\Theta}(y))$$

$$d(y, F_{\Theta}(y)) := \underbrace{\|y - F_{\Theta}(y)\|^2}_{\text{Autoencoder loss}} + \underbrace{2\sigma^2 \operatorname{div}_y F_{\Theta}}_{\text{Divergence-based penalty}}$$

Autoencoder loss

Divergence-based
penalty

**Any unified mathematical
framework?**

Tweedie's Formula for Exponential Family Distribution

B Efron, Journal of the American Statistical Association, 2011

- **Probability distribution of exponential family:**

$$p(y|\eta) = p_0(y) \exp(\eta^\top T(y) - \varphi(\eta))$$



- **Using the Bayes' rule, the posterior density:**

$$p(\eta|y) = \exp(\eta^\top T(y) - \lambda(y)) [p(\eta) e^{-\varphi(\eta)}], \quad \text{where} \quad \lambda(y) = \log \frac{p(y)}{p_0(y)}$$

Tweedie's Formula for Exponential Family Distribution

Kim and Ye, NeurIPS, 2021

- Tweedie's formula → **Posterior mean**
- **The closed form solution for the posterior mean:**

$$\begin{aligned}\hat{\eta}^\top T'(y) &= \lambda'(y) = -\nabla_y \log p_0(y) + \nabla_y \log p(y) \\ &= -l'_0(y) + l'(y)\end{aligned}$$

$$l'(y) = \nabla_y \log p(y)$$

Score function

Tweedie's Formula for Exponential Family Distribution

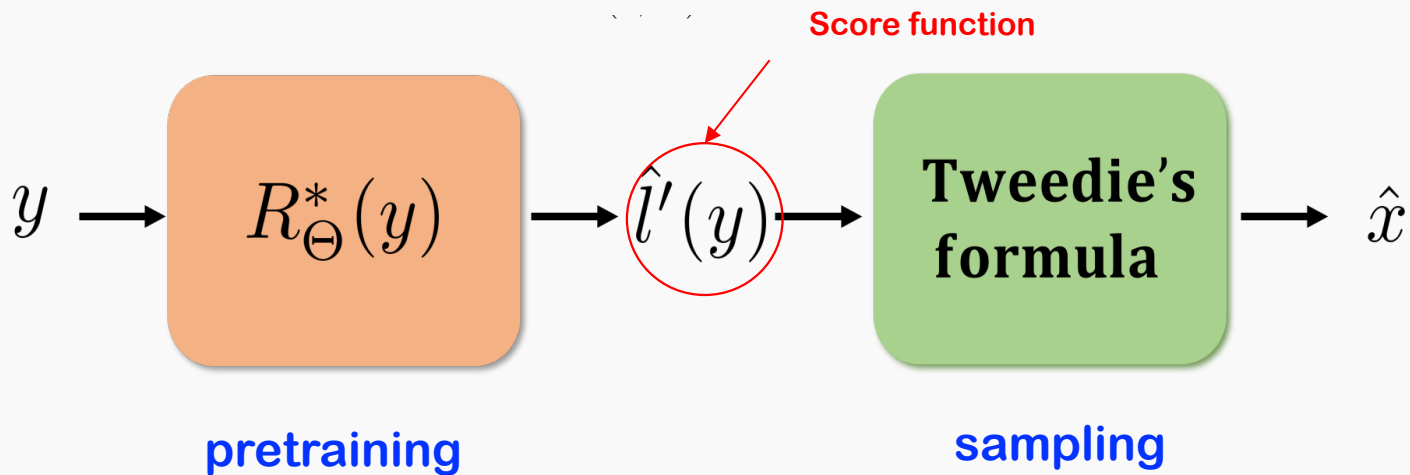
Kim and Ye, NeurIPS, 2021

Distribution	$p(y x)$	η	$T(y)$	$p_0(y)$	$l'_0(y)$	\hat{x}
Gaussian	$\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y-x)^2}{2\sigma^2}}$	x/σ^2	y	$\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{y^2}{2\sigma^2}}$	$-\frac{y}{\sigma^2}$	$y + \sigma^2 l'(y)$
Gaussian	$\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y-x)^2}{2\sigma^2}}$	$\left[\frac{x}{\sigma^2}, -\frac{1}{2\sigma^2}\right]^\top$	$[y, y^2]^\top$	$\frac{1}{\sqrt{2\pi}}$	0	$y + \sigma^2 l'(y)$
Poisson	$\frac{x^y e^{-x}}{y!}$	$\log(x)$	y	$\frac{1}{y!}$	$\simeq -\log\left(y + \frac{1}{2}\right)$	$\left(y + \frac{1}{2}\right) \exp(l'(y))$
Gamma(α, β)	$\frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{y}{x}\right)^{\alpha-1} e^{-\beta\frac{y}{x}}$	$\left[\alpha - 1, -\frac{\beta}{x}\right]^\top$	$[\log y, -y]^\top$	1	0	$\frac{\beta y}{(\alpha-1) - y l'(y)}$
Bernoulli	$x^y (1-x)^{(1-y)}$	$\log\left(\frac{x}{1-x}\right)$	y	1	0	$\frac{e^{l'(y)}}{1+e^{l'(y)}}$
Exponential	$x e^{-yx}, y \geq 0$	$-x$	y	1	0	$-l'(y)$

As long as we can compute the **score function**,
optimal denoising can be achieved by using Tweedie's formula.

Noise2Score

Kim and Ye, NeurIPS, 2021




**How to estimate the
score function?**

Score Function



$$p_{data}(x)$$



$$\nabla_x \log p_{data}(x)$$

Score function

Denosing Autoencoder (DAE)

Alain et al, JMLR, 2014


- DAE F_{Θ} is similar to **Noise2X**

$$\ell_{DAE}(\Theta) = \mathbb{E}_{\substack{y \sim P_Y \\ u \sim \mathcal{N}(0, I) \\ \sigma_a \sim \mathcal{N}(0, \delta^2)}} \|y - F_{\Theta}(y + \sigma_a u)\|^2$$

- DAE can be used to estimate the score function of data y

$$F_{\Theta^*}(y) = y + \sigma_a^2 l'(y) + o(\sigma_a^2)$$

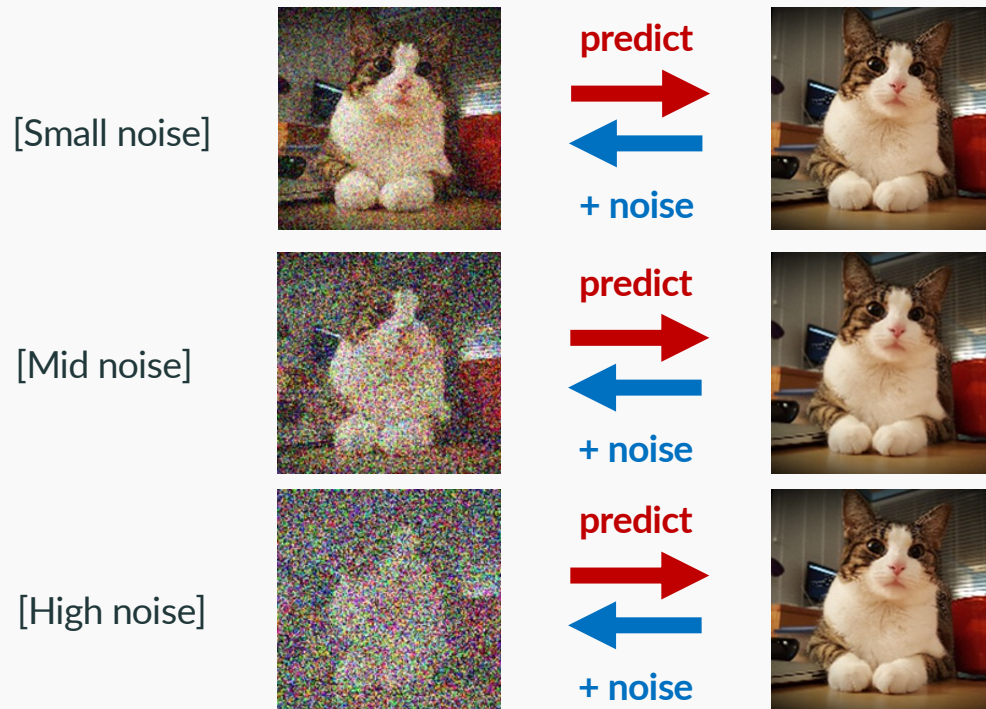
$$\sigma_a \rightarrow 0$$



$$F_{\Theta^*}(y) = y + \sigma_a^2 l'(y)$$

Equal to Tweedie's formula for Gaussian noises

Denoising Autoencoder (DAE)



Relation to SURE

By using residual form: $F_{\Theta}(y) = \sigma^2 R_{\Theta}(y) + y$

$$\hat{l}'(y) = \frac{F_{\Theta^*}(y) - y}{\sigma_a^2} = R_{\Theta}(y)$$

$$\begin{aligned} \ell_{SURE}(\Theta) &= \mathbb{E}_{y \sim P_Y} \left\{ \|y - F_{\Theta}(y)\|^2 + 2\sigma^2 \operatorname{div}_y F_{\Theta}(y) \right\} \\ &= \mathbb{E}_{y \sim P_Y} \left\{ \sigma^4 \|R_{\Theta}(y)\|^2 + 2\sigma^4 \operatorname{div}_y R_{\Theta}(y) \right\} + 2\sigma^2 \operatorname{dim}(y) \end{aligned}$$





$$\ell_{ISM}(\Theta) = \mathbb{E}_{y \sim P_Y} \left\{ \frac{1}{2} \|\Psi_{\Theta}(y)\|^2 + \operatorname{div}_y \Psi_{\Theta}(y) \right\}$$

Implicit Score matching cost by


Hyvärinen et al, JMLR, 2005

**Noise2X, SURE are
score-based approaches!**

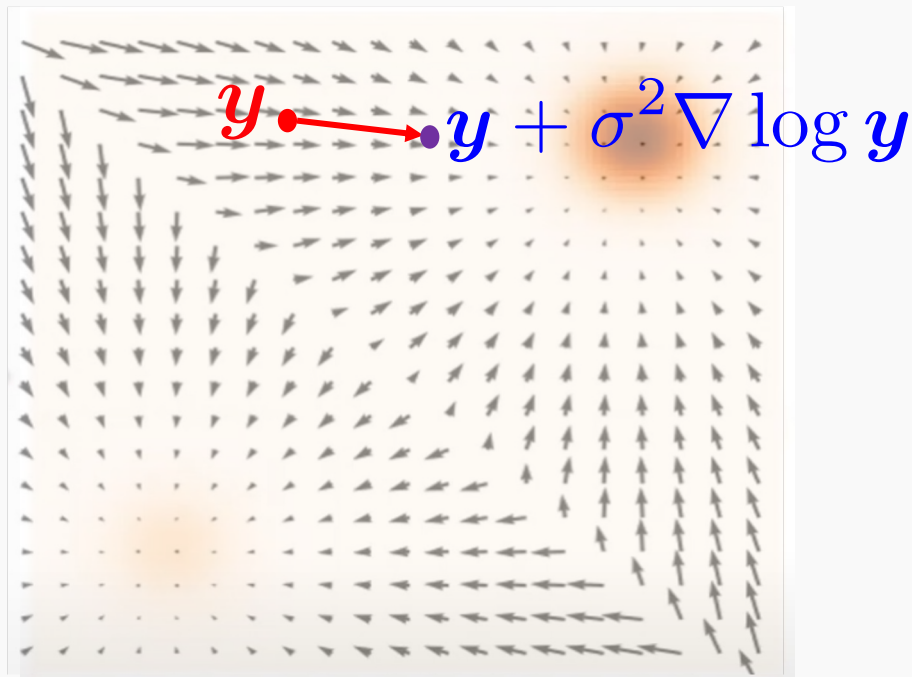
Gaussian ($\sigma = 50$)Poisson ($\zeta = 0.05$)Gamma ($k = 50$)

Generalization beyond the
one-step denoising?

Noise2Score vs Diffusion Models




Diffusion model



Noise2Score

Method1: Score-method with Langevin Dynamics (SMLD)



- Once the score model is trained to optimality,
 - i.e. $s_\theta(\mathbf{x}) \simeq \nabla_{\mathbf{x}} p(\mathbf{x})$
- Example: Use **Langevin dynamics** to draw samples

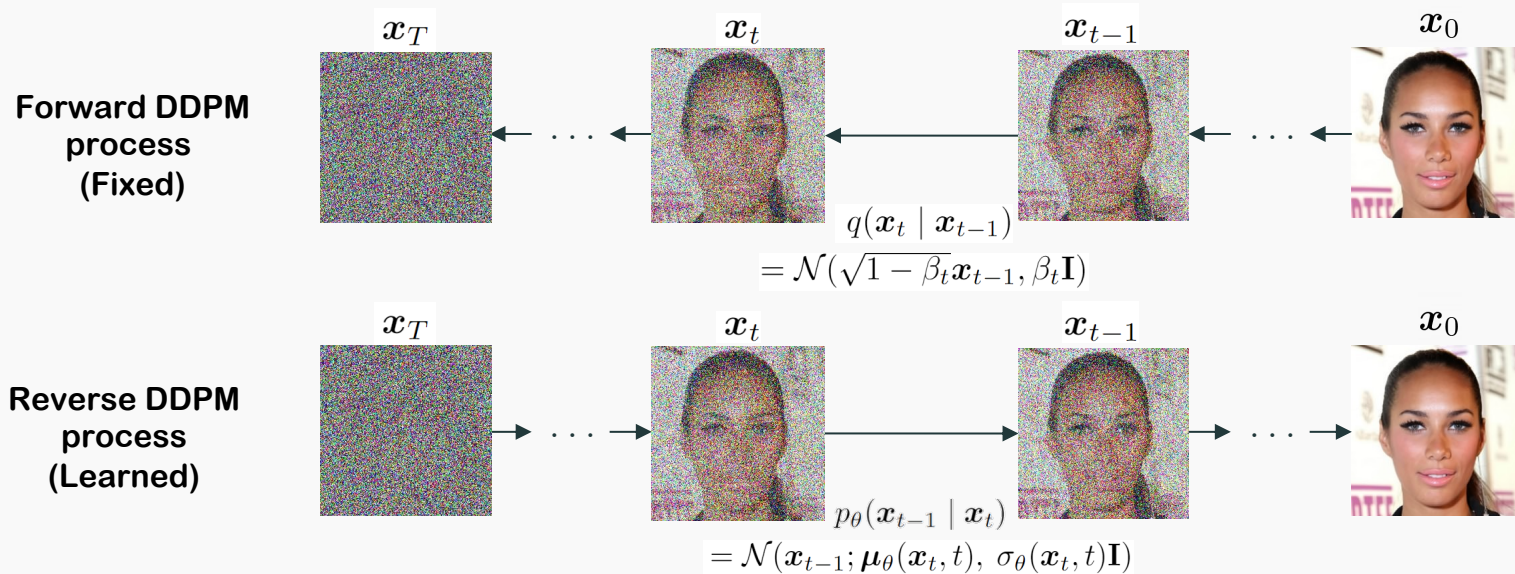
$$\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i + \epsilon \nabla_{\mathbf{x}} \log p(\mathbf{x}) + \sqrt{2\epsilon} \mathbf{z}_i$$

$$i = 0, 1, \dots, K$$

“Variance Exploding (VE)” parameterization

Method2: Diffusion Denoising Probabilistic Model (DDPM)

(Ho et al. NeurIPS 2020)



DDPM Training & Sampling

Algorithm 1 Training


- 1: **repeat**
- 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
- 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
- 4: $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 5: Take gradient descent step on
 $\nabla_{\theta} \left\| \epsilon - \epsilon_{\theta}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|^2$
- 6: **until** converged

Algorithm 2 Sampling

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 2: **for** $t = T, \dots, 1$ **do**
- 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$
- 5: **end for**
- 6: **return** \mathbf{x}_0


“Variance Preserving (VP)” parameterization

Cf. Noise2Score



**How to use for inverse
problems?**

Typical Setup of Inverse Problems



Biomedical imaging systems cast as the above **forward model**

- Acquiring the image x : **inverse problem**
- System naturally **ill-posed**: what is the best solution?
- Ex: microscopy, MRI, CT, optics, etc.

Score-based Diffusion Models for Accelerated MRI

Jung et al, MEDIA, 2022

$$\min_{\mathbf{x}} \|\mathbf{y} - A\mathbf{x}\|^2$$

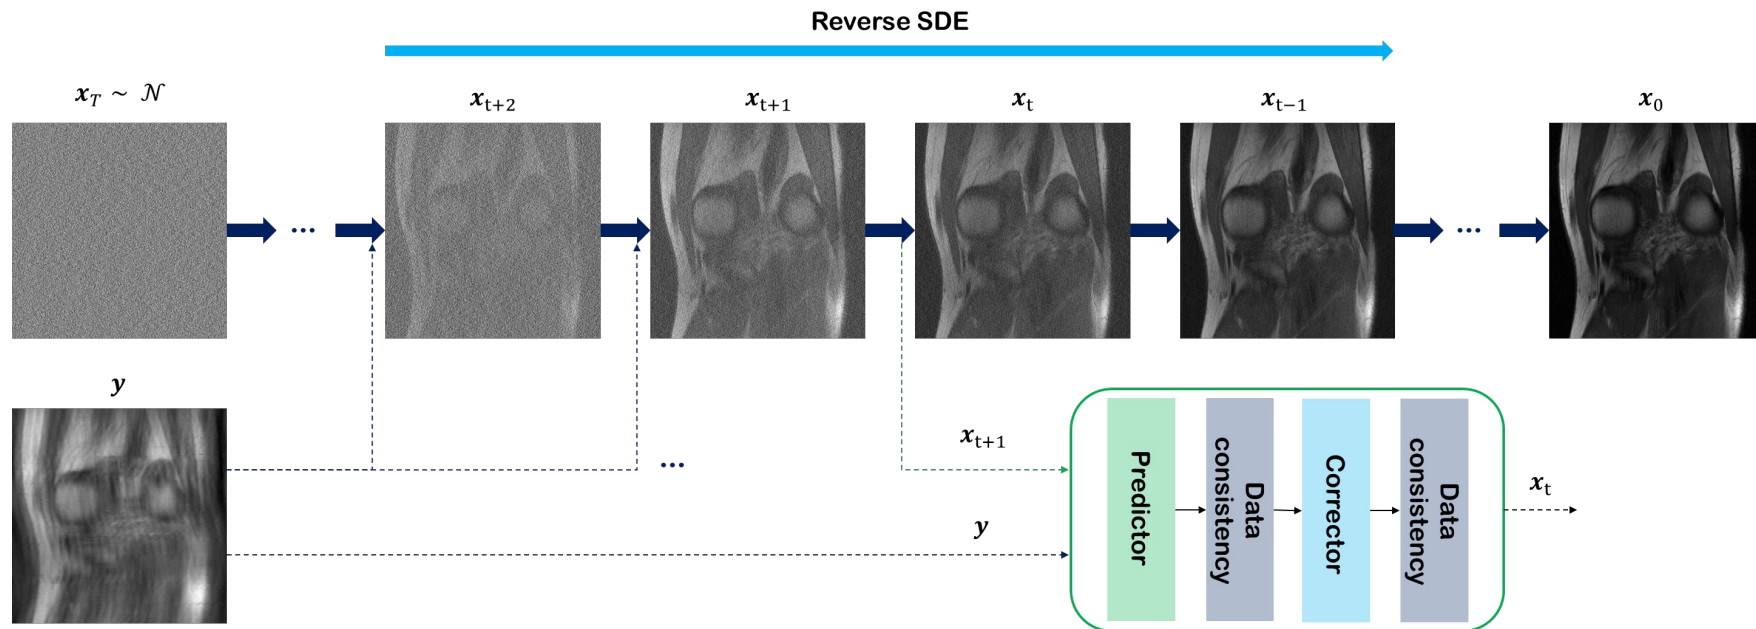
$$\mathbf{x}_i \leftarrow \mathbf{x}_{i+1} + \epsilon_i \mathbf{s}_\theta(\mathbf{x}_{i+1}, \sigma_{i+1}) + \sqrt{2\epsilon_i} \mathbf{z}$$

Denoising step (reverse SDE)

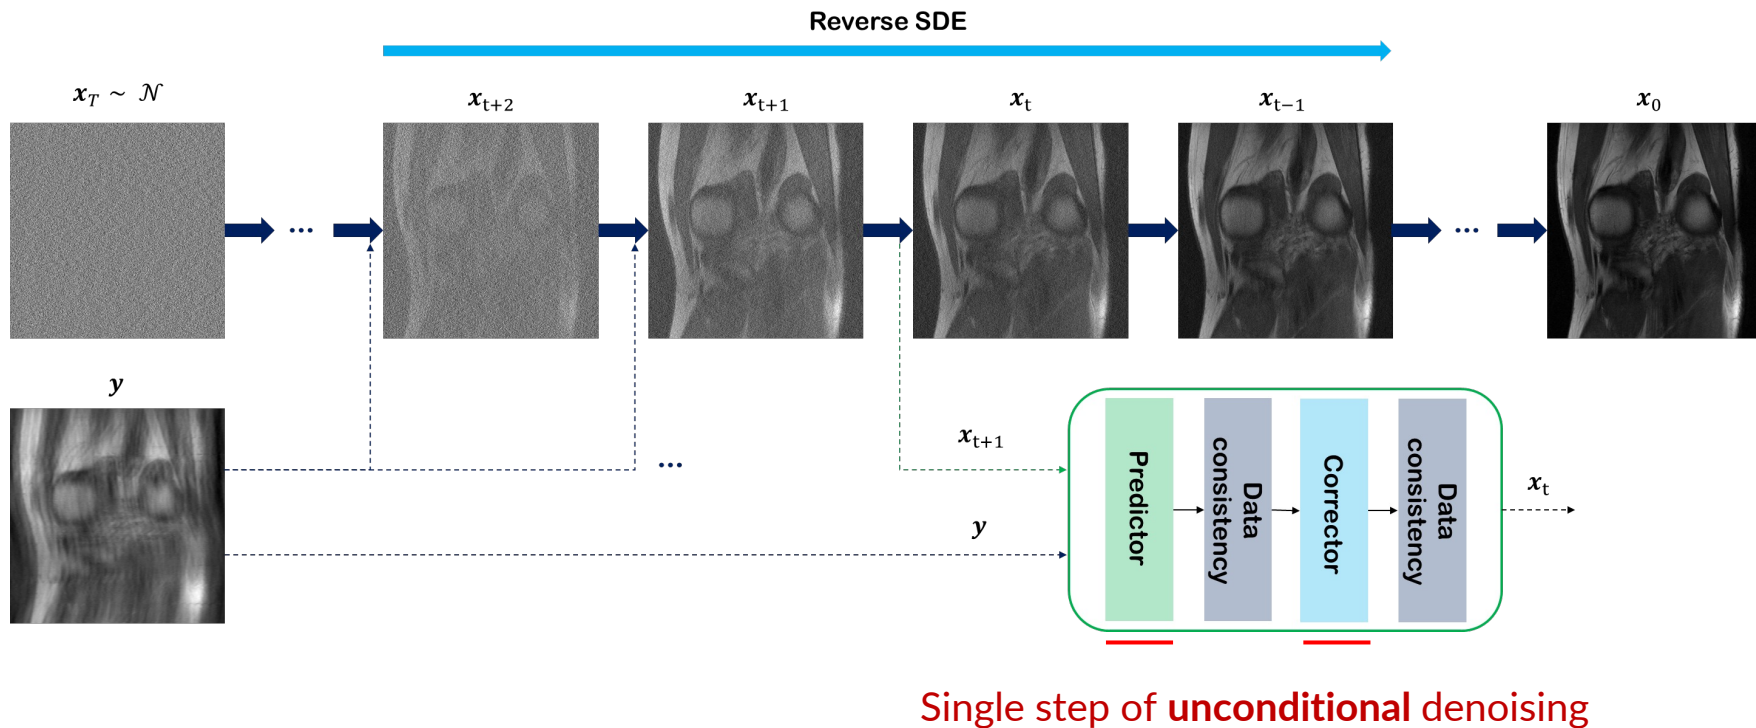
$$\mathbf{x}_i \leftarrow \mathbf{x}_i + \lambda A^*(\mathbf{y} - A\mathbf{x}_i),$$

Data consistency step (e.g. GD, POCS)

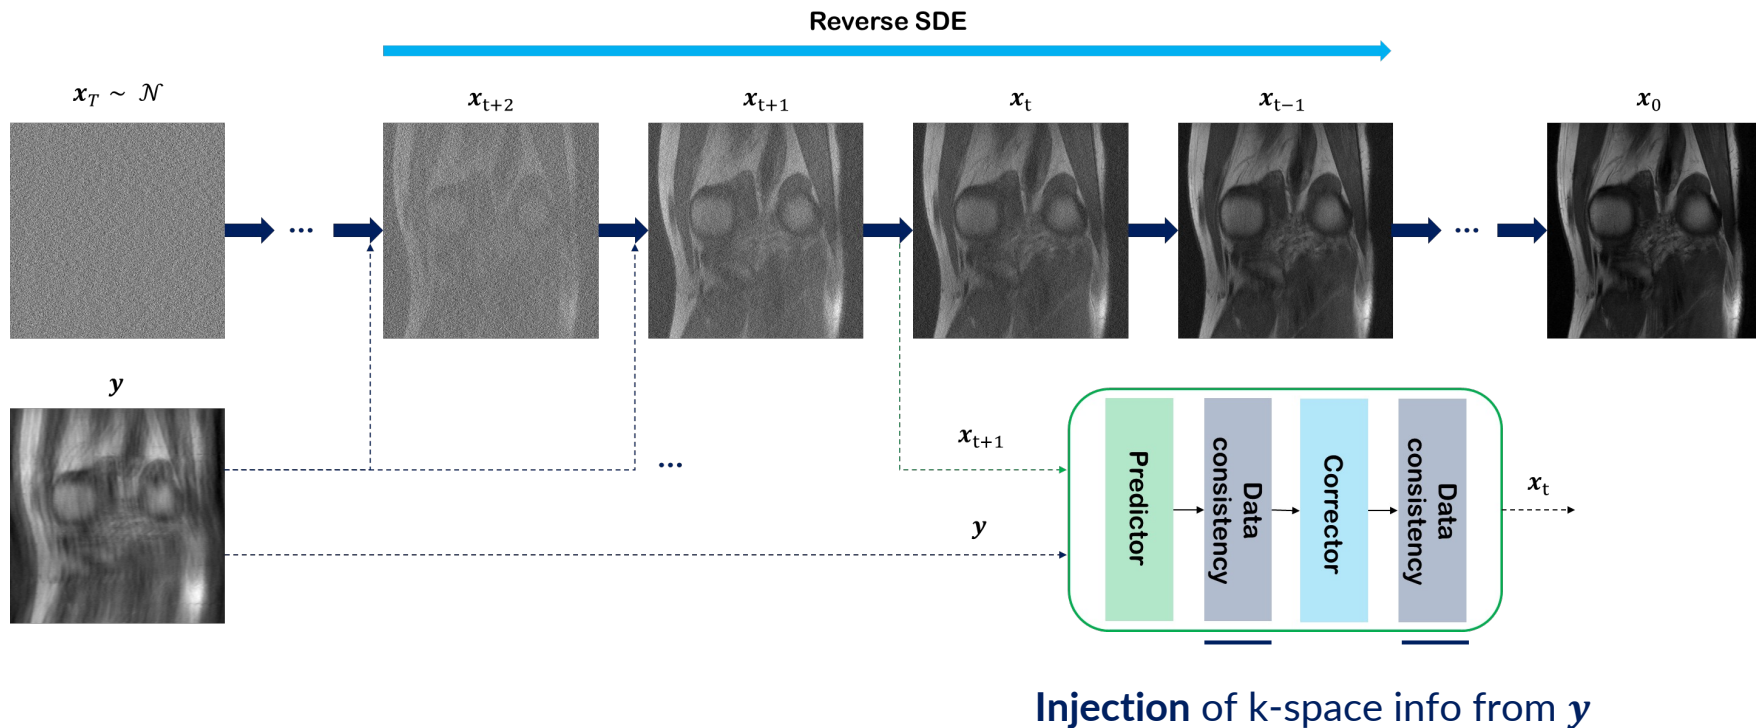
Score-based Diffusion Models for Accelerated MRI



Score-based Diffusion Models for Accelerated MRI

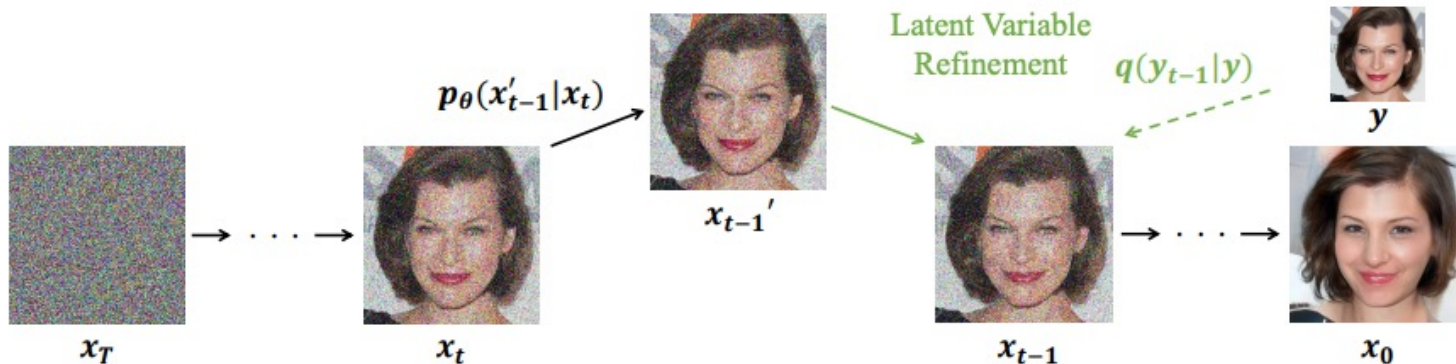


Score-based Diffusion Models for Accelerated MRI



ILVR: Conditioning Method for DDPM

Choi et al, ICLR, 2021

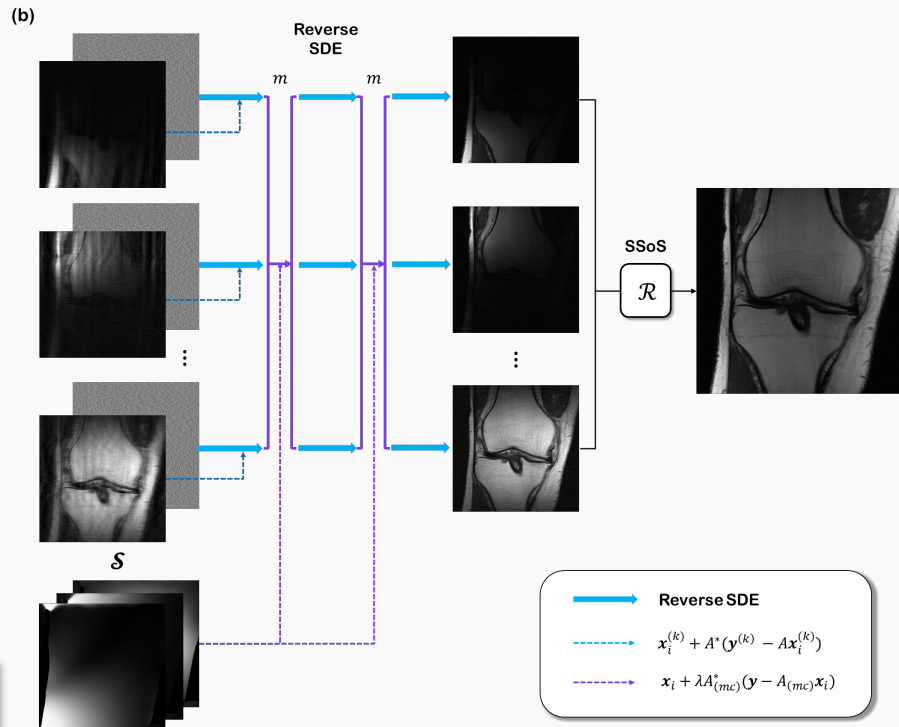
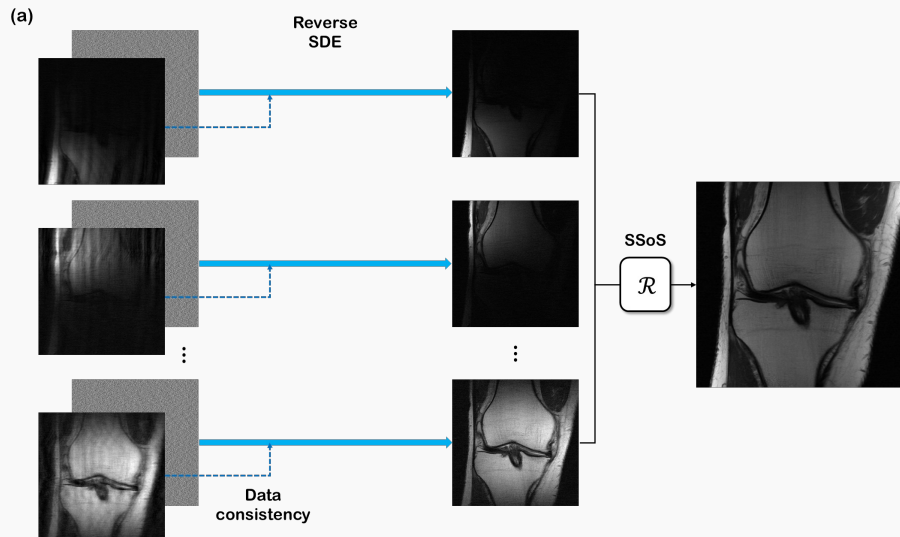


$$x'_{t-1} \sim p_\theta(x'_{t-1}|x_t),$$

$$x_{t-1} = \phi(y_{t-1}) + (I - \phi)(x'_{t-1}).$$

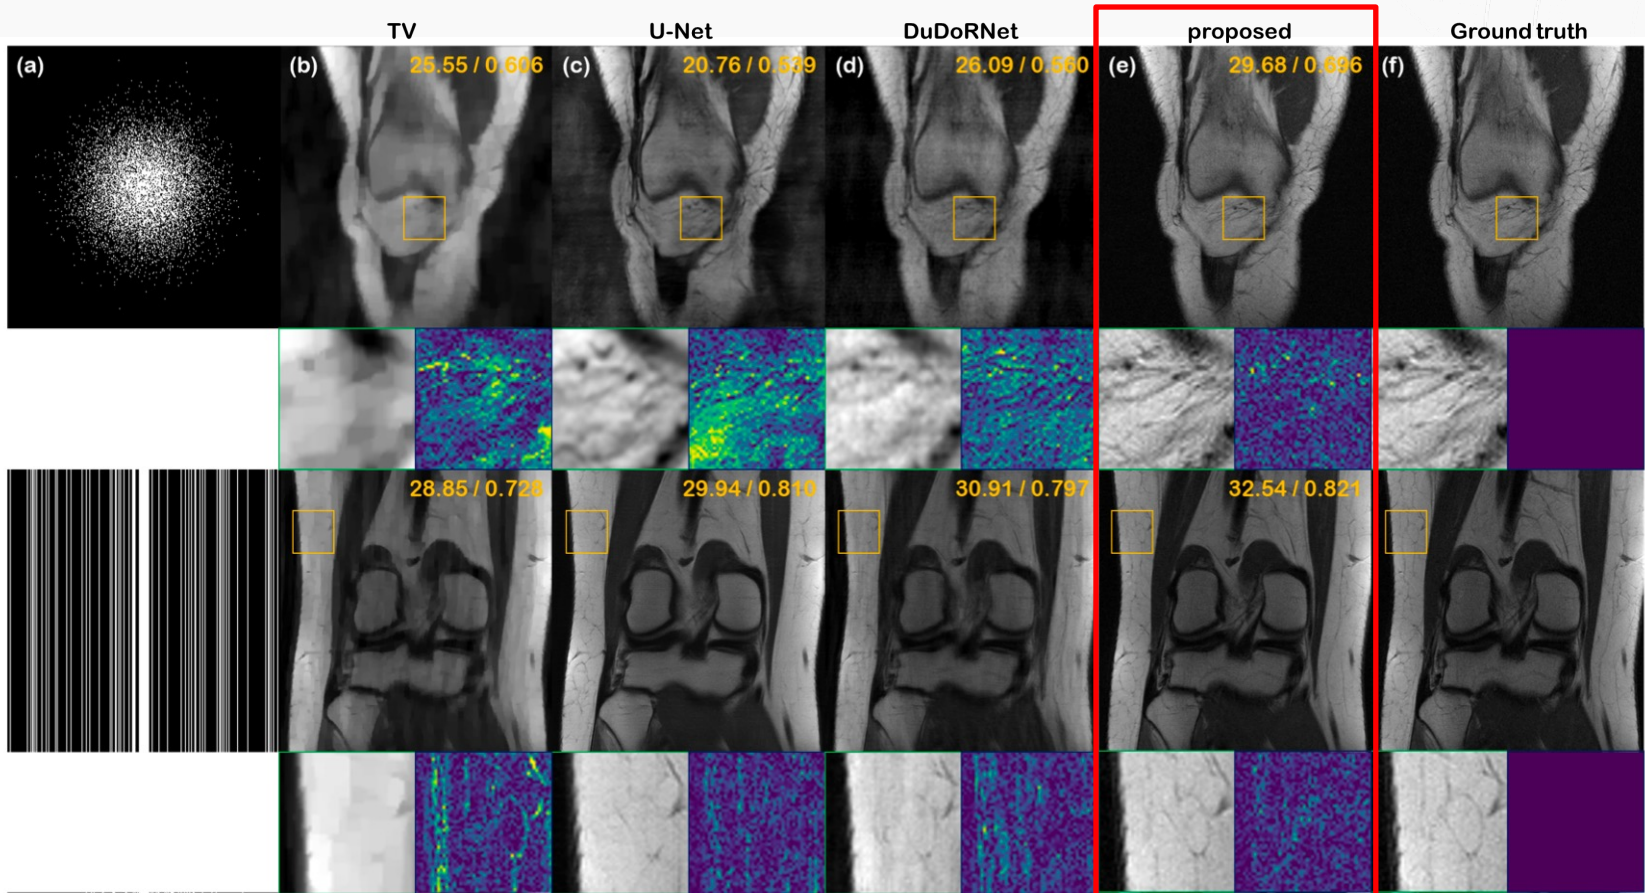
- Denoising step (reverse SDE)
- Data consistency step

Parallel Imaging

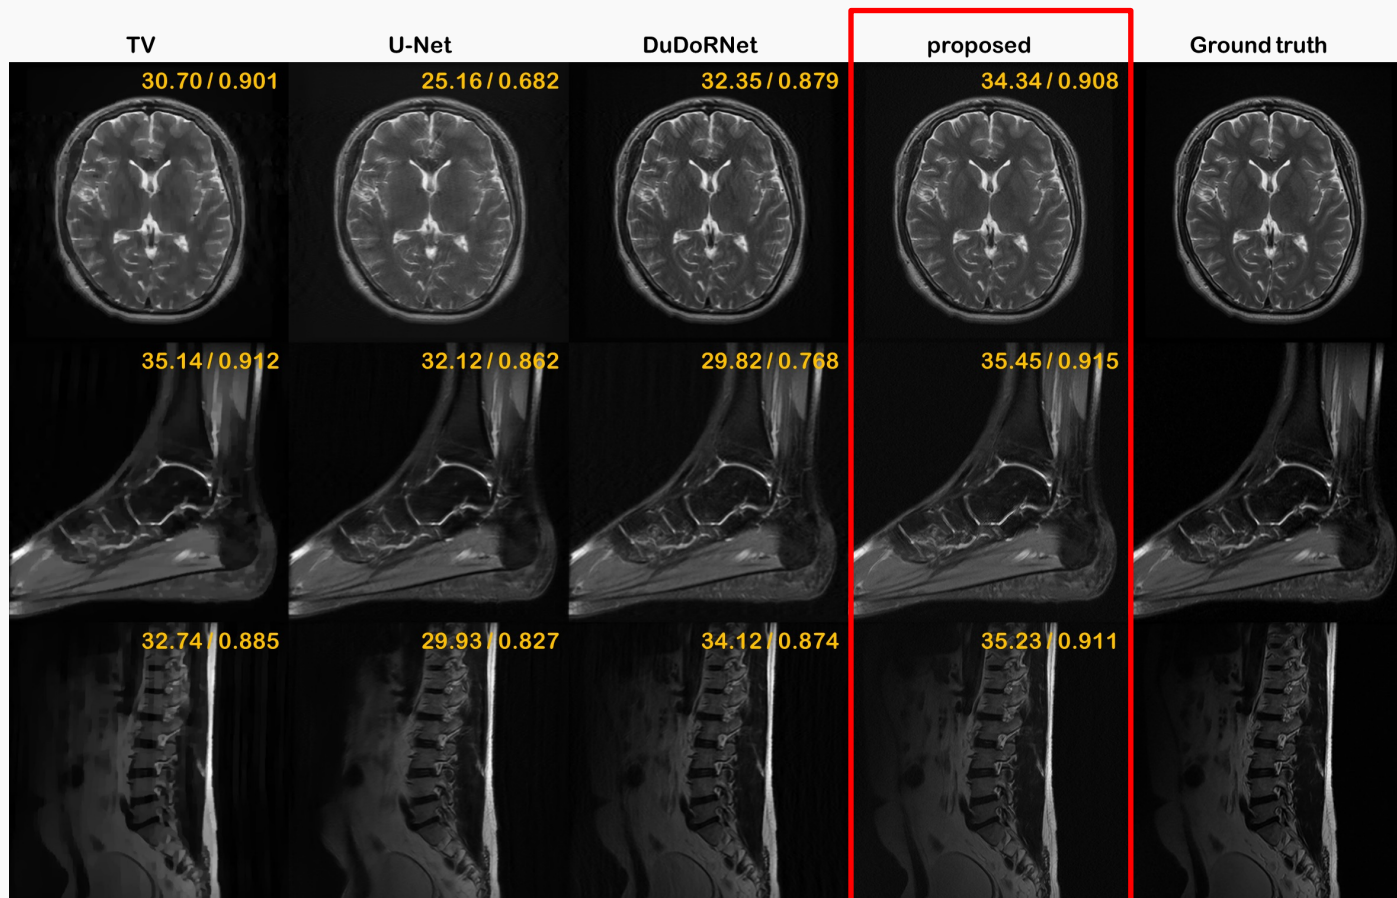


Diffusion model is trained using only DICOM files
--> no k-space data is necessary

State-of-the-art Performance



Generalization Capability



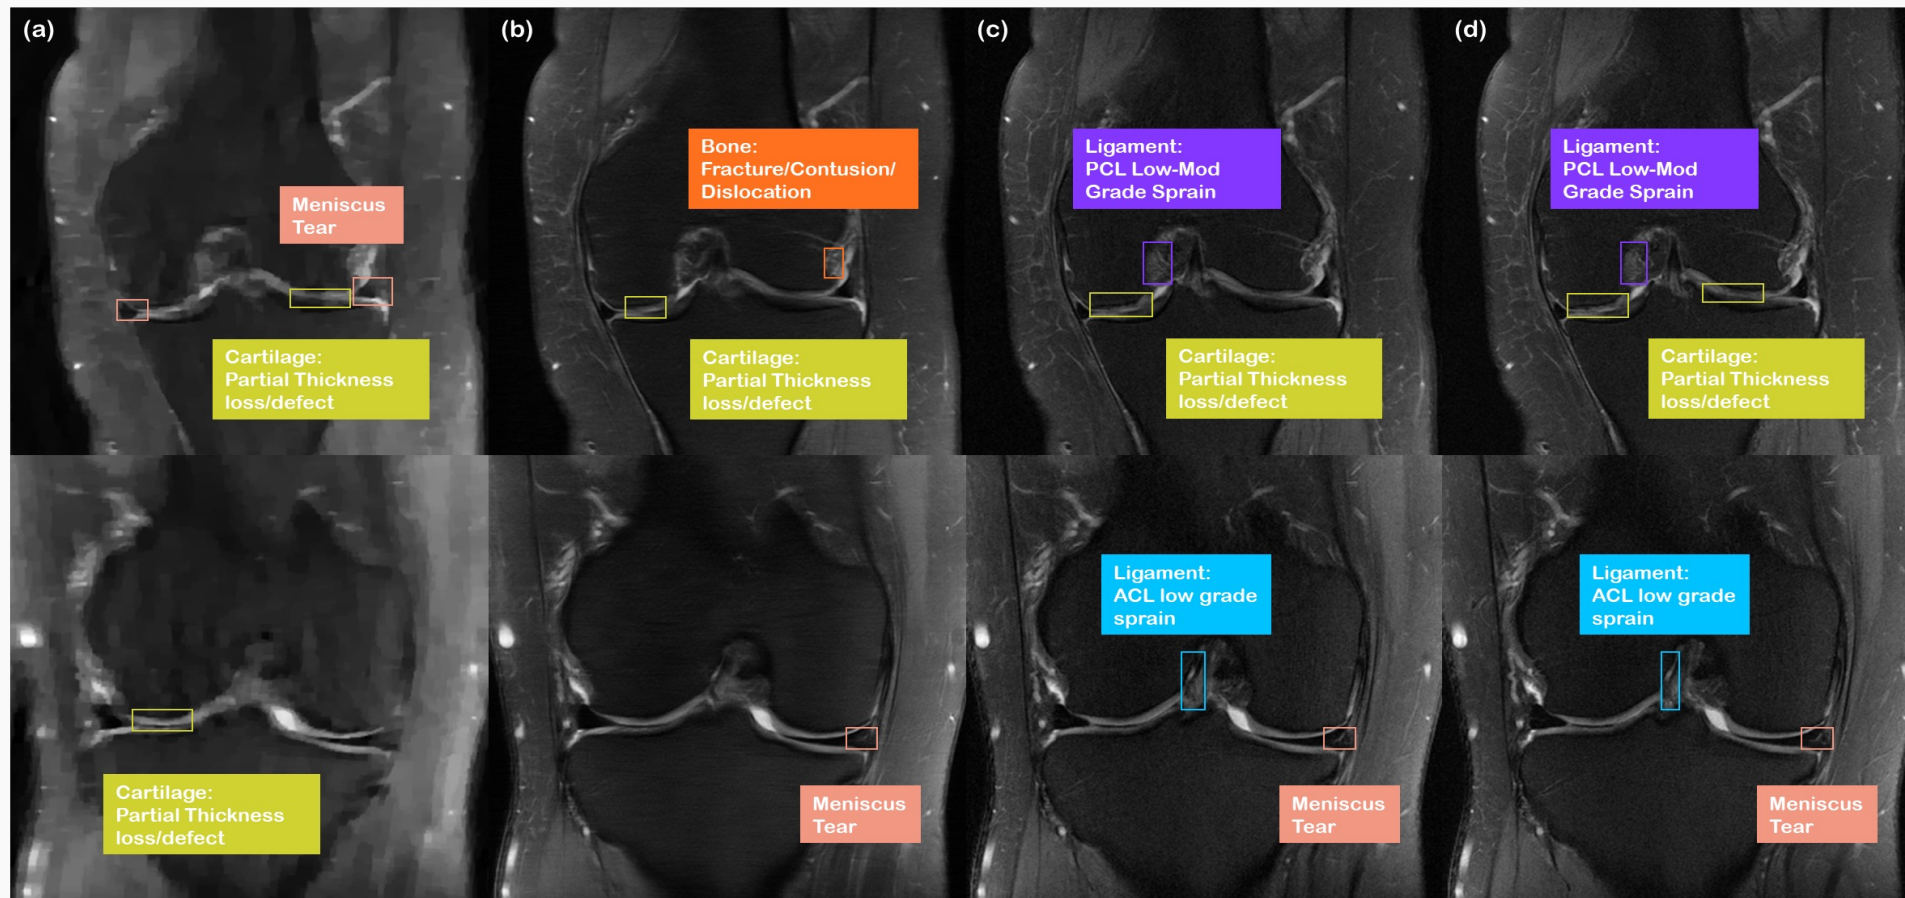
Diagnostic Capability

CS


DL

Diffusion

Fully-sampled



Uncertainty Quantification




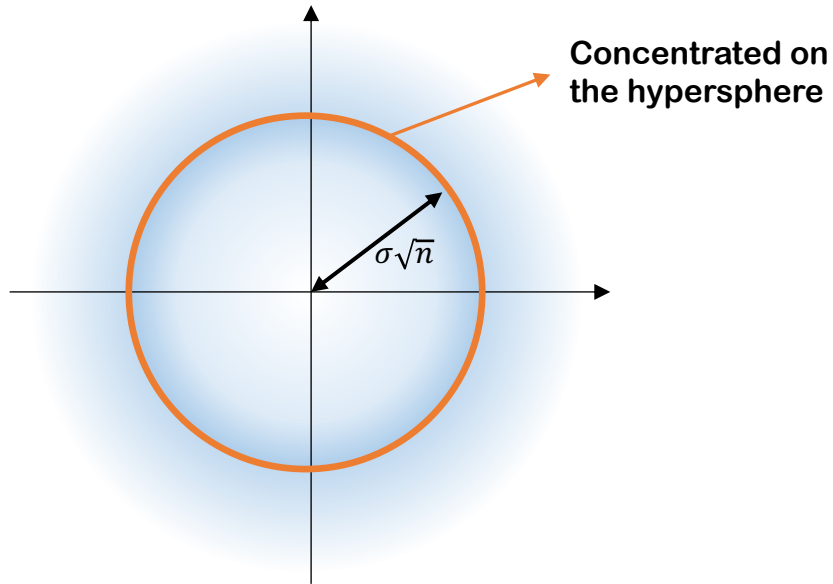
How to further improve?

Chung et al, NeurIPS, 2022; ICLR 2023

Concentration of Measures

RV: $X = (X_1, X_2, \dots, X_n)$, $X_i \sim \mathcal{N}(0, \sigma^2)$

$$\frac{\|X\|^2}{n} = \frac{X_1^2 + X_2^2 + \dots + X_n^2}{n} \rightarrow \sigma^2 \text{ (Law of Large number)}$$

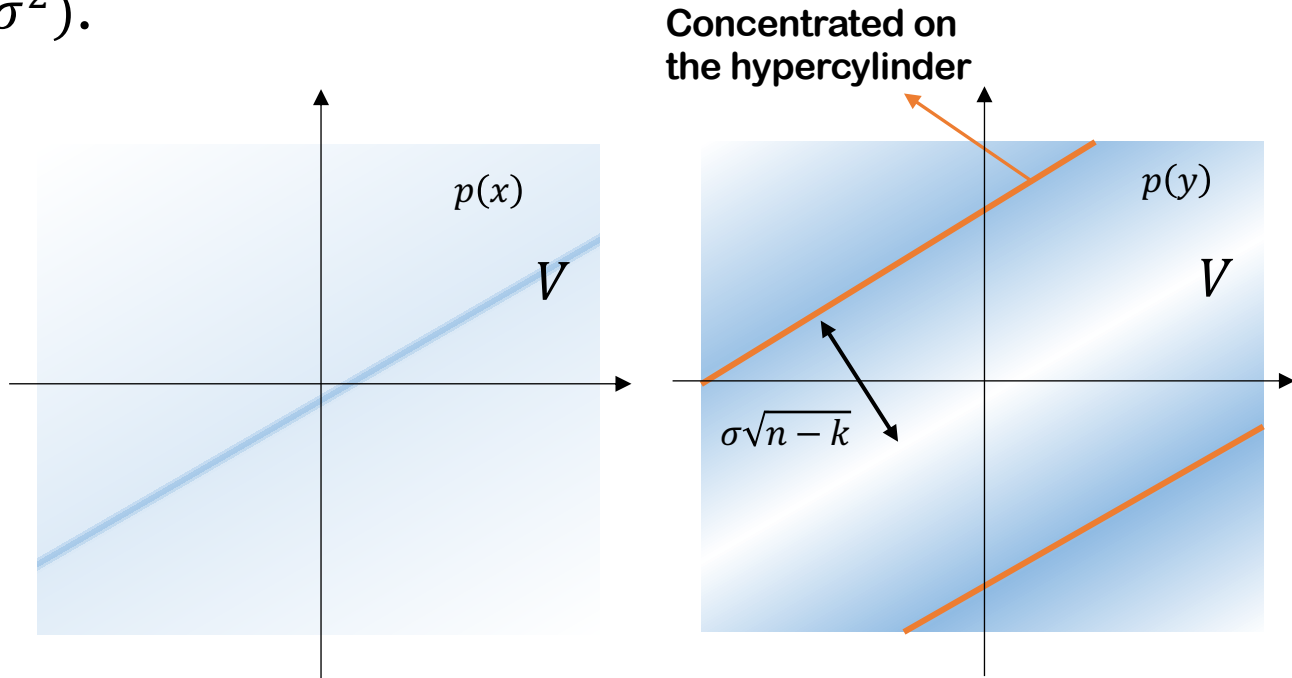


Concentration of Measures


$p(x) > 0 \leftrightarrow x \in V$, V : k dimensional linear space

When $p(y|x) = \mathcal{N}(x, \sigma^2)$.


Marginal $p(y)$?




Concentration on Manifolds




Forward Diffusion



Reverse Diffusion



Diffusion step is the transition from one manifold to another



DPS: DIFFUSION POSTERIOR SAMPLING

Chung et al, ICLR, 2023.


General Non-Linear Inverse Problems

Case 1. Linear


$$\mathcal{A}(\mathbf{x}) \equiv A\mathbf{x}$$

Case 2. Phase retrieval (Non-linear)

$$\mathcal{A}(\mathbf{x}) \equiv |\mathcal{F}\mathbf{x}|$$



$\mathcal{F}, |\cdot|$




Case 3. Neural network

$$\mathcal{A}(\mathbf{x}) \equiv G_{\phi}(\mathbf{x})$$



G_{ϕ}



Diffusion using Posterior Sampling


Suppose that our goal is to generate based on some condition

$$\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log q(\mathbf{y} | \mathbf{x}_t) \quad (\text{Bayes' rule})$$

Posterior sampling using diffusion model:

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2} \mathbf{x} - \beta(t) (\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \mathbf{x}_t)) \right] dt + \sqrt{\beta(t)} d\bar{\mathbf{w}},$$


The Devil is in the Likelihood



$$p(\mathbf{y}|\mathbf{x}_0) = \mathcal{N}(\mathbf{y}|\mathcal{A}\mathbf{x}_0, \sigma^2\mathbf{I})$$

$$\log p(\mathbf{y}|\mathbf{x}_0) = -\|\mathbf{y} - \mathcal{A}\mathbf{x}_0\|_2^2/\sigma^2$$

The Devil is in the Likelihood



$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$$

Intractable!

$$p(\mathbf{y}|\mathbf{x}_0) = \mathcal{N}(\mathbf{y}|\mathcal{A}\mathbf{x}_0, \sigma^2\mathbf{I})$$

$$\log p(\mathbf{y}|\mathbf{x}_0) = -\|\mathbf{y} - \mathcal{A}\mathbf{x}_0\|_2^2/\sigma^2$$

Tweedie's Formula for Approximation

$$p(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \int p(\mathbf{y}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \mathbb{E}_{p(\mathbf{x}_0|\mathbf{x}_t)}[p(\mathbf{y}|\mathbf{x}_0)]$$

Tweedie's Formula for Approximation


$$p(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \int p(\mathbf{y}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \mathbb{E}_{p(\mathbf{x}_0|\mathbf{x}_t)}[p(\mathbf{y}|\mathbf{x}_0)]$$

$$\simeq p(\mathbf{y}|\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$$

Push
expectation
inside



Tweedie's Formula for Approximation

$$p(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \int p(\mathbf{y}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$$

$$= \mathbb{E}_{p(\mathbf{x}_0|\mathbf{x}_t)}[p(\mathbf{y}|\mathbf{x}_0)]$$

$$\simeq p(\mathbf{y}|\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$$

$$= p(\mathbf{y}|\hat{\mathbf{x}}_0)$$



Jensen Bound for the Approximation

Theorem 1. *For the given measurement model (6) with $\mathbf{n} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, we have*

$$p(\mathbf{y}|\mathbf{x}_t) \simeq p(\mathbf{y}|\hat{\mathbf{x}}_0), \quad (13)$$

where the approximation error can be quantified with the Jensen gap, which is upper bounded by

$$\mathcal{J} \leq \frac{d}{\sqrt{2\pi\sigma^2}} e^{-1/2\sigma^2} \|\nabla_{\mathbf{x}} \mathcal{A}(\mathbf{x})\|_{m_1}, \quad (14)$$

where $\|\nabla_{\mathbf{x}} \mathcal{A}(\mathbf{x})\| := \max_{\mathbf{x}} \|\nabla_{\mathbf{x}} \mathcal{A}(\mathbf{x})\|$ and $m_1 := \int \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\| p(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0$.

Diffusion Posterior Sampling

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t)$$

Diffusion Posterior Sampling

$$\begin{aligned}\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) &= \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t) \\ &\simeq \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0)\end{aligned}$$

Theorem 1.

Diffusion Posterior Sampling

$$\begin{aligned}\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) &= \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t) \\ &\simeq \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0) \\ &\simeq s_{\theta^*}(\mathbf{x}_t, t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0)\end{aligned}$$

1. Gaussian

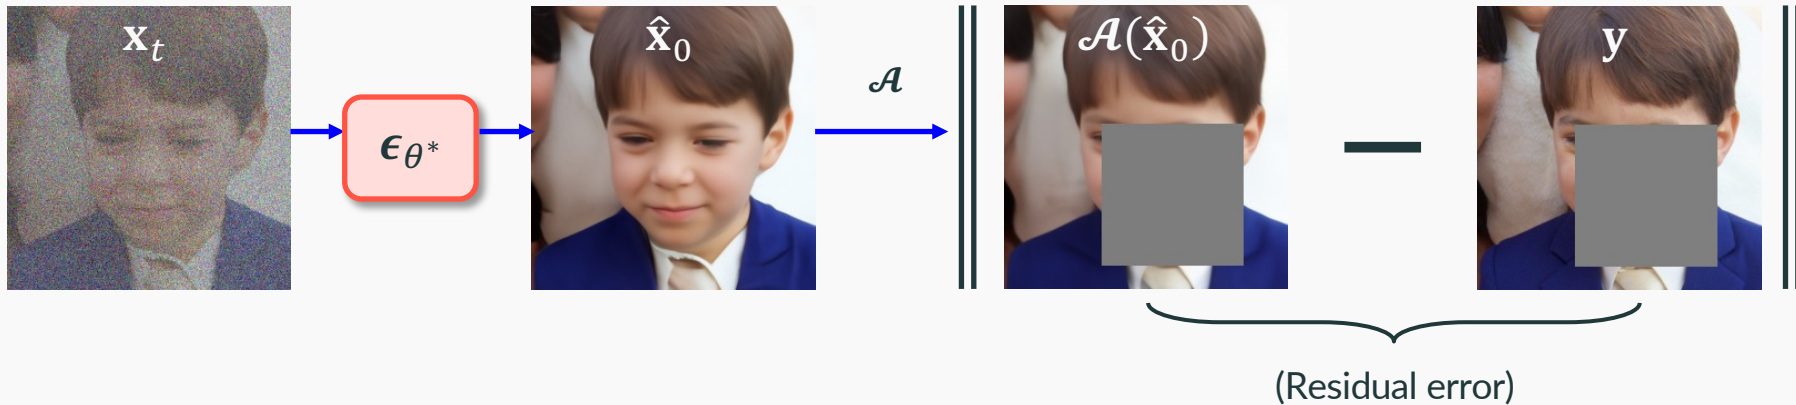
$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0) = -\rho \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$$

2. Poisson

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0) \simeq -\rho \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_{\Lambda}^2$$

Diffusion Posterior Sampling

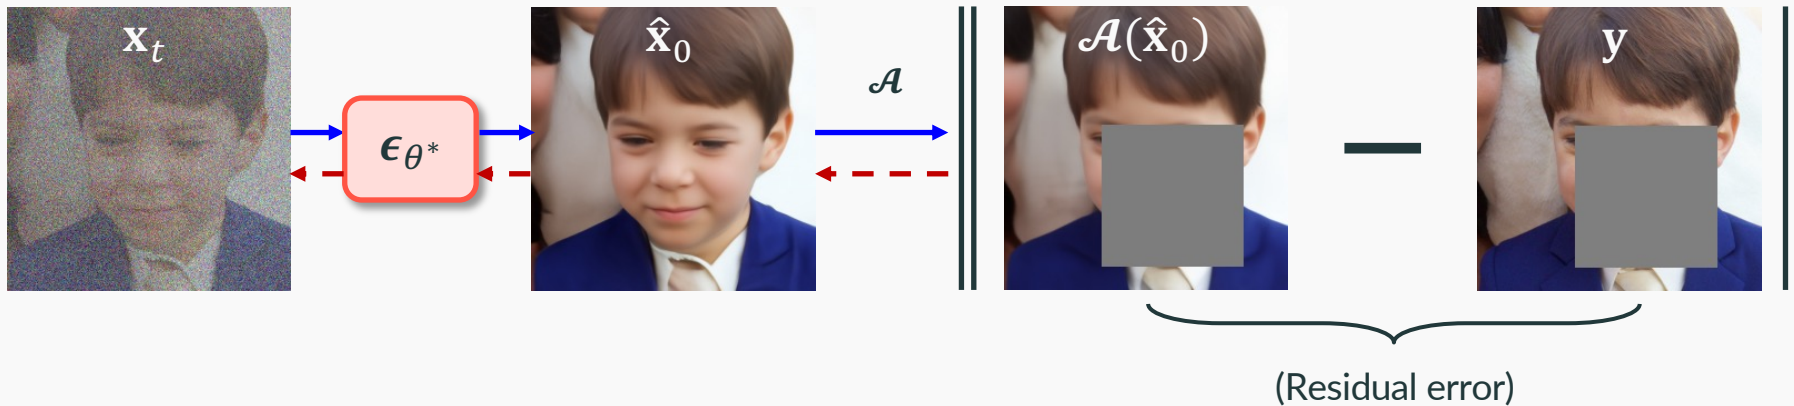
$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0) = -\rho \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$$



Forward pass

Diffusion Posterior Sampling

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \hat{\mathbf{x}}_0) = -\rho \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$$




Forward pass



Backward propagation


Manifold Constrained Gradient (MCG)

Chung et al, NeurIPS, 2022; Arxiv 2023



$$\frac{\partial}{\partial x_i} \|\mathbf{W}(y - \mathbf{H}\hat{x}_0)\|_2^2 = -2\mathbf{J}_{Q_i}^T \mathbf{H}^T \mathbf{W}^T \mathbf{W} (y - \mathbf{H}\hat{x}_0) \in T_{\hat{x}_0} \mathcal{M}$$

Geometry of DPS



Some Examples






Figure 6: Results on solving nonlinear inverse problems with Gaussian noise ($\sigma = 0.05$).



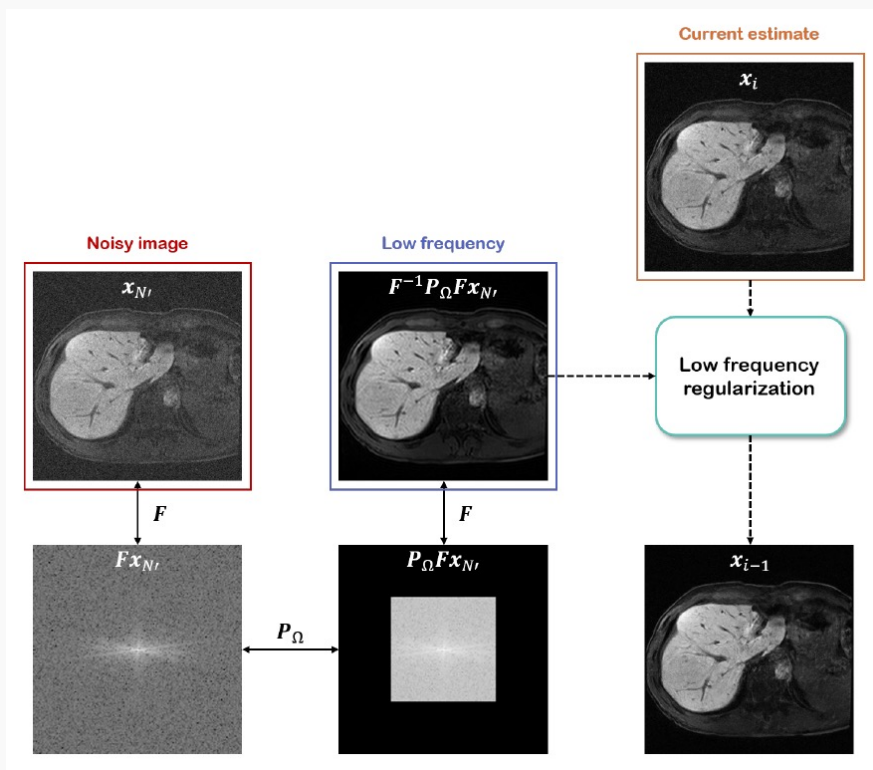
Reverse Diffusion Denoising + Super-resolution



Hijack the reverse diffusion!

Reverse Diffusion Denoising + Super-resolution

1. Low-frequency regularization

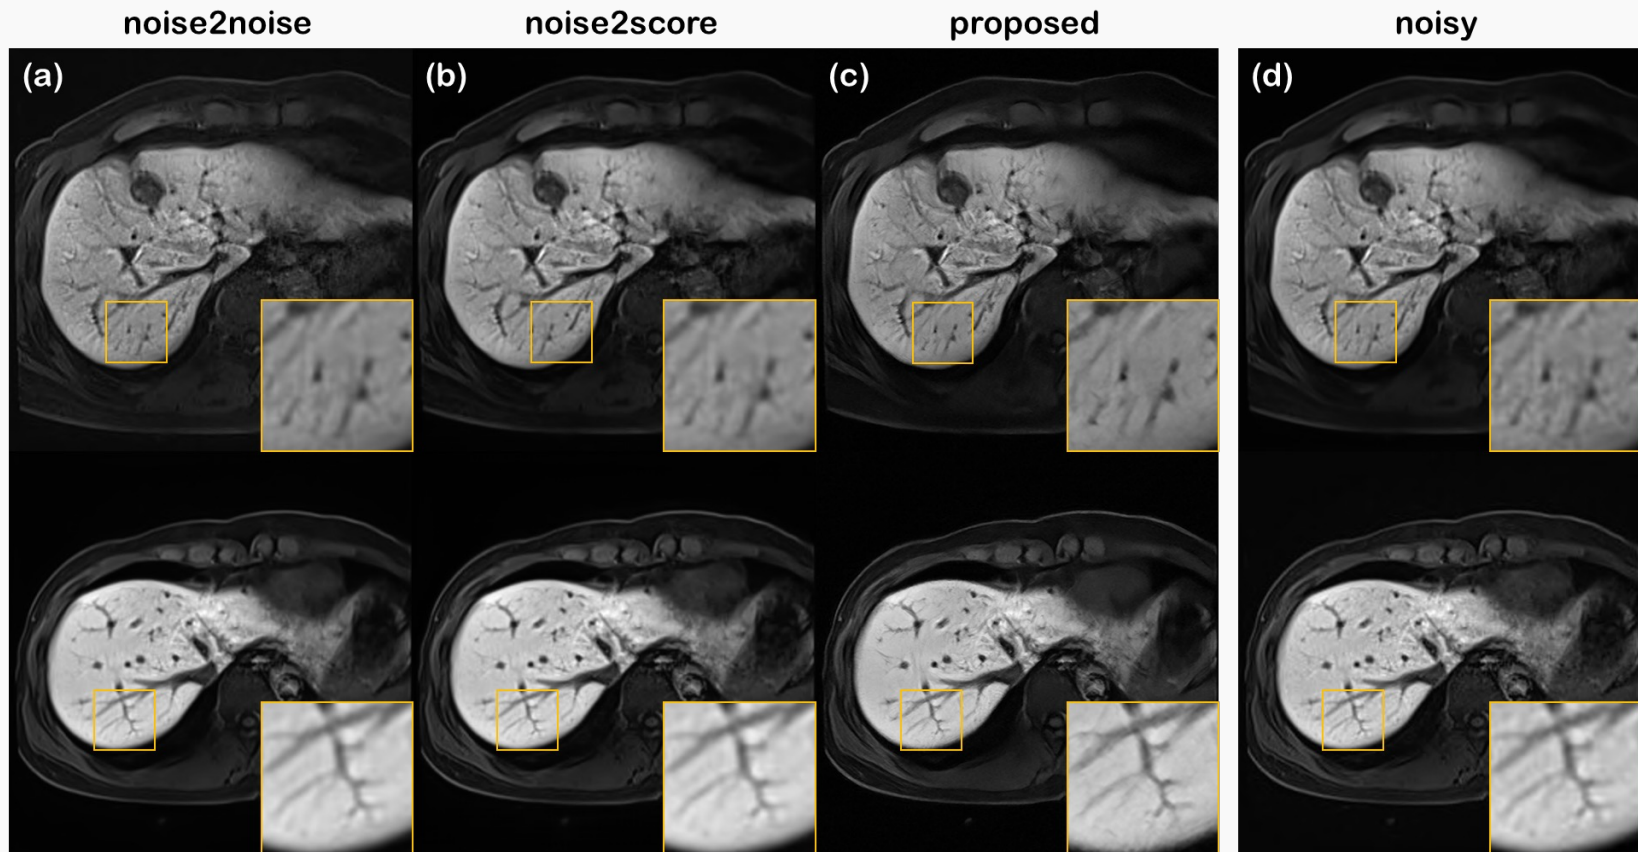


2. Reconstruction Guidance (MCG)


$$\hat{x}_0 = x_{i+1} + \sigma_{i+1}^2 s_\theta(x_{i+1}, \sigma_{i+1})$$
$$x_i'' = x_i' - \gamma \nabla_{x_{i+1}} \|x_i' - \hat{x}_0\|_2^2,$$

MCG

Reverse Diffusion Denoising + Super-resolution



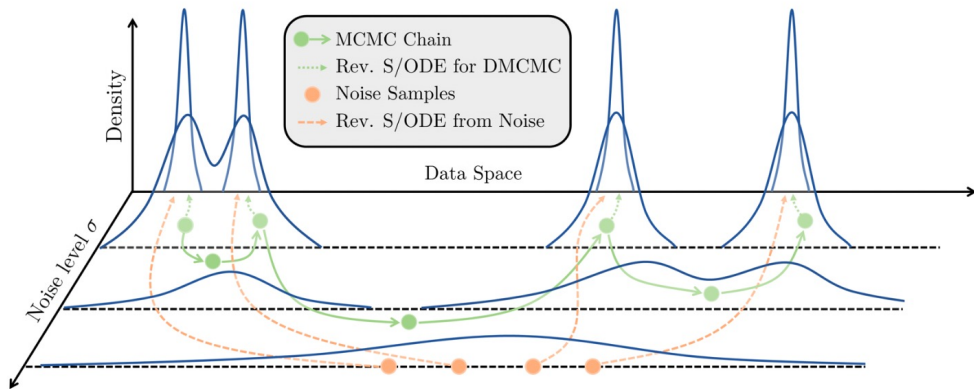
Uncertainty Quantification, User Controllability



How to accelerate?

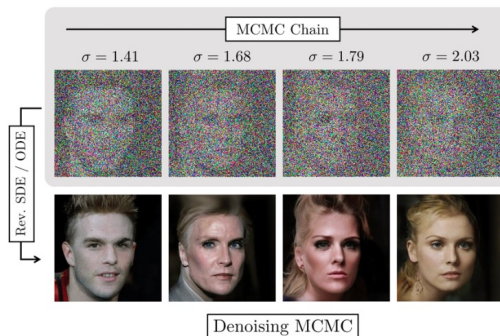
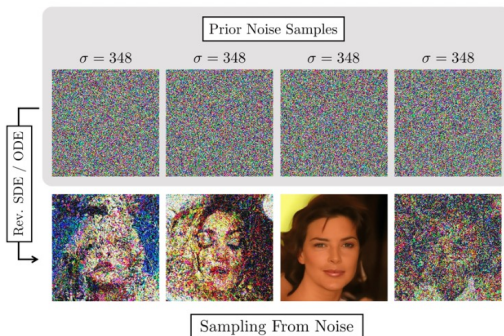
Denoising Markov-Chain Monte Carlo (DMCMC)

Kim et al, ICML 2023




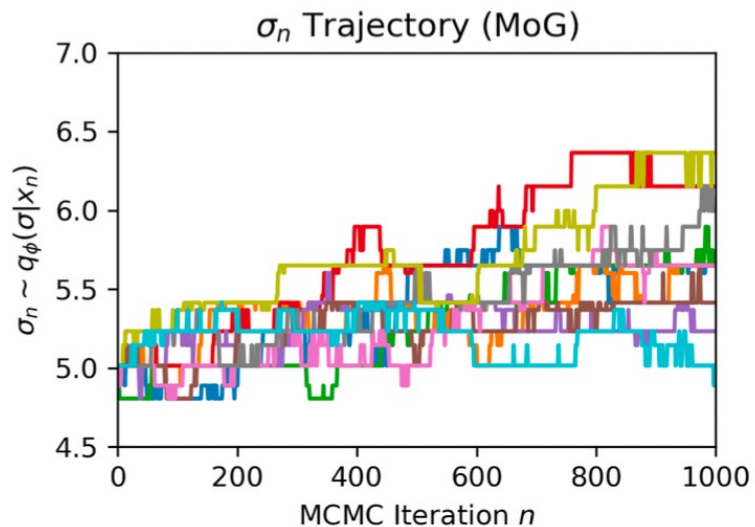
MCMC on $\mathcal{X} \times \mathcal{S}$
Langevin Gibbs Sampling

- ① $\mathbf{x}_{n+1} = \mathbf{x}_n + (\eta/2) \cdot s_\theta(\mathbf{x}_n, \sigma_n) + \sqrt{\eta} \cdot \epsilon$
- ② $\sigma_{n+1} \sim \hat{p}(\sigma | \mathbf{x}_{n+1})$
- ③ integrating the reverse-S/ODE



Denoising Markov-Chain Monte Carlo (DMCMC)

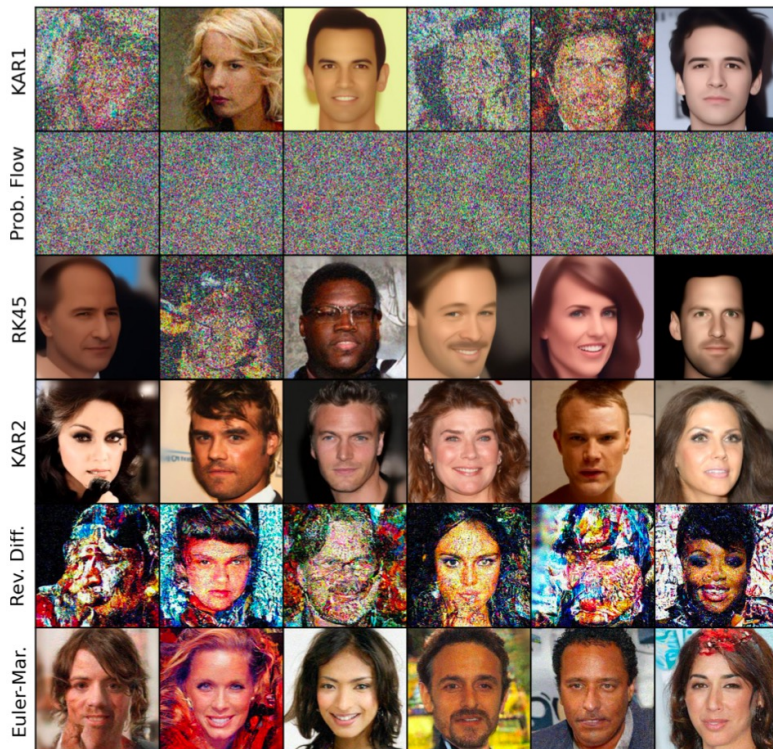
Kim et al, ICML 2023




Denoising Markov-Chain Monte Carlo (DMCMC)

Kim et al, ICML 2023

Without DLG ($NFE \approx 400$)



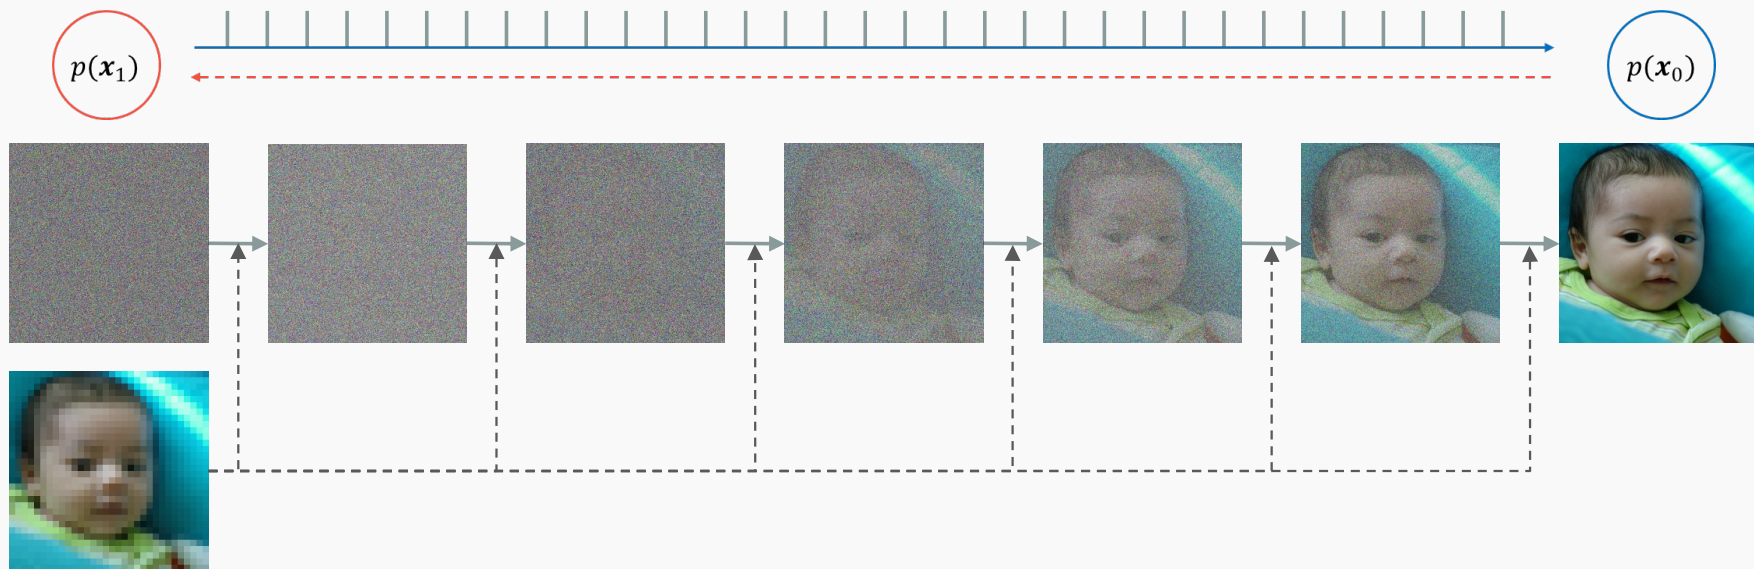
With DLG ($NFE < 150$)



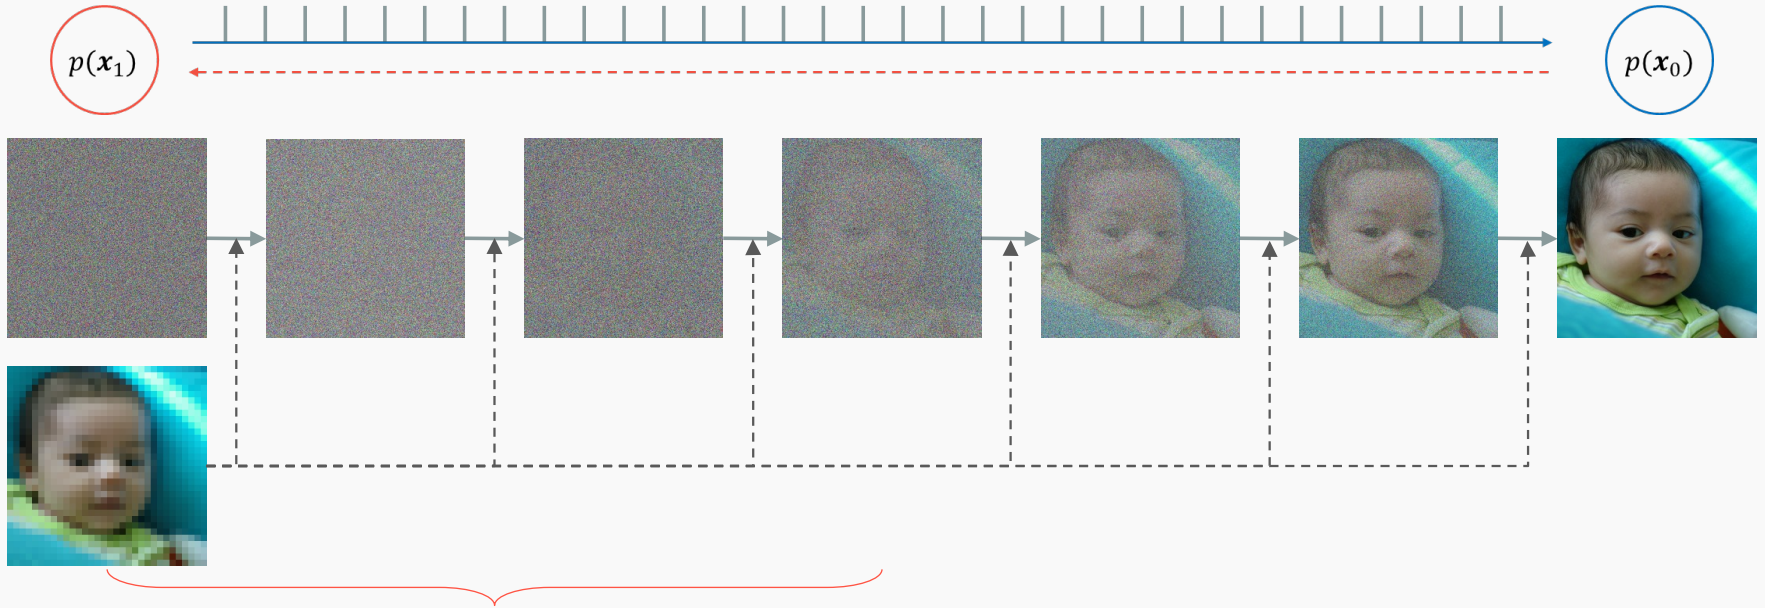
ACCELERATION OF CONDITIONAL DIFFUSION

Chung et al, CVPR, 2022

Intuition: Why use the whole process?



Intuition: Why use the whole process?

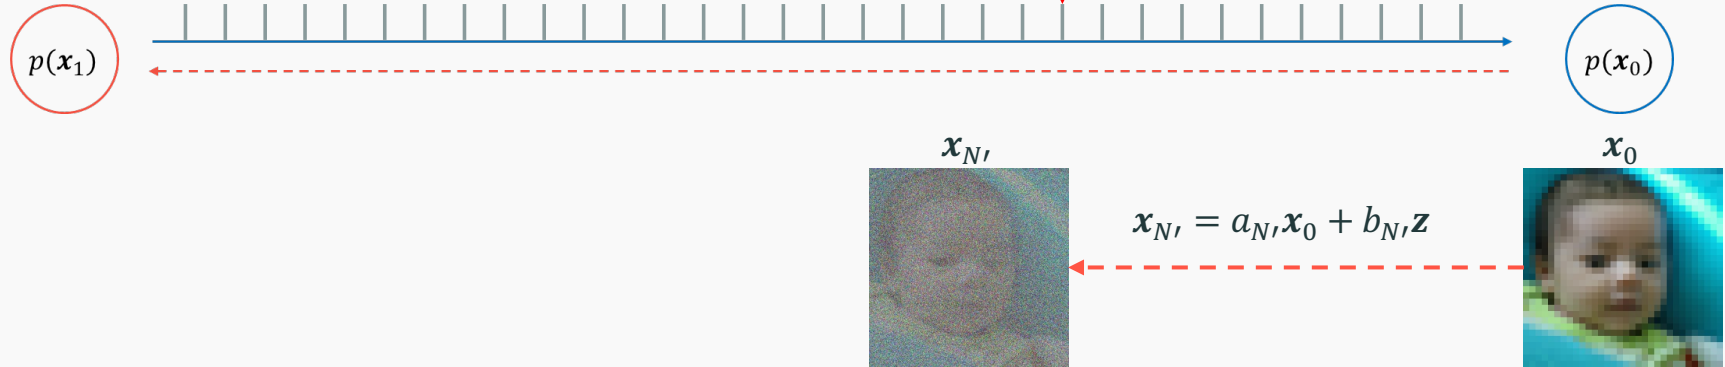


Is this part necessary?

Intuition: Why use the whole process?

- $t_0 = 0.3$
- $N' = t_0 N = 300$

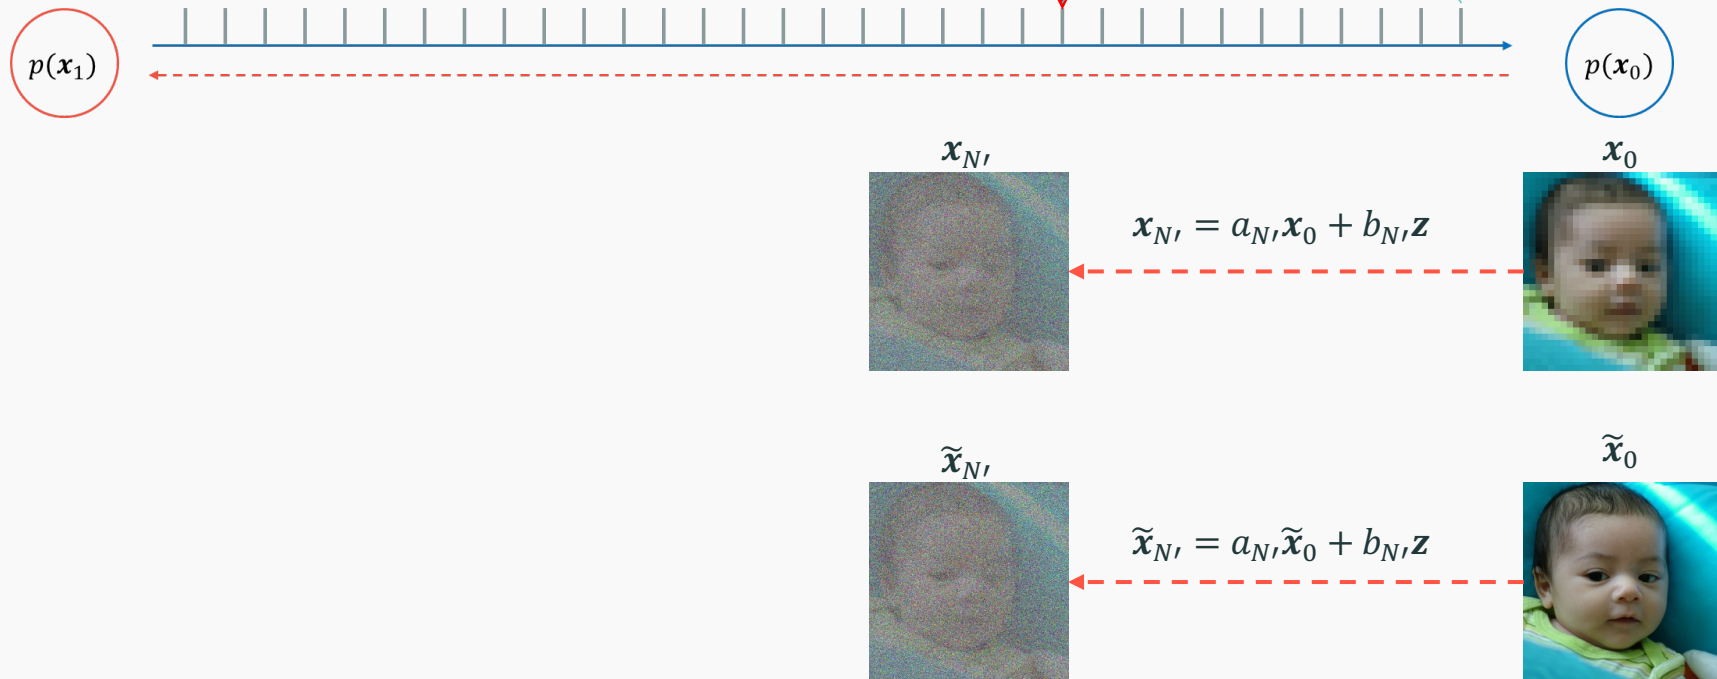
use this much



Intuition: Why use the whole process?

- $t_0 = 0.3$
- $N' = t_0 N = 300$

use this much



CCDF: The Algorithm




Algorithm 1 Accelerated Super-resolution / inpainting (VP, Markov)

Require: $x_0, \hat{x}_0, N', \{\alpha_i\}_{i=1}^{N'}, \{\sigma_i\}_{i=1}^{N'}, s_\theta$

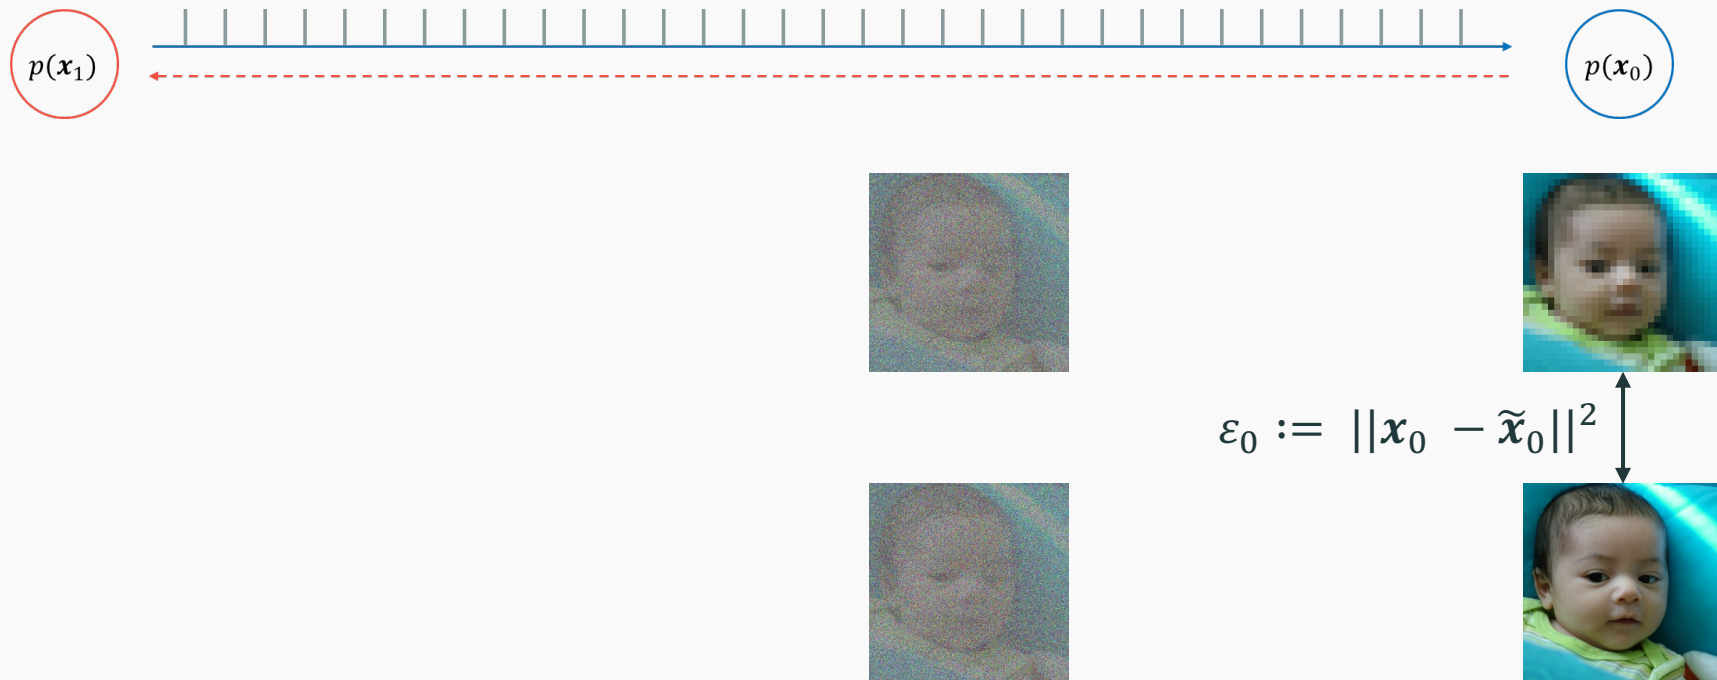
- 1: $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 2: $x_{N'} \leftarrow \sqrt{\bar{\alpha}_{N'}}x_0 + \sqrt{1 - \bar{\alpha}_{N'}}z$ \triangleright Forward diffusion
- 3: **for** $i = N'$ to 1 **do** \triangleright Reverse diffusion
- 4: $x'_{i-1} \leftarrow \frac{1}{\sqrt{\alpha_i}}(x_i + (1 - \alpha_i)s_\theta(x_i, i))$
- 5: $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 6: $x_{i-1} \leftarrow x'_{i-1} + \sigma_i z$ \triangleright Unconditional update
- 7: $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 8: $\hat{x}_i \leftarrow \sqrt{\bar{\alpha}_i}\hat{x}_0 + \sqrt{1 - \bar{\alpha}_i}z$
- 9: $x_{i-1} = (\mathbf{I} - \mathbf{P})x_{i-1} + \hat{x}_i$ \triangleright Measurement consistency

10: **end for**

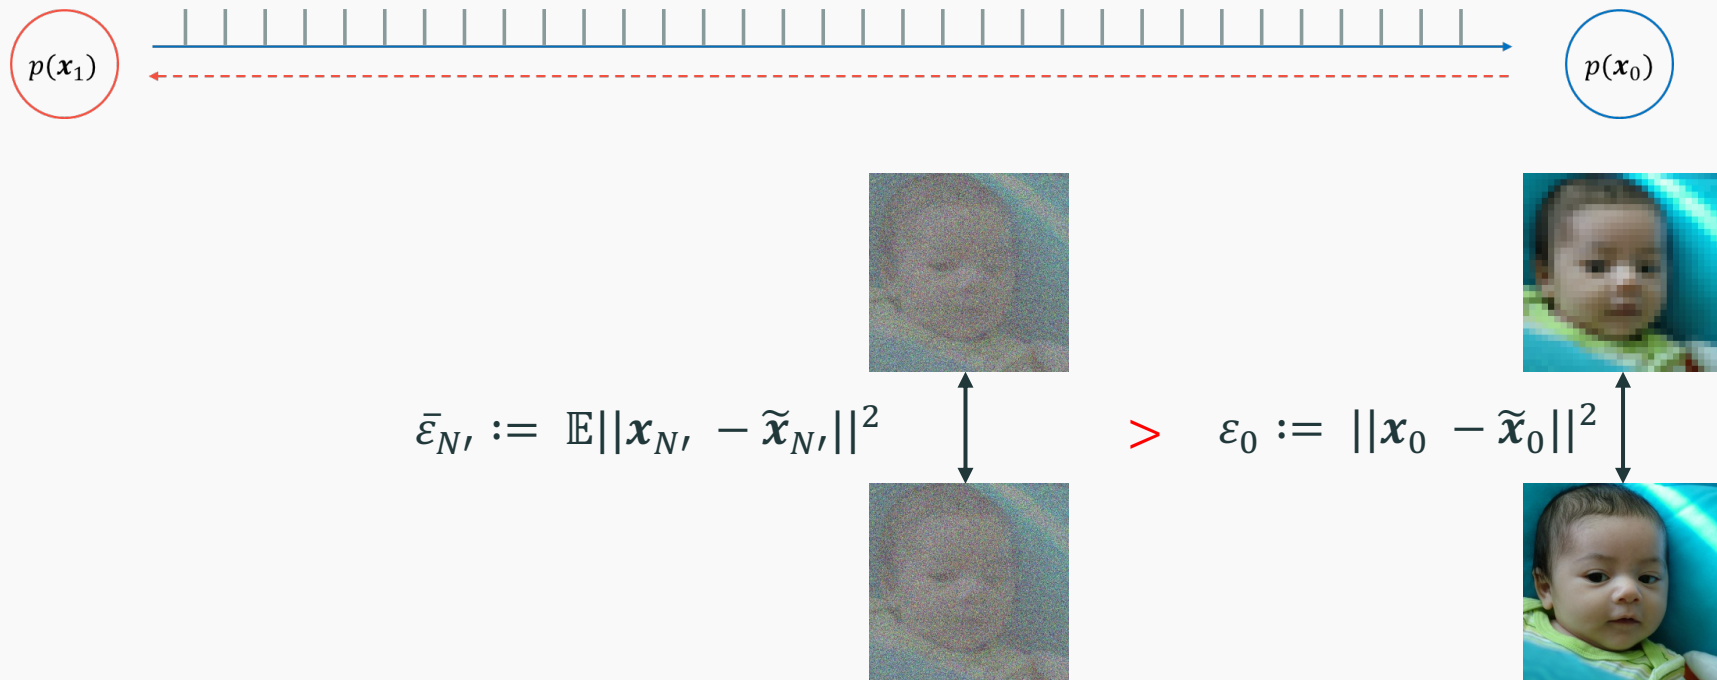
11: **return** x_0





Theoretical Challenges





Theoretical Challenges



Theoretical Challenges



$$\bar{\epsilon}_{0,r} := \mathbb{E} \|\mathbf{x}_{0,r} - \tilde{\mathbf{x}}_{0,r}\|^2 < \epsilon_0 ?$$



Theoretical Findings

Chung et al, CVPR, 2022

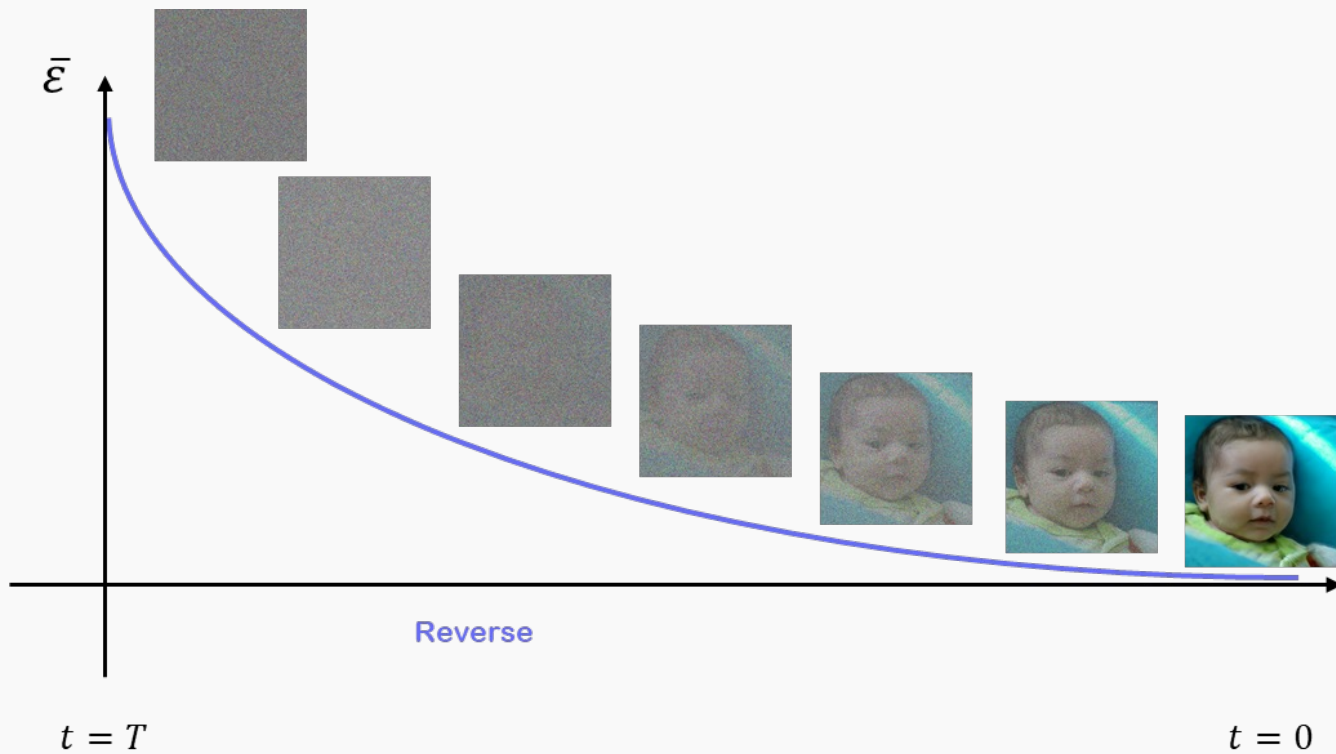
Theorem 2. (shortcut path)

- For any $0 < \mu \leq 1$, there exists a **minimum** N' s.t.

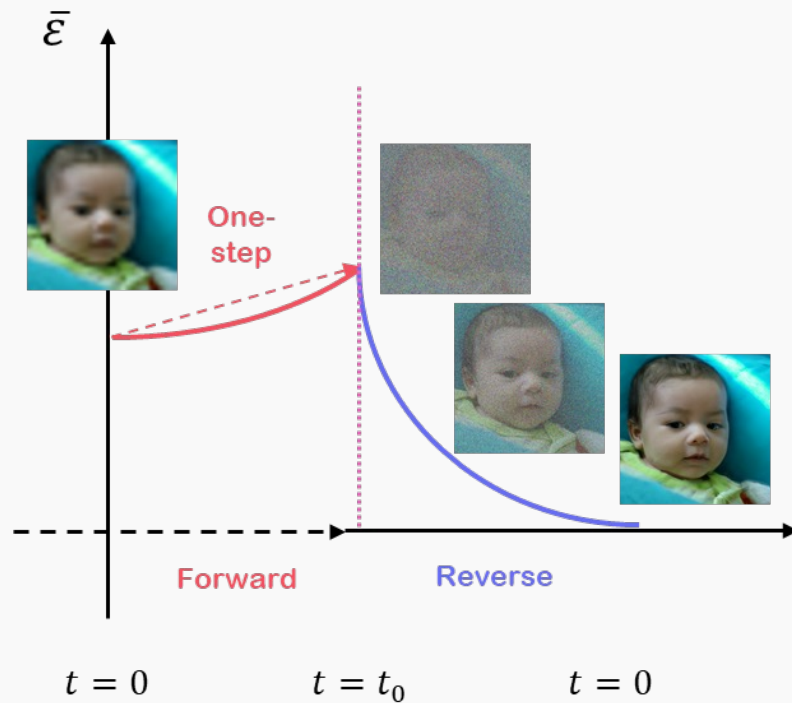
$$\bar{\varepsilon}_{0,r} \leq \mu \varepsilon_0$$

- **Optimal** N' **decreases** as ε_0 **gets smaller**

Come Closer Diffuse Faster (CCDF)

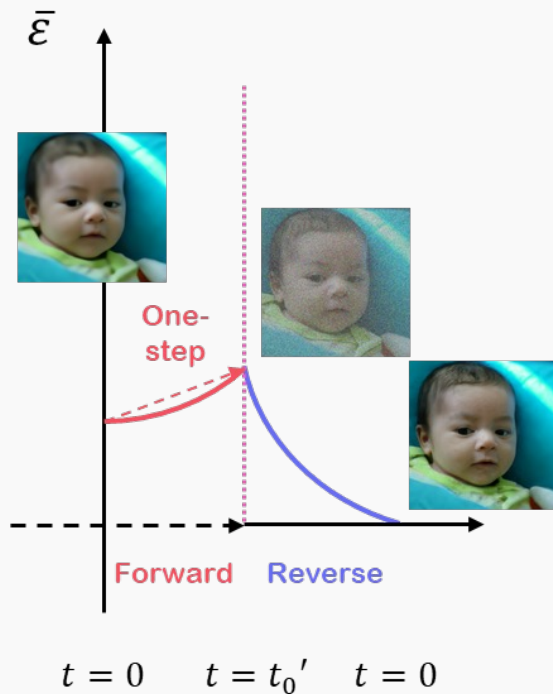


Come Closer Diffuse Faster (CCDF)



Come Closer Diffuse Faster (CCDF)

Feed-forward
network correction



Reverse Diffusion is Contracting!

Stochastic Contraction

Theorem 1.

$$\bar{\varepsilon}_{0,r} \leq \frac{2C\tau}{1 - \lambda^2} + \lambda^{2N'} \bar{\varepsilon}_{N'}$$

$$\tau = \frac{\text{Tr}(\mathbf{A}^T \mathbf{A})}{n}$$

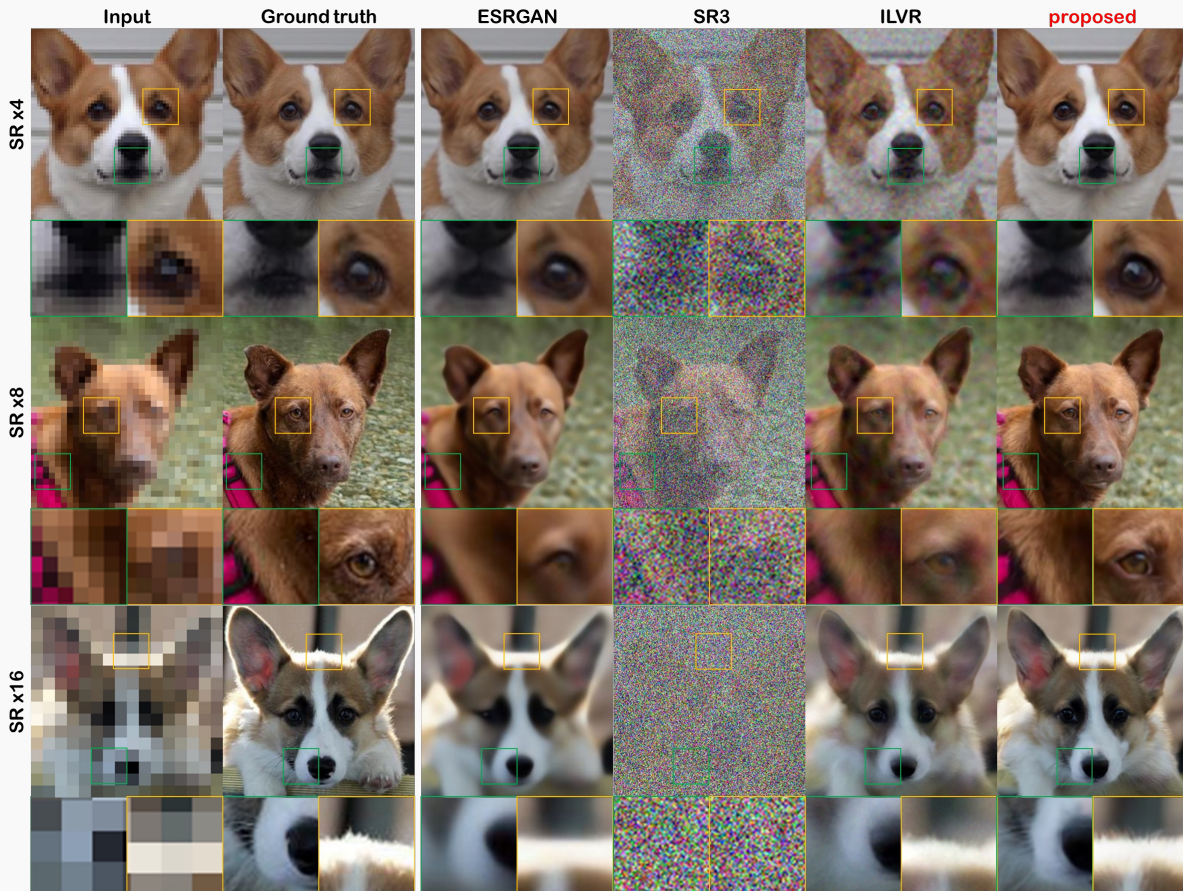
Error decreases **exponentially**
with reverse diffusion!

Reverse diffusion is a
contracting mapping with λ

$$\lambda = \begin{cases} \max_{i \in [N']} \sqrt{\alpha_i} \left(\frac{1 - \bar{\alpha}_{i-1}}{1 - \bar{\alpha}_i} \right) & (DDPM) \\ \max_{i \in [N']} \frac{\sigma_{i-1}^2 - \sigma_0^2}{\sigma_i^2 - \sigma_0^2} & (SMLD) \\ \max_{i \in [N']} \frac{\sigma_{i-1}}{\sigma_i} & (DDIM) \end{cases}$$

$$C = \begin{cases} n(1 - \alpha_N) & (DDPM) \\ n \max_{i \in [N']} \sigma_i^2 - \sigma_{i-1}^2 & (SMLD) \\ 0 & (DDIM) \end{cases}$$

Experimental Results: SR



20 step diffusion

- ILVR, SR3

$$N = 20, \quad t_0 = 1.0$$


- proposed

$$N = 100, \quad t_0 = 0.2$$

t_0	0.05	0.1	0.2	0.5	0.75	1.0 [5]
SR $\times 4$	63.90	60.90	<u>60.91</u>	64.04	64.14	63.31
SR $\times 8$	85.21	78.13	75.76	79.34	79.67	<u>77.34</u>
SR $\times 16$	116.37	101.79	92.59	88.09	92.12	<u>88.49</u>

Table 1. FID(\downarrow) scores on FFHQ test set for SR task with $N = 1000$, and varying t_0 values. $t_0 = 1.0$ is the baseline method without any acceleration used in [5]. Numbers in boldface, and underline indicate the best, and the second best scores.

Experimental Results: Inpainting



20 step diffusion

- **Score-SDE**

$$N = 20, \quad t_0 = 1.0$$

- **proposed**

$$N = 100, \quad t_0 = 0.2$$

What if We Do Not Know the Forward Model?

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \boldsymbol{\eta}$$

unknown

What if We Do Not Know the Forward Model?


$$\mathbf{y} = \mathcal{A}_{\phi}(\mathbf{x}) + \boldsymbol{\eta}$$

unknown

What if We Do Not Know the Forward Model?


$$\mathbf{y} = \mathcal{A}_{\phi}(\mathbf{x}) + \eta$$

unknown



\mathcal{A}_{ϕ}

Blind deconvolution
(deblurring)



Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) = \mathcal{N}(\mathbf{y}|\mathbf{k}_0 * \mathbf{x}_0, \sigma^2 I)$$

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) = \mathcal{N}(\mathbf{y}|\mathbf{k}_0 * \mathbf{x}_0, \sigma^2 I)$$

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

Image
prior

Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) = \mathcal{N}(\mathbf{y}|\mathbf{k}_0 * \mathbf{x}_0, \sigma^2 I)$$

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

kernel prior

Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) = \mathcal{N}(\mathbf{y}|\mathbf{k}_0 * \mathbf{x}_0, \sigma^2 \mathbf{I})$$

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

Choices? Diffusion prior!

Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) = \mathcal{N}(\mathbf{y}|\mathbf{k}_0 * \mathbf{x}_0, \sigma^2 I)$$

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

Trained independently
with DSM

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}) \simeq s_{\theta}^I(\mathbf{x})$$

$$\nabla_{\mathbf{k}} \log p(\mathbf{k}) \simeq s_{\theta}^k(\mathbf{k})$$

Diffusion Model for **Blind** Deconvolution

$$p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0) \propto p(\mathbf{y}|\mathbf{x}_0, \mathbf{k}_0)p(\mathbf{x}_0)p(\mathbf{k}_0)$$

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)[\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\hat{\mathbf{x}}_0, \hat{\mathbf{k}}_0)] \right] dt + \sqrt{\beta(t)}d\bar{\mathbf{w}}$$

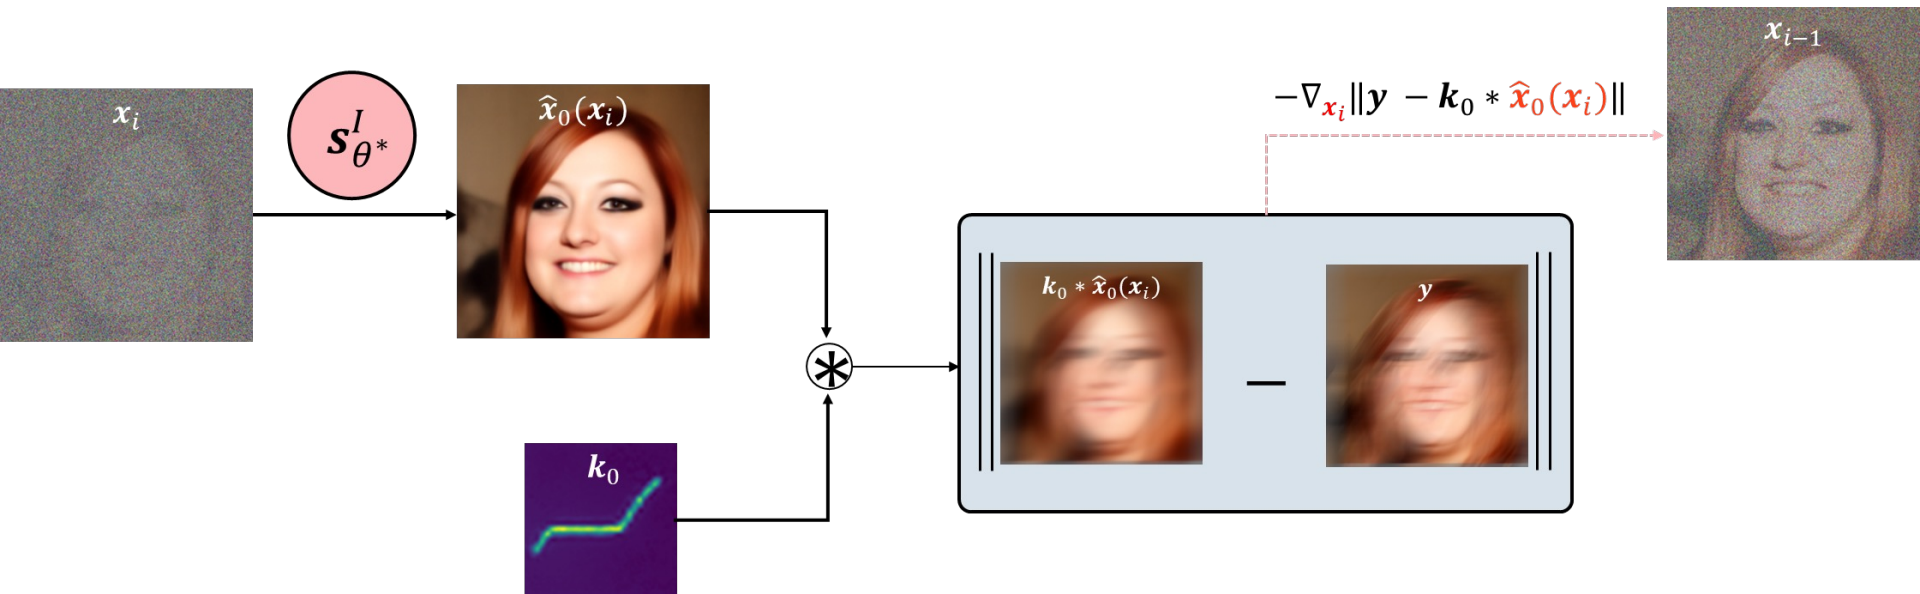
$$d\mathbf{k} = \left[-\frac{\beta(t)}{2}\mathbf{k} - \beta(t)[\nabla_{\mathbf{k}_t} \log p(\mathbf{k}_t) + \nabla_{\mathbf{k}_t} \log p(\mathbf{y}|\hat{\mathbf{x}}_0, \hat{\mathbf{k}}_0)] \right] dt + \sqrt{\beta(t)}d\bar{\mathbf{w}}$$

Theorem. Under similar conditions as in DPS,

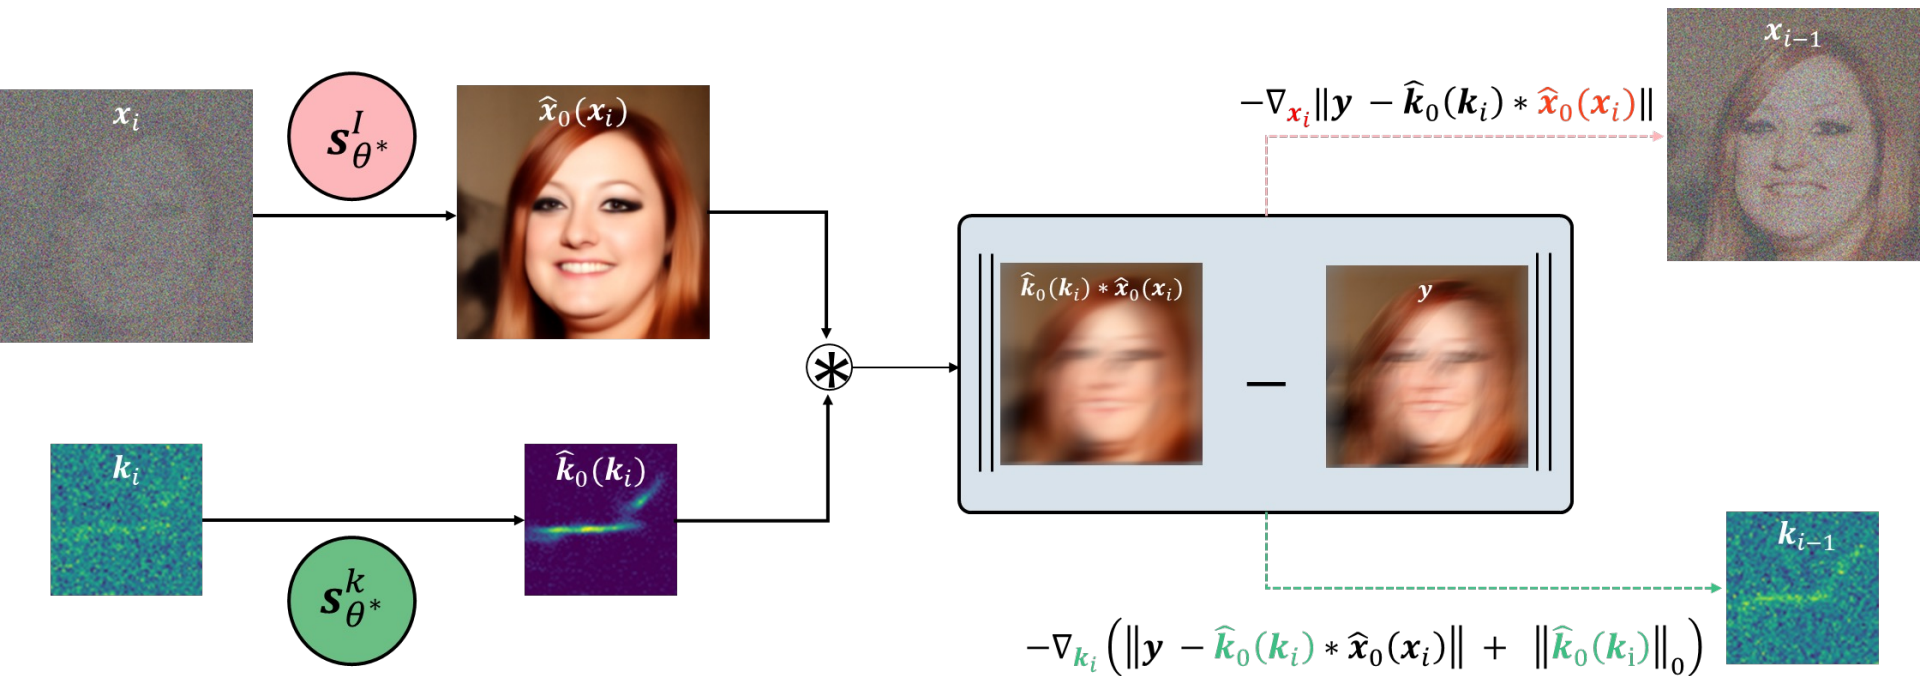
$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t, \mathbf{k}_t) \simeq \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\hat{\mathbf{x}}_0(\mathbf{x}_t), \hat{\mathbf{k}}_0(\mathbf{k}_t))$$

$$\nabla_{\mathbf{k}_t} \log p(\mathbf{y}|\mathbf{x}_t, \mathbf{k}_t) \simeq \nabla_{\mathbf{k}_t} \log p(\mathbf{y}|\hat{\mathbf{x}}_0(\mathbf{x}_t), \hat{\mathbf{k}}_0(\mathbf{k}_t))$$

Single Diffusion Model for Non-blind Deconvolution




Parallel Diffusion Model for **Blind** Deconvolution




Applicability to Diverse **Blind** Inverse Problems

1. Blind deconvolution



2. Imaging through turbulence



Results: Blind Deblurring

Measurement



Ours




Ground truth



Results: Imaging through Turbulence

Measurement



Ours

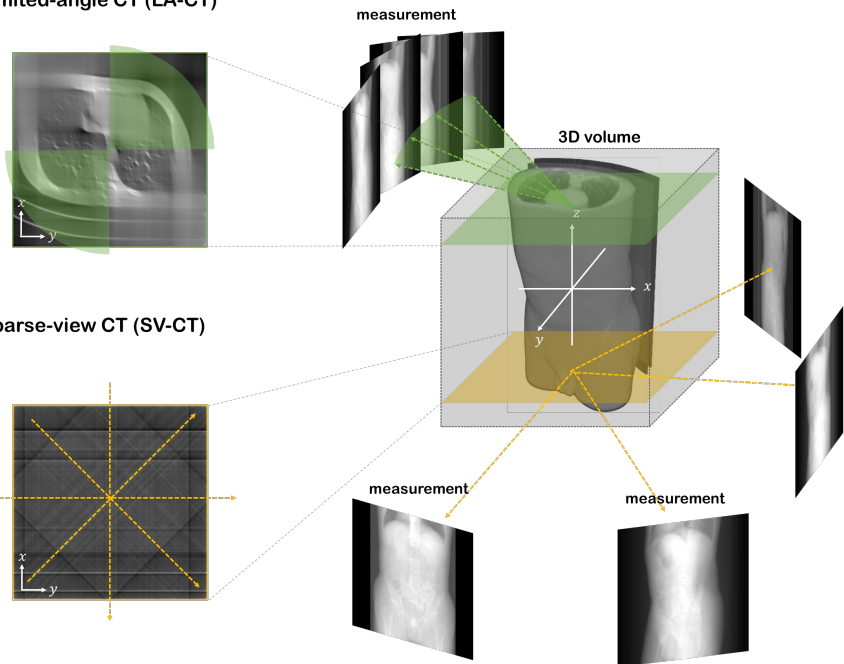


Ground truth

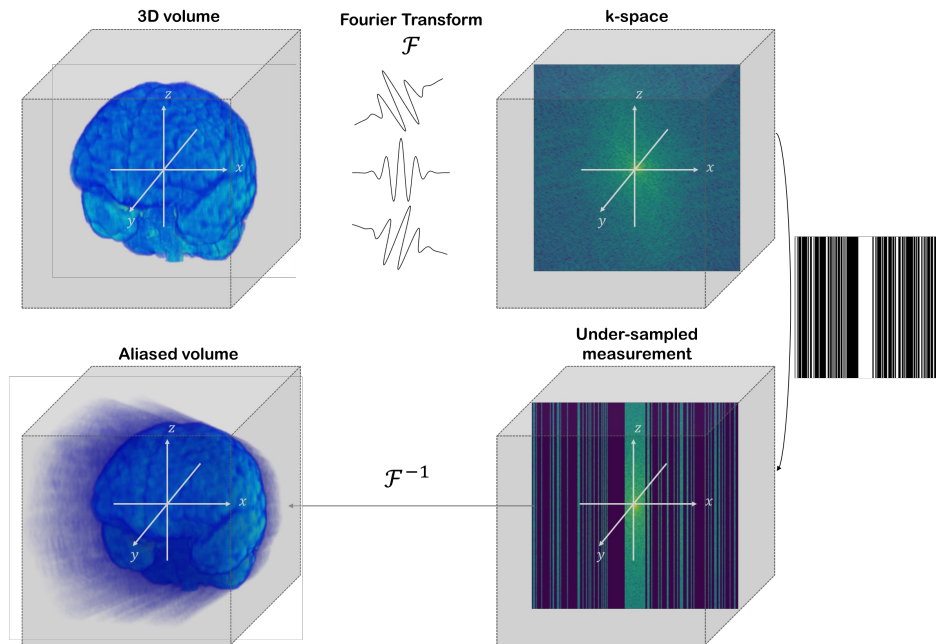


Scaling Up to 3D Inverse Problems

(a) Limited-angle CT (LA-CT)



(c) Compressed-sensing MRI



Naïve application induce **incoherent** reconstructions

Scaling Up to 3D Inverse Problems

- 3D representations are **memory heavy**
 - **Voxels**: Hard to deal with $> 128^3$ data
 - **Point clouds**: Sparse representation, but not suitable for medical imaging inverse problems
- 3D voxel diffusion?
 - The whole diffusion process **stays in the data dimension**
 - **Computationally too heavy**

Augmenting 2D Diffusion Prior with Model-based Prior

Model-based prior (TV)

$$TV(\mathbf{x}) := \left\| \left[D_x \mathbf{x}, D_y \mathbf{x}, D_z \mathbf{x} \right] \right\|_1$$

Diffusion prior

$$\nabla_{\mathbf{x}} \log p(\mathbf{x})$$

Augmenting 2D Diffusion Prior with Model-based Prior

Model-based prior (TV: z)

$$TV_z(\mathbf{x}) := \|D_z \mathbf{x}\|_1$$

Diffusion prior (\mathbf{x}_y)

$$\nabla_{\mathbf{x}_{\mathbf{y}}} \log p(\mathbf{x}_{\mathbf{y}})$$

Augmenting 2D Diffusion Prior with Model-based Prior

1. Denoising with score function (**parallel**)

$$\mathbf{x}'_{i-1} \leftarrow (\sigma_i^2 - \sigma_{i-1}^2) \mathbf{s}_{\theta^*}(\mathbf{x}_t, t) + \sqrt{\sigma_i^2 - \sigma_{i-1}^2} \boldsymbol{\epsilon}$$


2. Data consistency + TV prior augmenting (**joint**)

$$\mathbf{x}_{i-1} \leftarrow \underset{\mathbf{x}'_{i-1}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}'_{i-1}\|_2^2 + \|\mathbf{D}_z \mathbf{x}'_{i-1}\|_1$$

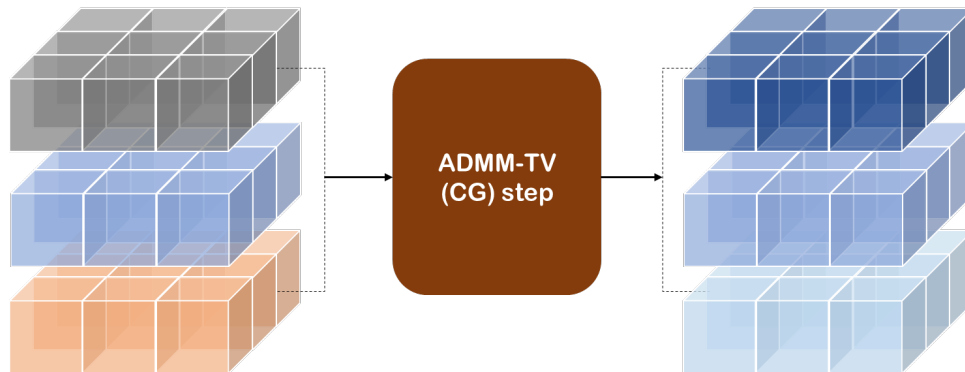
Effectively solved with, e.g. ADMM

Augmenting 2D Diffusion Prior with Model-based Prior



1. Score function denoising (parallel)



2. ADMM-TV (joint)



Results: 8-view 3D SV-CT



Coherent results across the **whole volume**



Results: General 3D Problems in Medical Imaging

Sparse-view tomography



In-distribution

Out-of-distribution



Measurement






Ours






Ground Truth



Limited-angle tomography


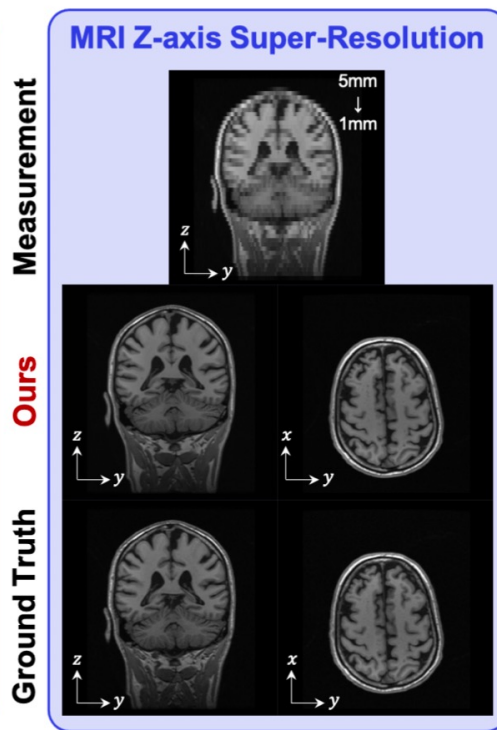
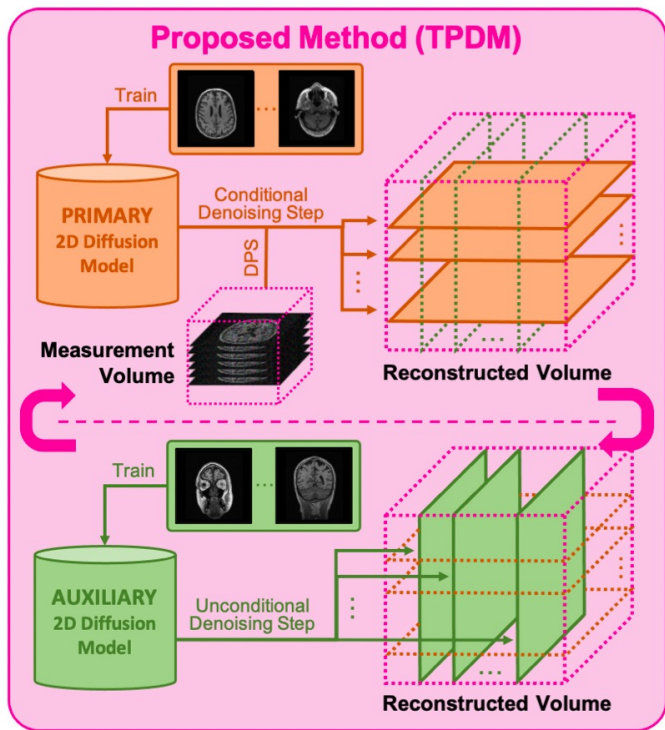



Compressed-sensing MRI




3D Diffusion using Perpendicular 2D Diffusion Model

Lee et al, arXiv:2303.08440 (2023)







Summary

- **Score-based approaches**: Exciting new path for solving inverse problems without labels
- **Universal solver** without knowledge about the problem a priori
- **Diffusion models**: Great **generalization** capacity
- Acceleration through **stochastic contraction** theory
- Understanding **diffusion geometry** helps optimizing algorithms



Questions?