Test-time Adaptation in Diffusion Models for Medical Image Reconstruction

Symposium on AI and Reconstruction for Biomedical Imaging (SyneRBI)

Hyungjin Chung, Ph.D.

Incoming Assistant Professor, Dept. of CSE, Korea University

Lead AI Research Scientist, EverEx

2026.03.09

Advances in Diffusion Models

Advances in Diffusion Models for Inverse Problems

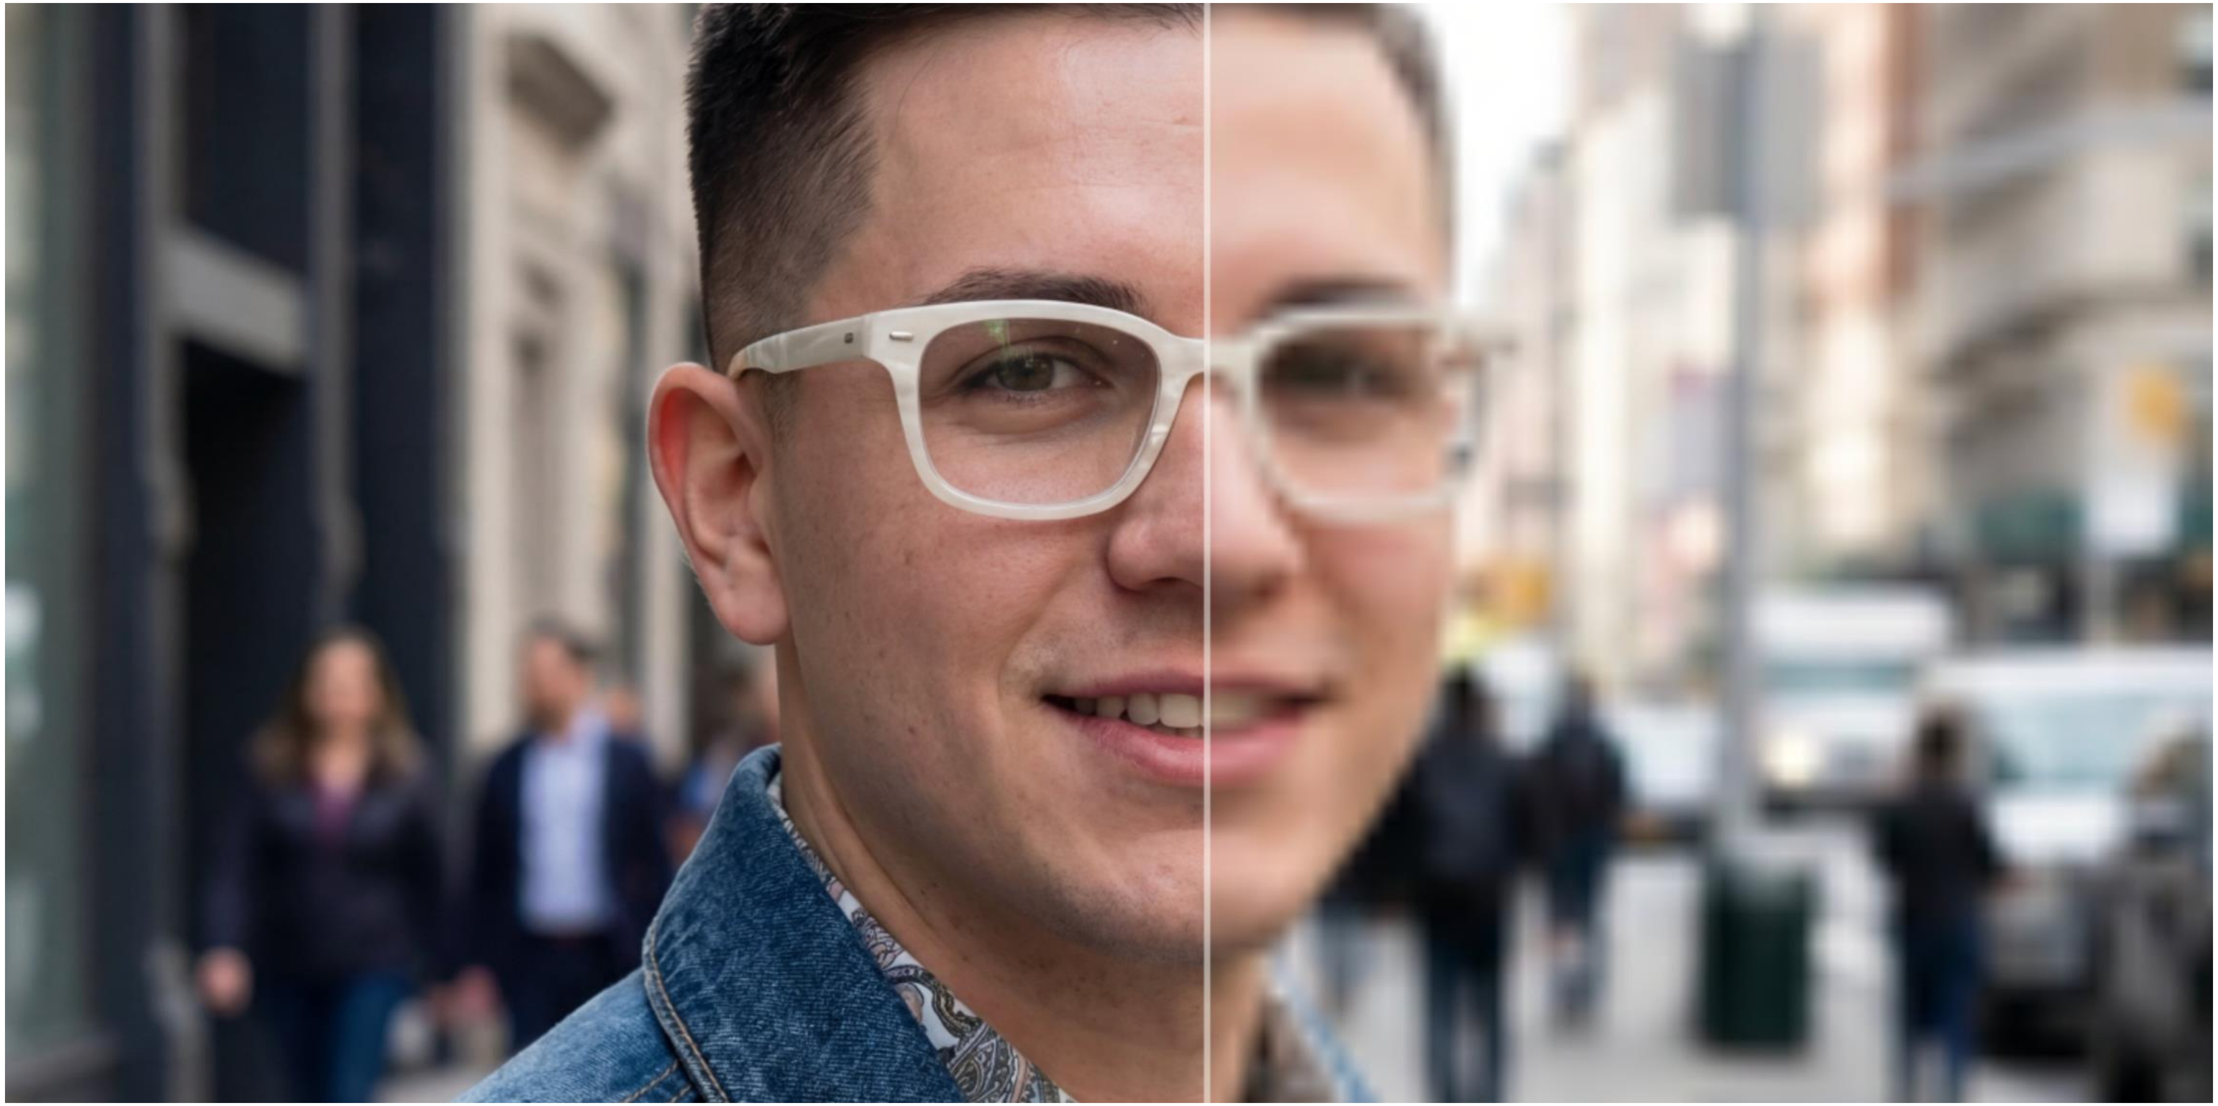

Generative Models are Pattern Matchers

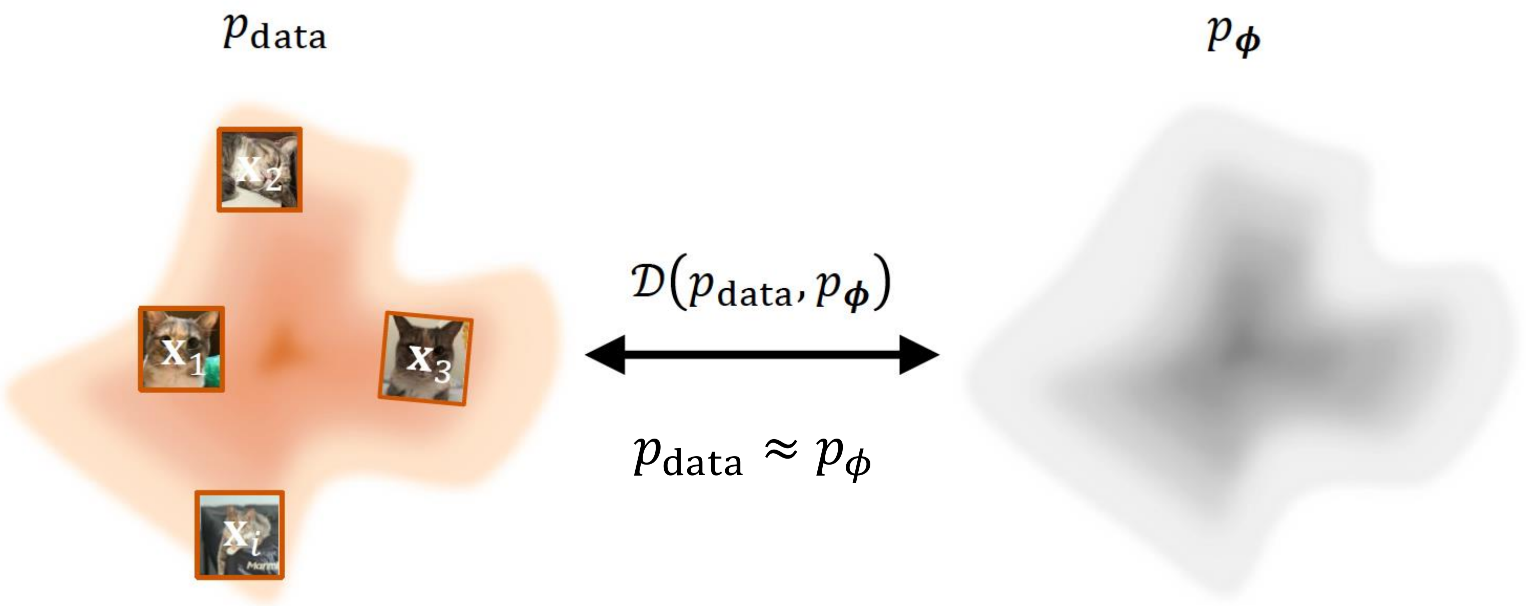

$$\phi^* \in \underset{\phi}{\operatorname{argmin}} \mathcal{D}(p_{\text{data}}, p_\phi)$$

Denoising Diffusion Models

Forward diffusion process

Gradually add noise to input, fixed

Data

Noise

Reverse denoising process

Generate data from noise by denoising, learned

[1] Sohl-Dickstein et al., Deep Unsupervised Learning using Nonequilibrium Thermodynamics, ICML 2015
[2] Ho et al., Denoising Diffusion Probabilistic Models, NeurIPS 2020
[3] Song et al., Score-Based Generative Modeling through Stochastic Differential Equations, ICLR 2021

Denoising Autoencoder Perspective

$$\|\mathbf{x}_0 - \underbrace{D_\theta(\mathbf{x}_t; t)}_{\text{“predict } \mathbf{x}_0 \text{ from } \mathbf{x}_t\text{”}}\|_2^2$$

[Small noise]

[Mid noise]

[High noise]

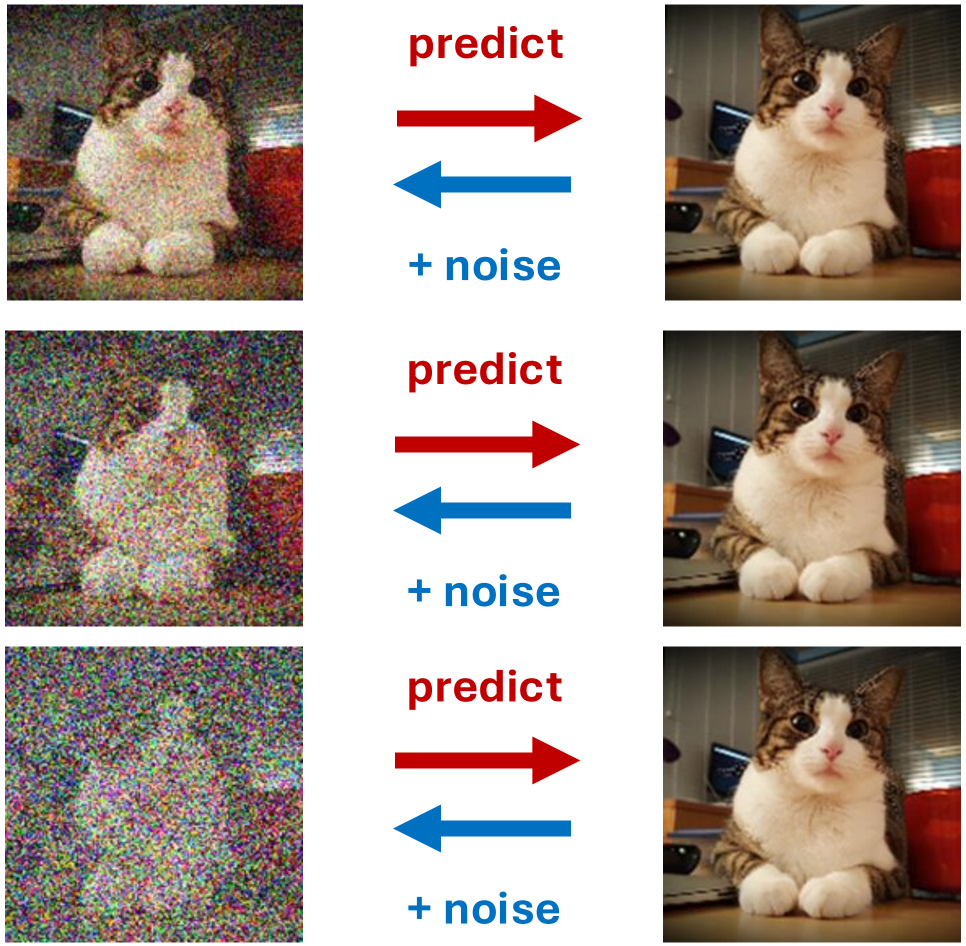

Denoising Score Matching and Autoencoders

$$\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t|\mathbf{x}_0) = -\frac{\mathbf{x}_t - \mathbf{x}_0}{\sigma_t^2}$$

(from definition of Gaussian distribution)

$$s_\theta(\mathbf{x}_t, t) := -\frac{\mathbf{x}_t - D_\theta(\mathbf{x}_t, t)}{\sigma_t^2}$$

(by design)

Then denoising score matching objective $\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)q(\mathbf{x}_0)}\left[\|\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t|\mathbf{x}_0) - s_\theta(\mathbf{x}_t, t)\|_2^2\right]$

is equivalent to the diffusion model objective! $\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)q(\mathbf{x}_0)}\left[\|\mathbf{x}_0 - D_\theta(\mathbf{x}_t, t)\|_2^2\right]$

Diffusion models are learning the scores of the smoothed $p(\mathbf{x})$

Visualization of Forward/Backward Process

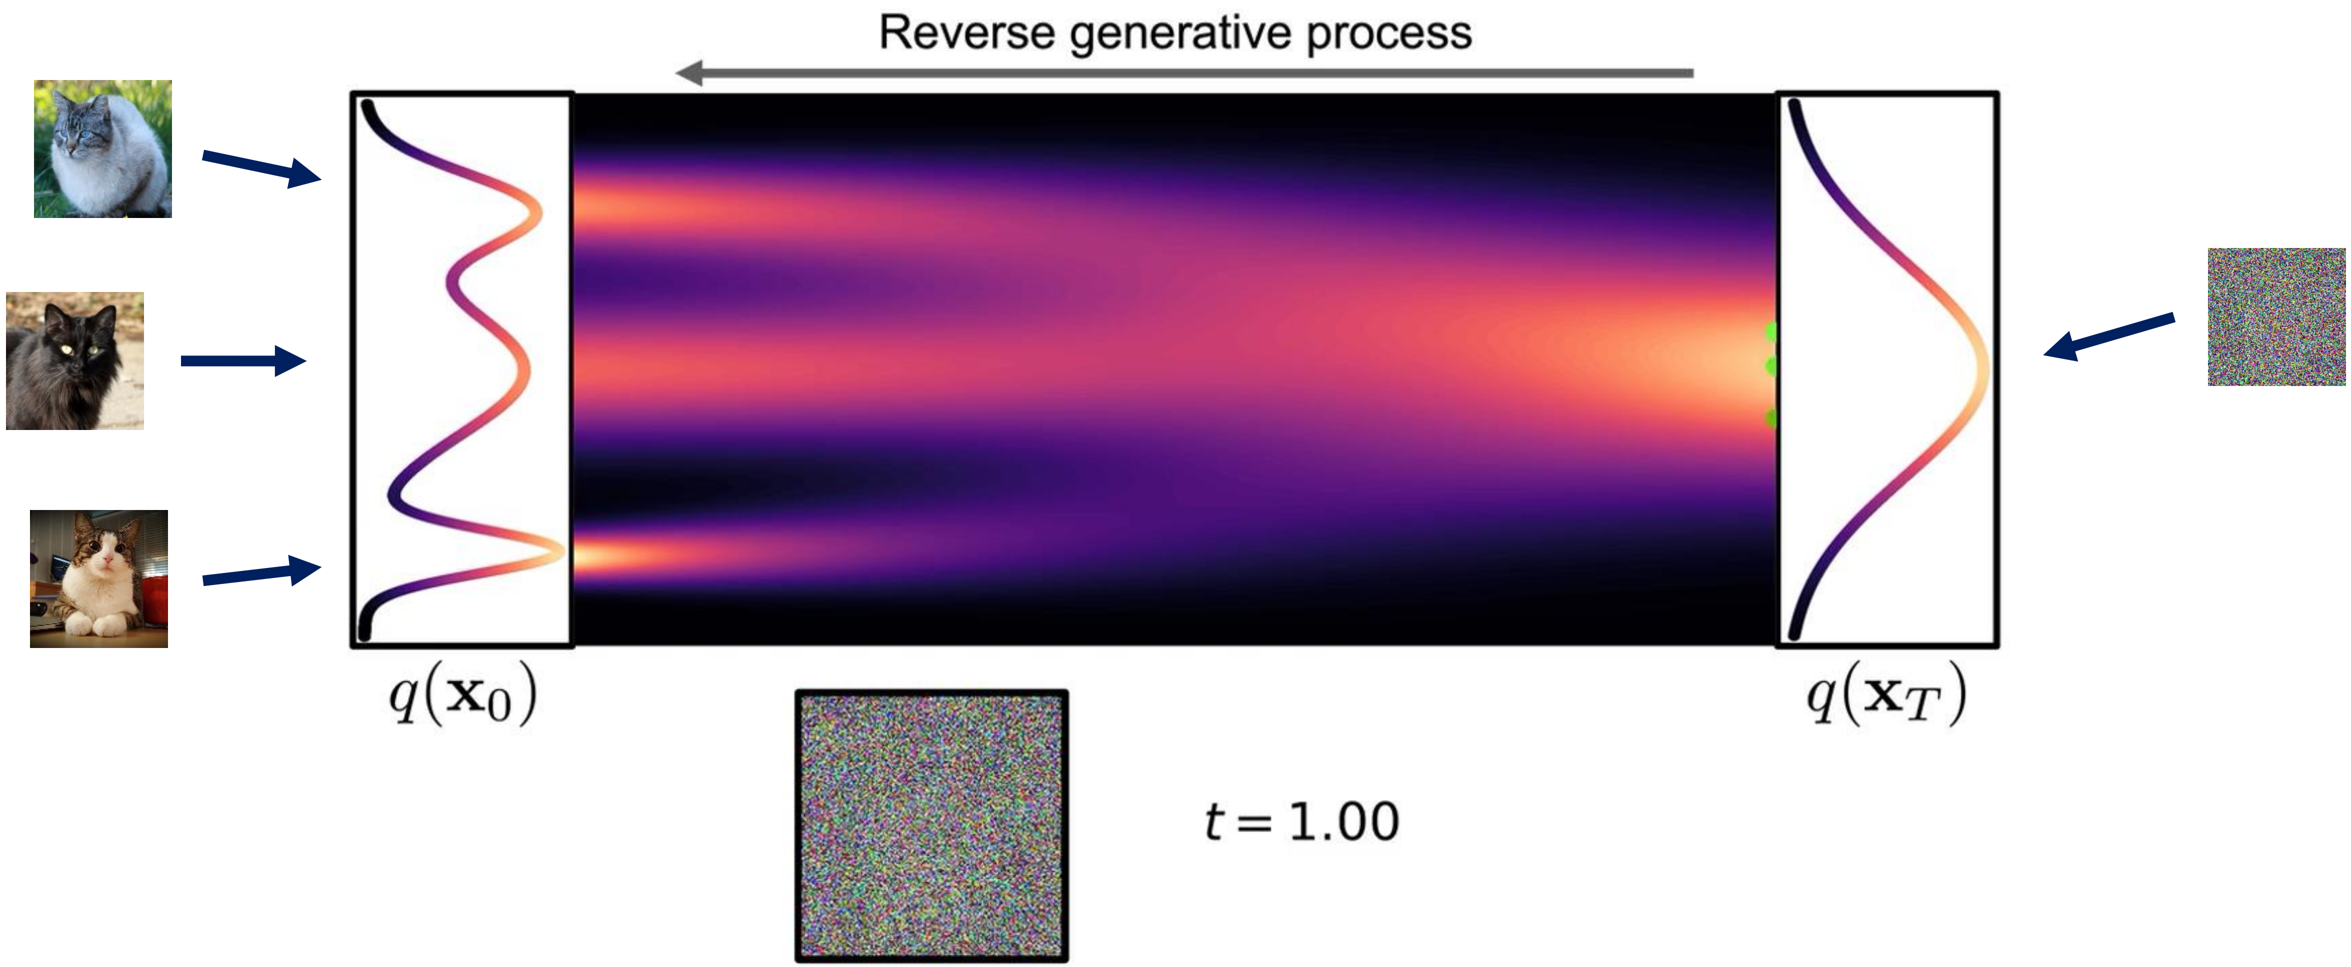

Inverse Problems in Biomedical Imaging

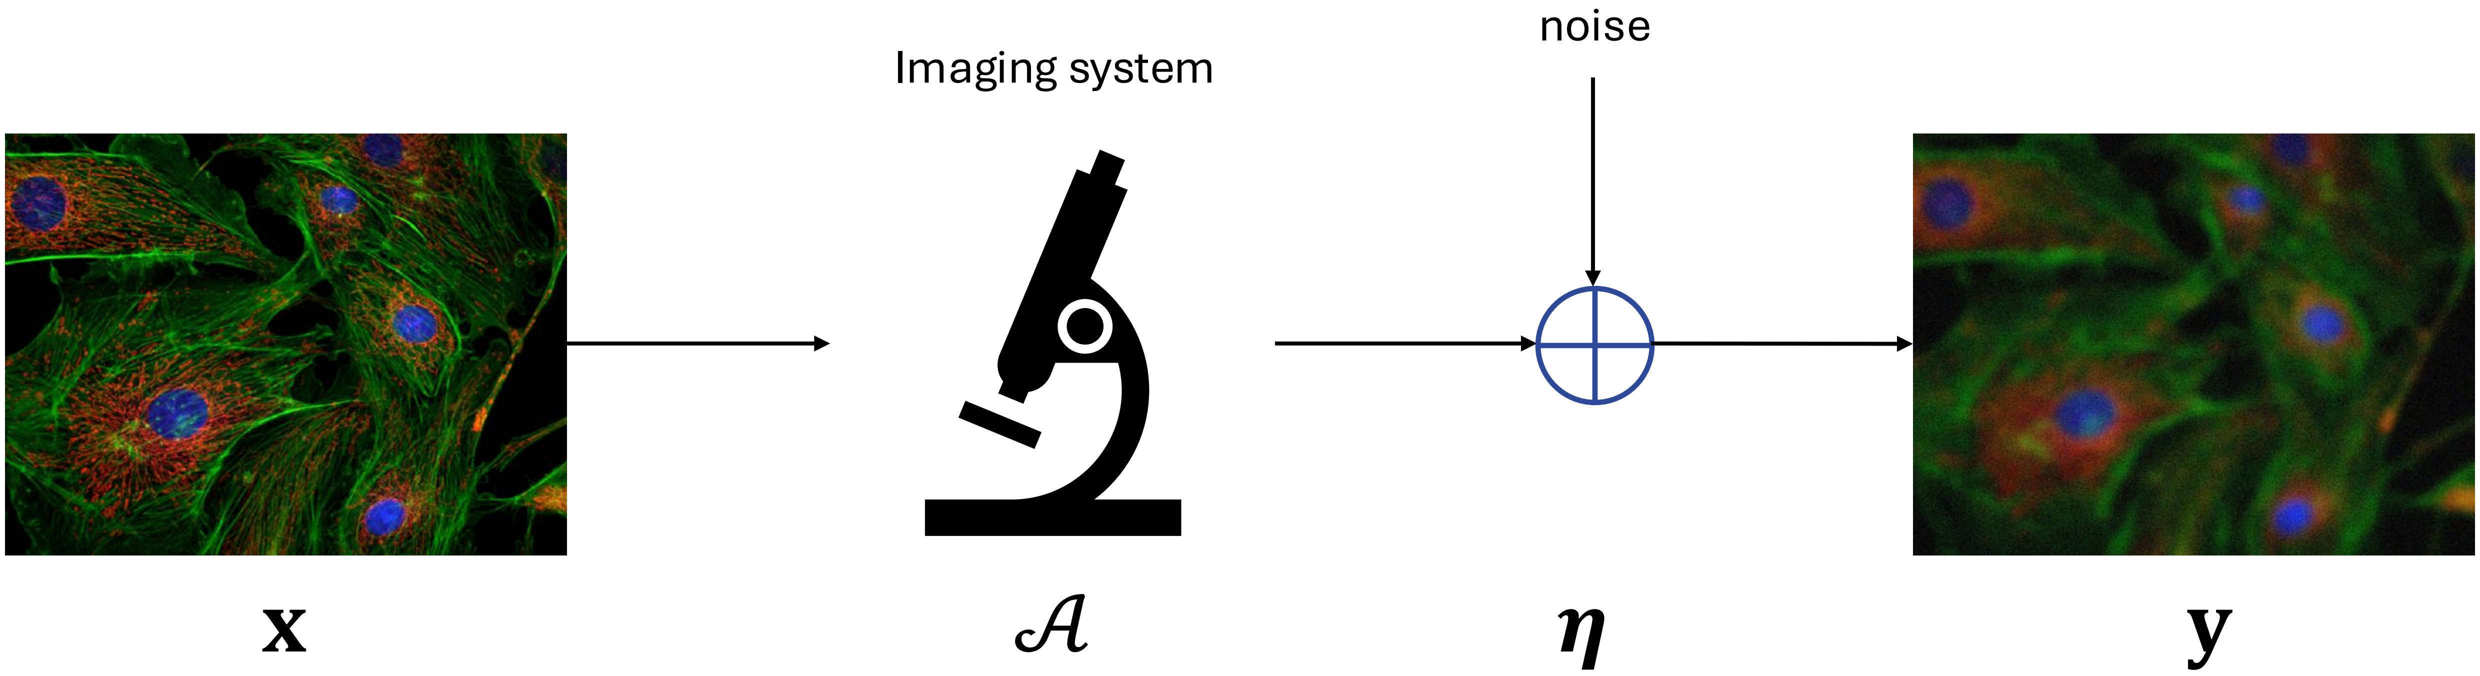

Biomedical imaging systems are cast as the above **forward model**

- Acquiring the image $\mathbf{x}$: **inverse problem**
- System naturally **ill-posed**: what is the best solution?
- Examples include: microscopy, MRI, CT, optics, etc.

Inverse Problems in Medical Image Reconstruction

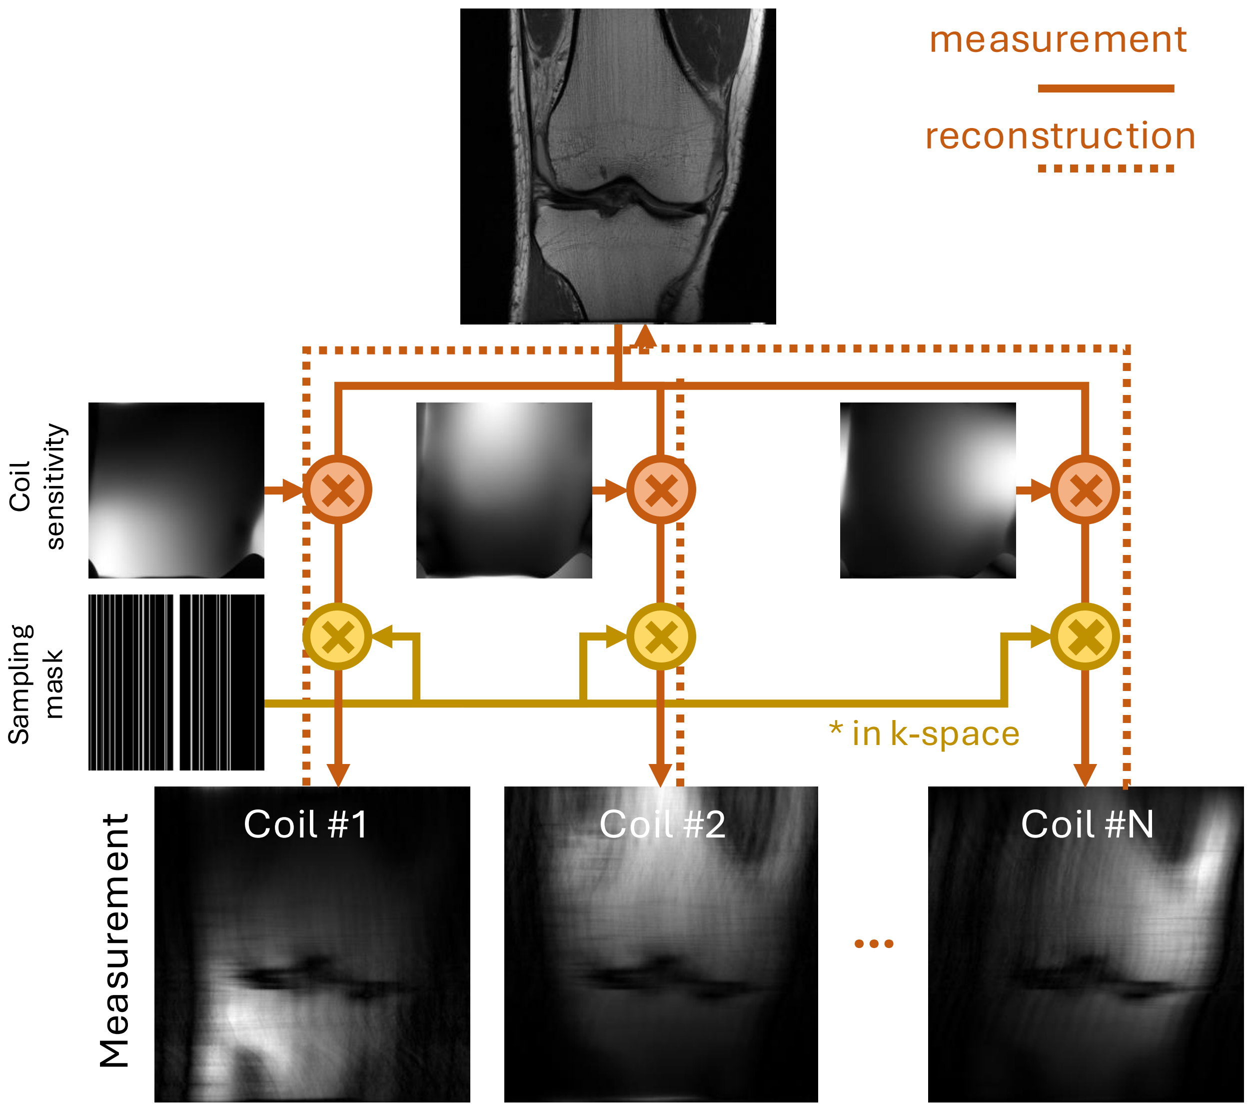

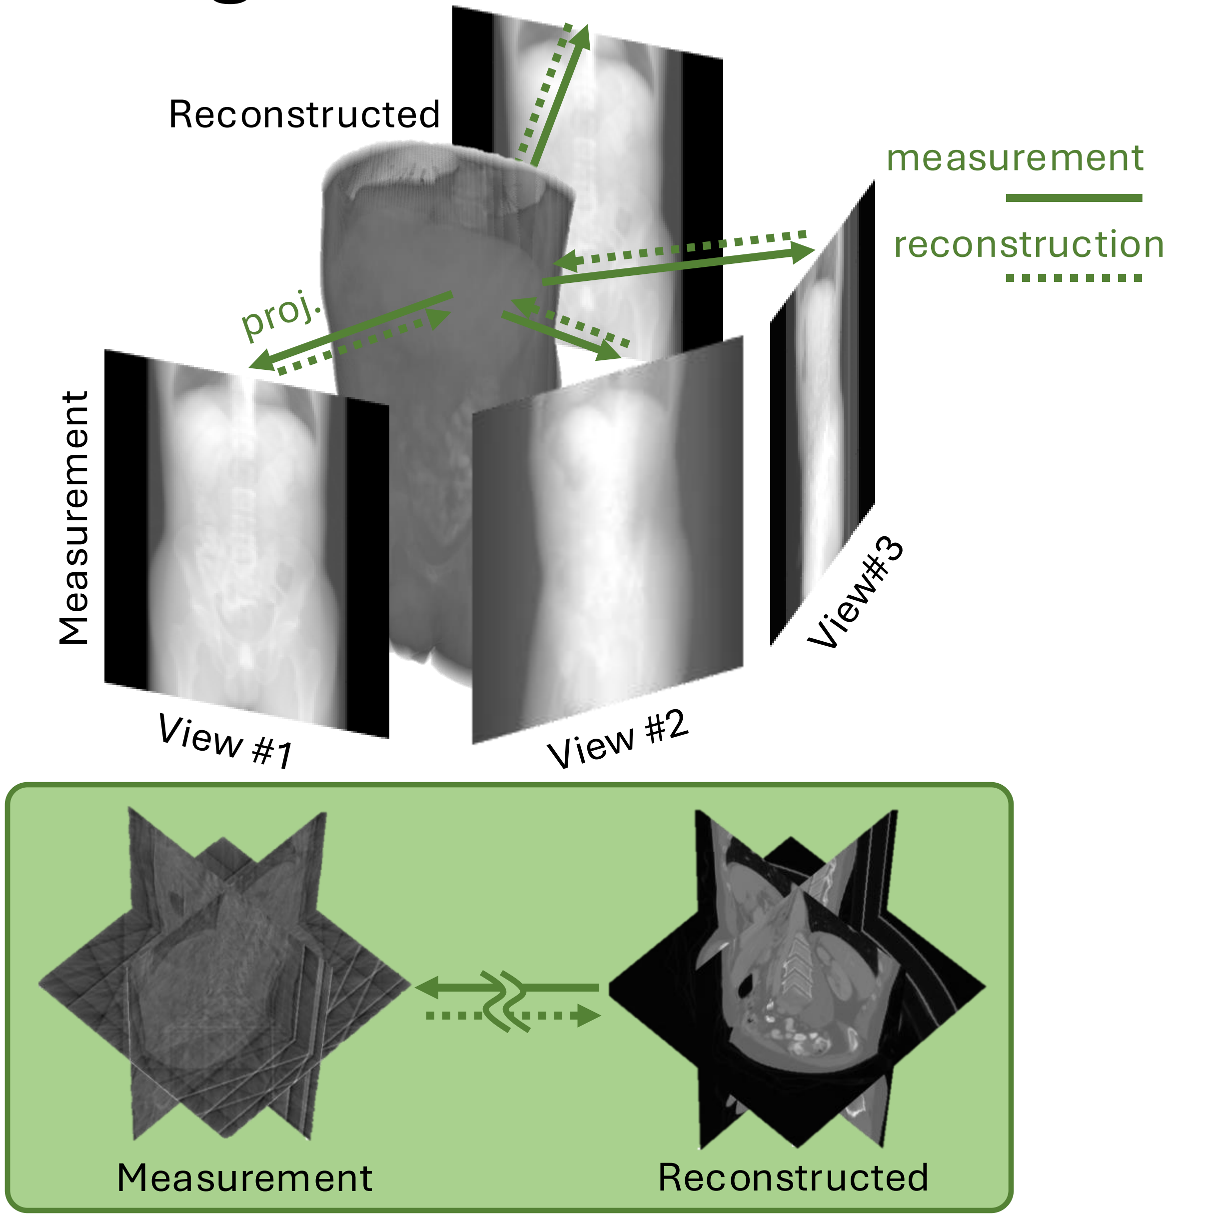

Sampling from the Prior with Diffusion Models

PF-ODE

$$d\boldsymbol{x}_t = -t \boxed{\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t)}\, dt$$

$$\approx s_\theta(\boldsymbol{x}_t, t)$$

Trained through (denoising) score matching

Sampling from the Posterior with Diffusion Models

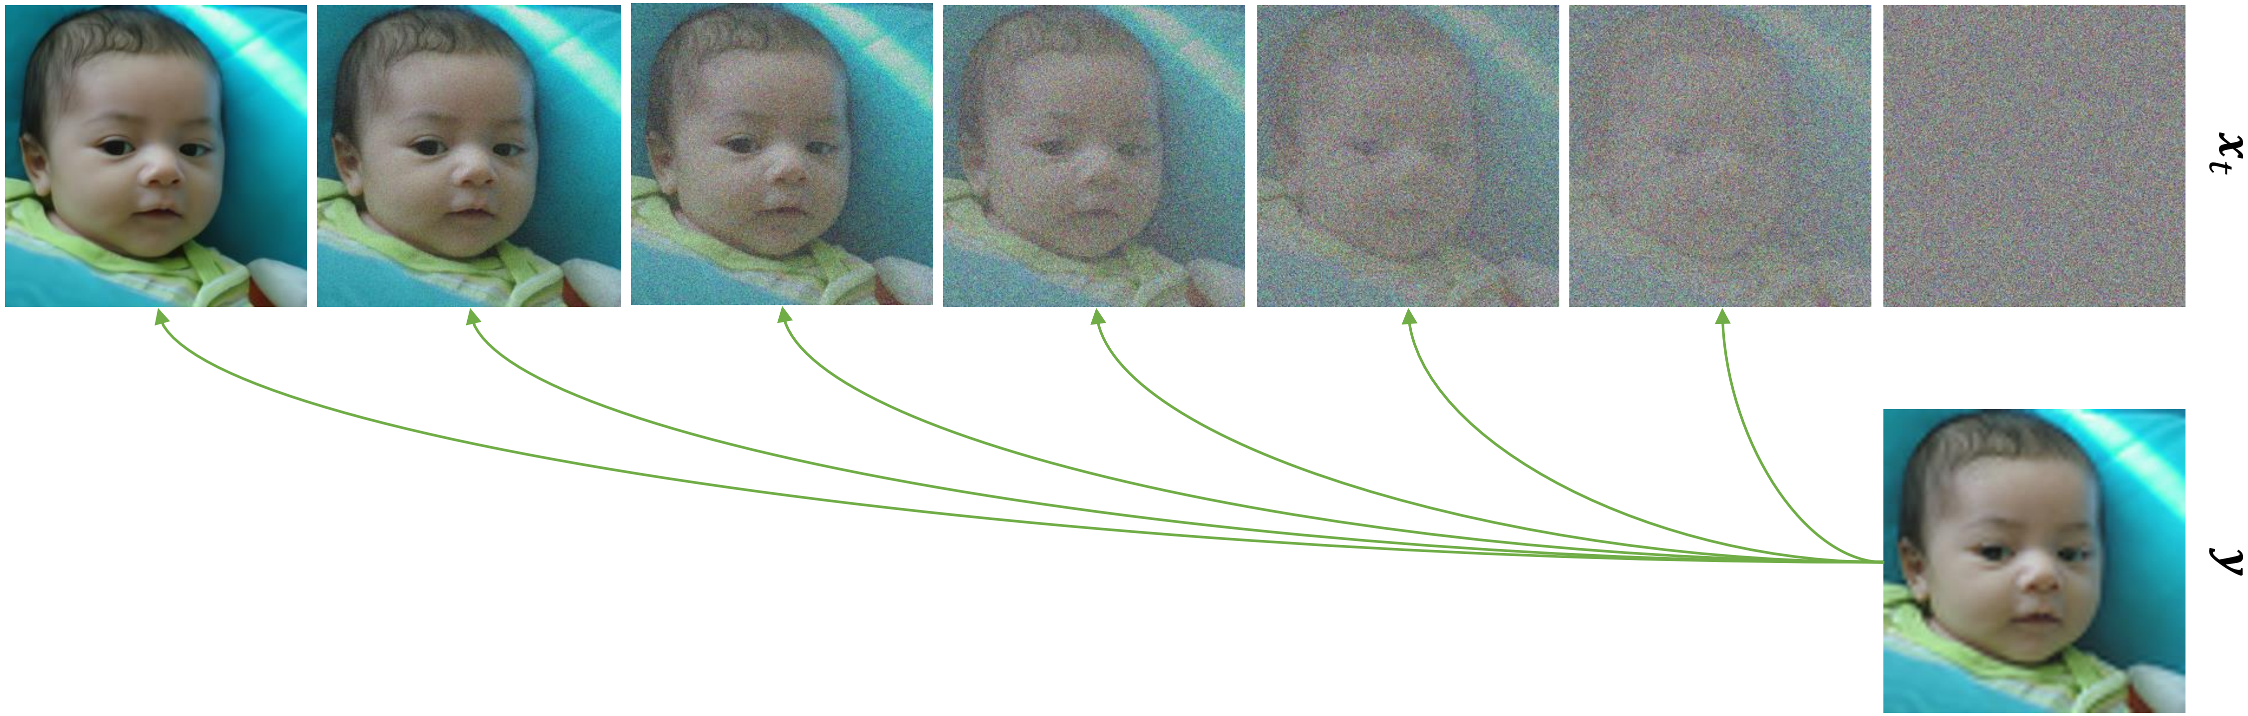

$$d\boldsymbol{x}_t = -t\boxed{\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t|\boldsymbol{y})}\,dt = \frac{\boldsymbol{x}_t - \boxed{\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}]}}{t}\,dt$$

$$\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}] = \boldsymbol{x}_t + t^2\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t|\boldsymbol{y})$$

Time-dependent Bayes Rule

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

Time-dependent Bayes Rule

$$\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t|\boldsymbol{y}) = s_{\theta^*}(\boldsymbol{x}_t) + \boxed{\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{y}|\boldsymbol{x}_t)}$$

Diffusion Posterior Sampling

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

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

Push expectation inside $\approx p(\boldsymbol{y}|\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t])$

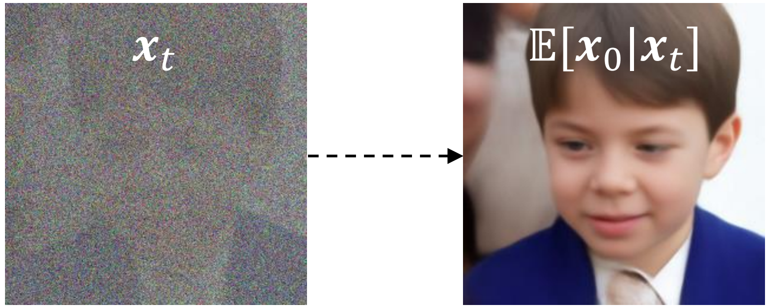

Diffusion Posterior Sampling

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

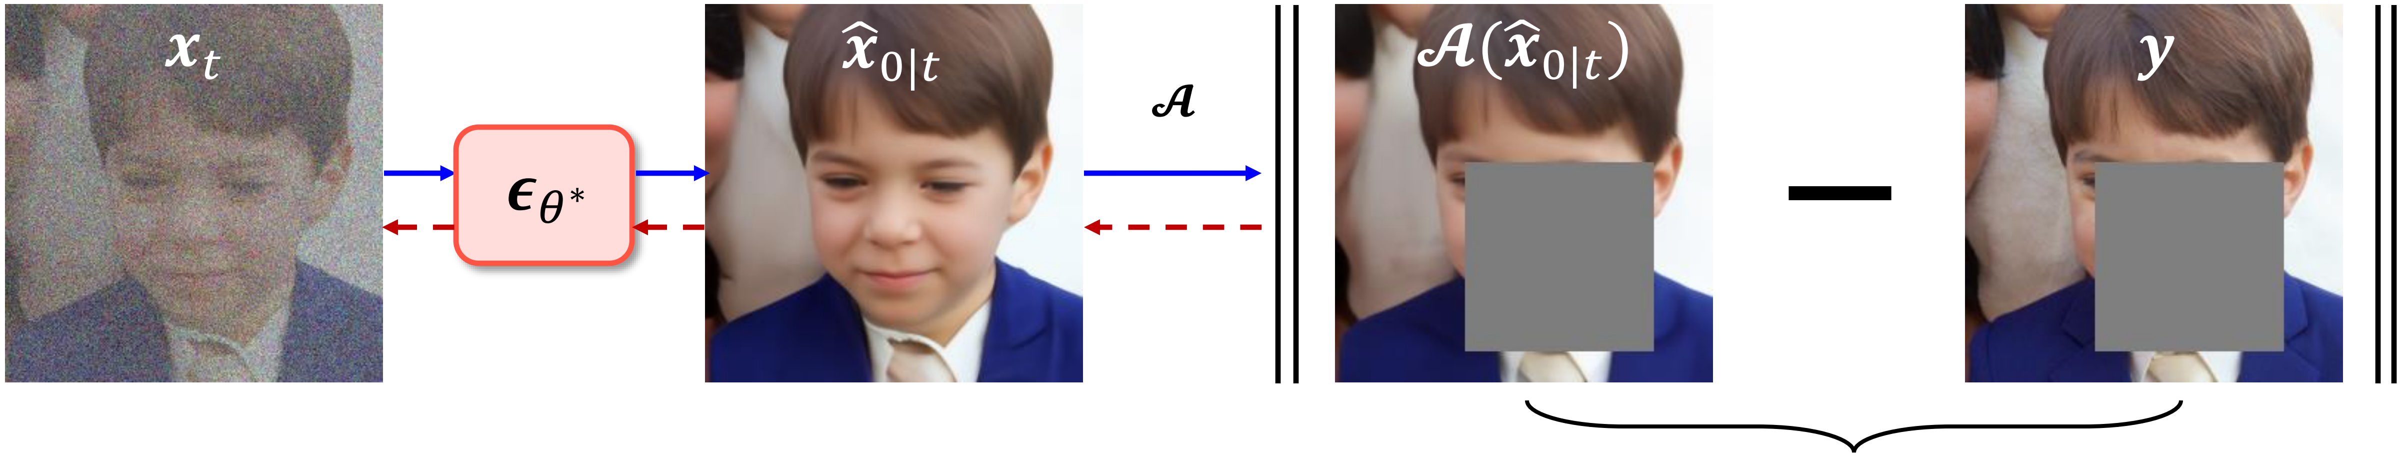

(Residual error)

Forward pass

Backward propagation

Diffusion Posterior Sampling

Algorithm 1 DPS - Gaussian

Require: $N, \boldsymbol{y}, \alpha$

$\boldsymbol{x}_N \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$
for $i = N - 1$ **to** 0 **do**
  $s \leftarrow s_\theta(\boldsymbol{x}_i, i)$
  $\hat{\boldsymbol{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_i}}(\boldsymbol{x}_i + \sqrt{1 - \bar{\alpha}_i})s$
  $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$
  $\boldsymbol{x}'_{i-1} \leftarrow \frac{1}{\sqrt{\alpha_i}}(\boldsymbol{x}_i + (1 - \alpha_i)s) + \sigma_i \boldsymbol{z}$
  $\boldsymbol{x}_{i-1} \leftarrow \boldsymbol{x}'_{i-1} - \alpha \nabla_{\boldsymbol{x}_i} \|\boldsymbol{y} - \mathcal{A}(\hat{\boldsymbol{x}}_0)\|_2^2$
end for
return $\hat{\mathbf{x}}_0$

Algorithm 2 DPS - Poisson

Require: $N, \boldsymbol{y}, \alpha$

$\boldsymbol{x}_N \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$
for $i = N - 1$ **to** 0 **do**
  $s \leftarrow s_\theta(\boldsymbol{x}_i, i)$
  $\hat{\boldsymbol{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_i}}(\boldsymbol{x}_i + \sqrt{1 - \bar{\alpha}_i})s$
  $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$
  $\boldsymbol{x}'_{i-1} \leftarrow \frac{1}{\sqrt{\alpha_i}}(\boldsymbol{x}_i + (1 - \alpha_i)s) + \sigma_i \boldsymbol{z}$
  $\boldsymbol{x}_{i-1} \leftarrow \boldsymbol{x}'_{i-1} - \alpha \nabla_{\boldsymbol{x}_i} \|\boldsymbol{y} - \mathcal{A}(\hat{\boldsymbol{x}}_0)\|_{\boldsymbol{\Lambda}}^2$
end for
return $\hat{\mathbf{x}}_0$

Chung et al., "*Diffusion Posterior Sampling for General Noisy Inverse Problems*", ICLR 2023

Two Interpretations of DPS

1. Conditional score function

$$\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t | \boldsymbol{y}) \approx \nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t) + \nabla_{\boldsymbol{x}_t} \log p\left(\boldsymbol{y} \middle| \hat{\boldsymbol{x}}_{0|t}\right)$$

2. Conditional posterior mean

$$\mathbb{E}[\boldsymbol{x}_0 | \boldsymbol{x}_t, \boldsymbol{y}] \approx \mathbb{E}[\boldsymbol{x}_0 | \boldsymbol{x}_t] + \zeta_t^2 \nabla_{\boldsymbol{x}_t} \left\| \boldsymbol{y} - A\hat{\boldsymbol{x}}_{0|t} \right\|^2$$

Chung et al., "*Diffusion Posterior Sampling for General Noisy Inverse Problems*", ICLR 2023

Decomposed Diffusion Sampler

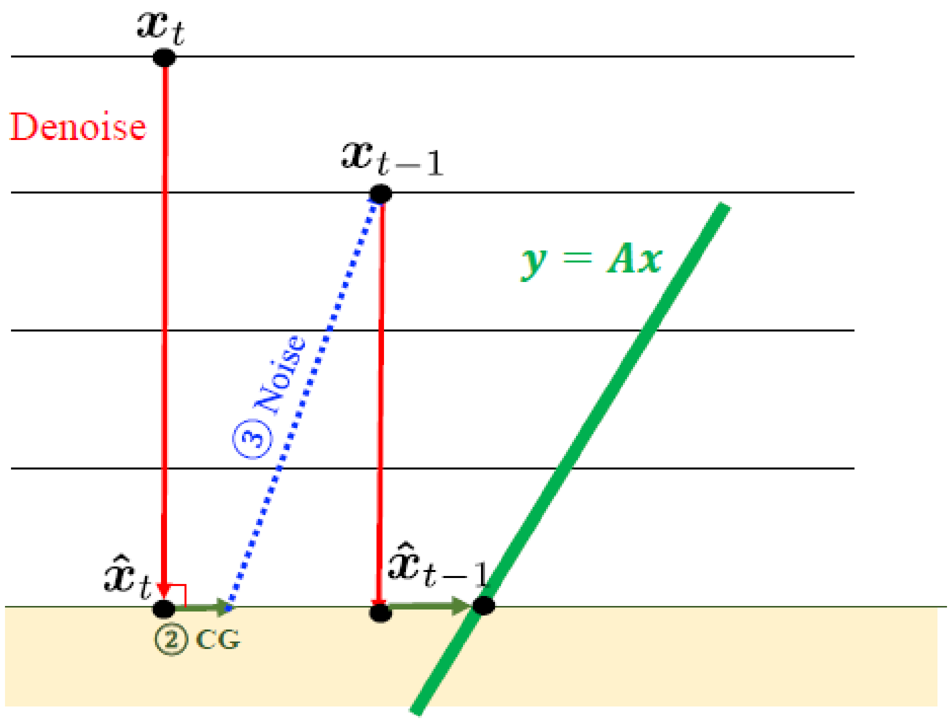

1. Posterior mean estimation

$$\hat{\boldsymbol{x}}_{0|t} = \mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}}(\boldsymbol{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t))$$

2. Conditional posterior mean (optimization with CG)

$$\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}] \approx \operatorname*{argmin}_{\boldsymbol{x}} \frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{y}\|^2 + \frac{\lambda}{2}\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{0|t}\|^2$$

$$= \mathrm{CG}(A^TA + \lambda I, A^Ty + \lambda\hat{\boldsymbol{x}}_{0|t})$$

3. Update to next level

$$\boldsymbol{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}] + \sqrt{1-\bar{\alpha}_{t-1}}\boldsymbol{\epsilon}$$

Chung et al., "*Decomposed Diffusion Sampler for Large-scale Inverse Problems*", ICLR 2024

Results: Multi-coil MRI

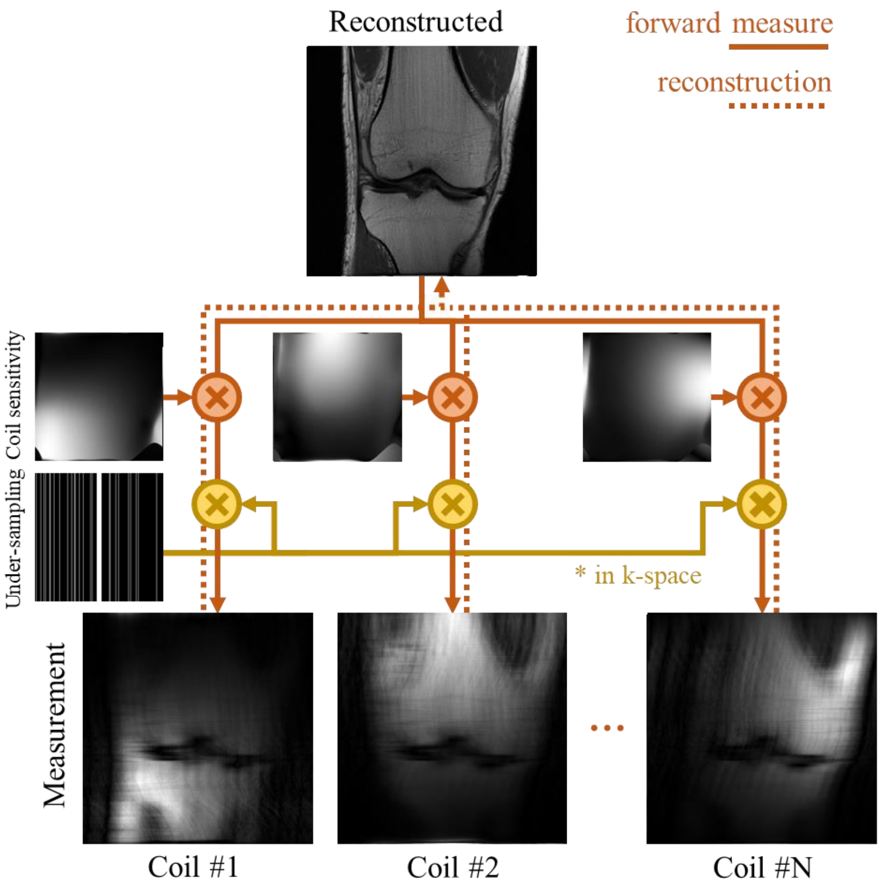

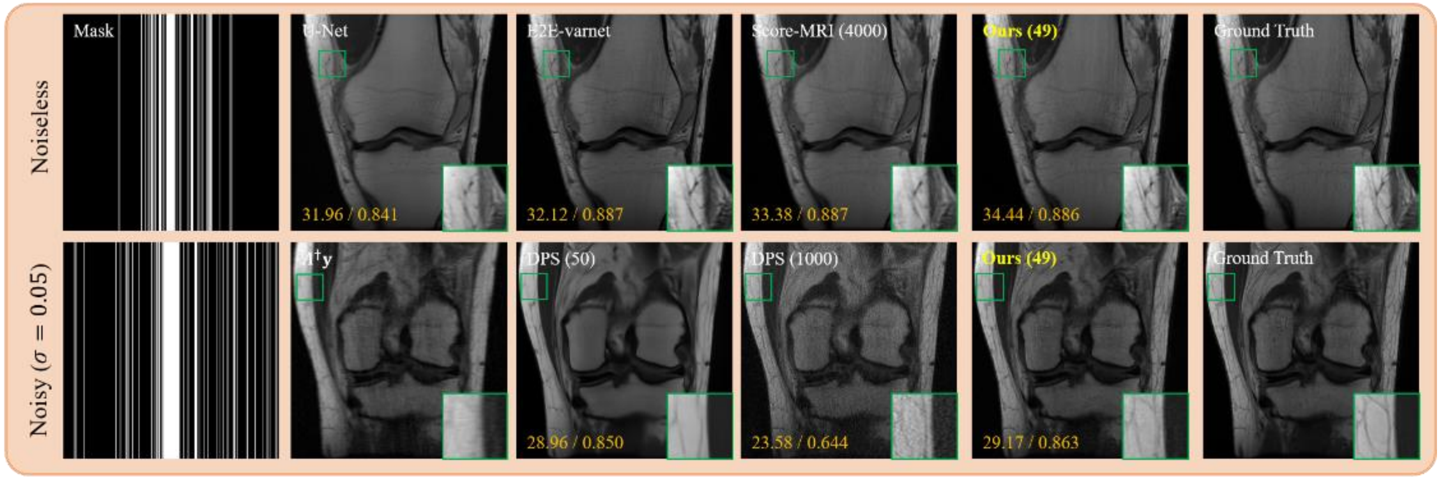

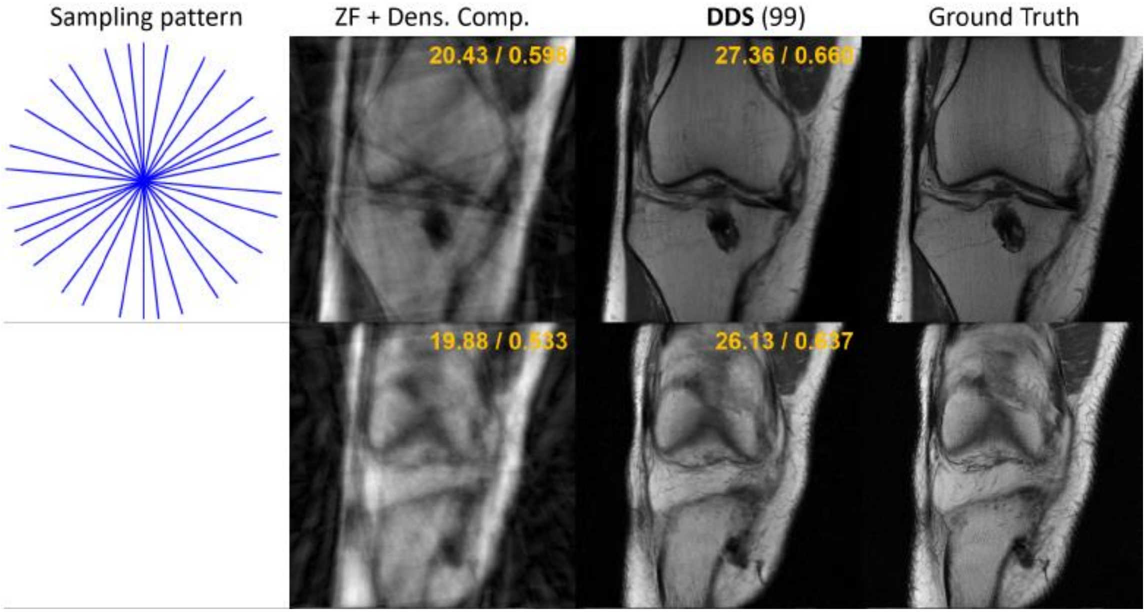

Chung et al., "*Decomposed Diffusion Sampler for Large-scale Inverse Problems*", ICLR 2024

3D Prior from 2D Diffusion Prior

Independent of z-dim.
→ **Independent slices**

Independent of xy-dim.
→ **Operates on** z-dim only

$$p_\theta(\mathbf{x}) := \overbrace{p_\theta^{xy}(\mathbf{x})}\,\overbrace{p^z(\mathbf{x})}$$

Proposal: factored prior → Reduced computation!

3D Prior from 2D Diffusion Prior

$$\nabla_{\mathbf{x}} \log p_{\theta}^{xy}(\mathbf{x}) \simeq s_{\theta^*}(\mathbf{x})$$

Generative diffusion prior (xy): Pre-trained 2D diffusion model

$$-\log p^{z}(\mathbf{x}) \simeq TV^{z}(\mathbf{x})$$

Model-based prior (z): Total-variation to impose smoothness

3D Prior from 2D Diffusion Prior

1. Posterior mean estimation

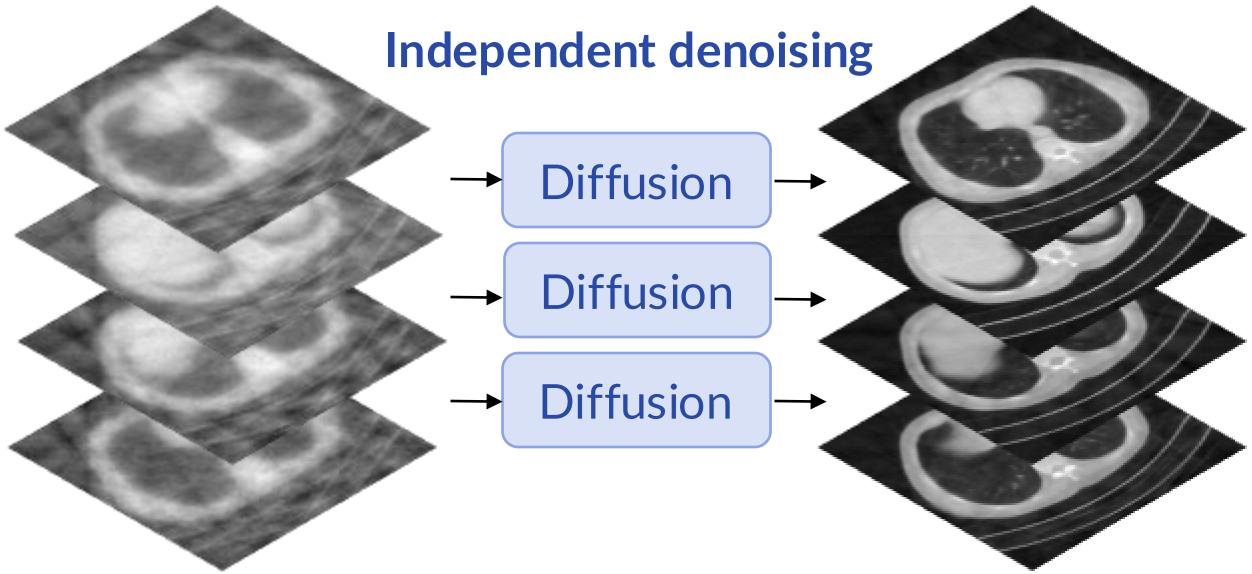

$$\mathbb{E}\left[\boldsymbol{x}_0^{(i)} \middle| \boldsymbol{x}_t^{(i)}\right] = \frac{1}{\sqrt{\bar{\alpha}_t}}\left(\boldsymbol{x}_t^{(i)} - \sqrt{1-\bar{\alpha}_t}\,\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t^{(i)})\right)$$

2. Conditional posterior mean

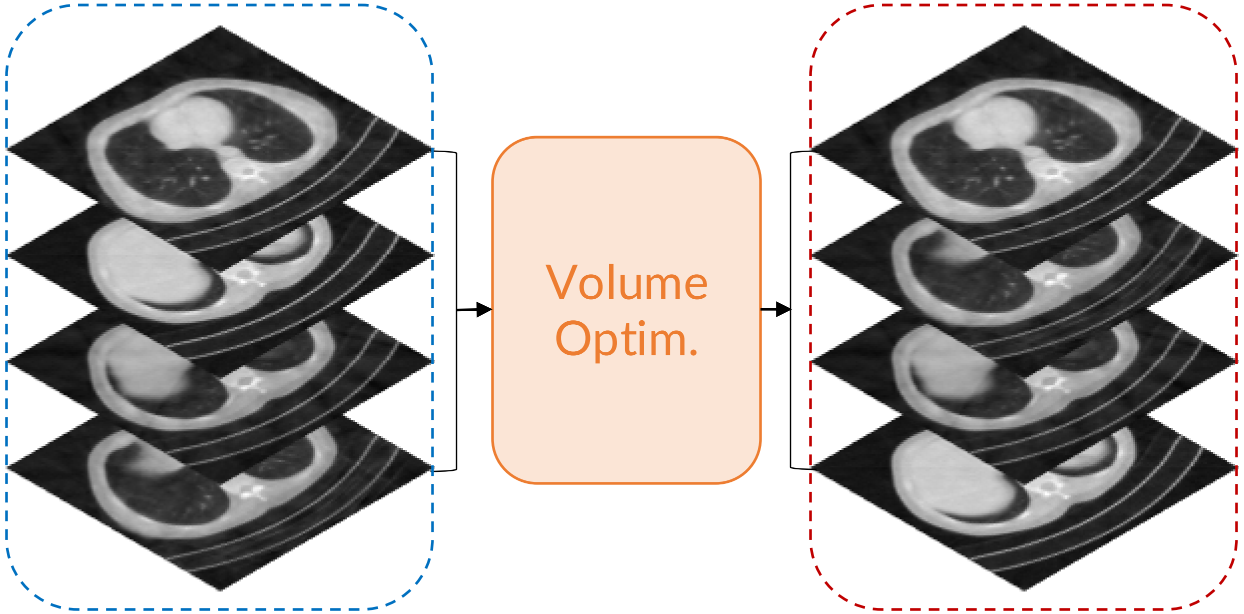

$$\mathbb{E}[\boldsymbol{X}_0|\boldsymbol{X}_t,\boldsymbol{Y}] \approx \operatorname*{argmin}_{\boldsymbol{X}} \frac{1}{2}\|A\boldsymbol{X}-\boldsymbol{Y}\|_2^2 + \frac{\lambda}{2}\|D_z\boldsymbol{X}\|_1$$

$$\boldsymbol{X} := [\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \cdots, \boldsymbol{x}^{(N)}]$$

Chung et al., "*Solving 3D Inverse Problems using Pre-trained 2D Diffusion Models*", CVPR 2023

Results: 3D Sparse-view CT

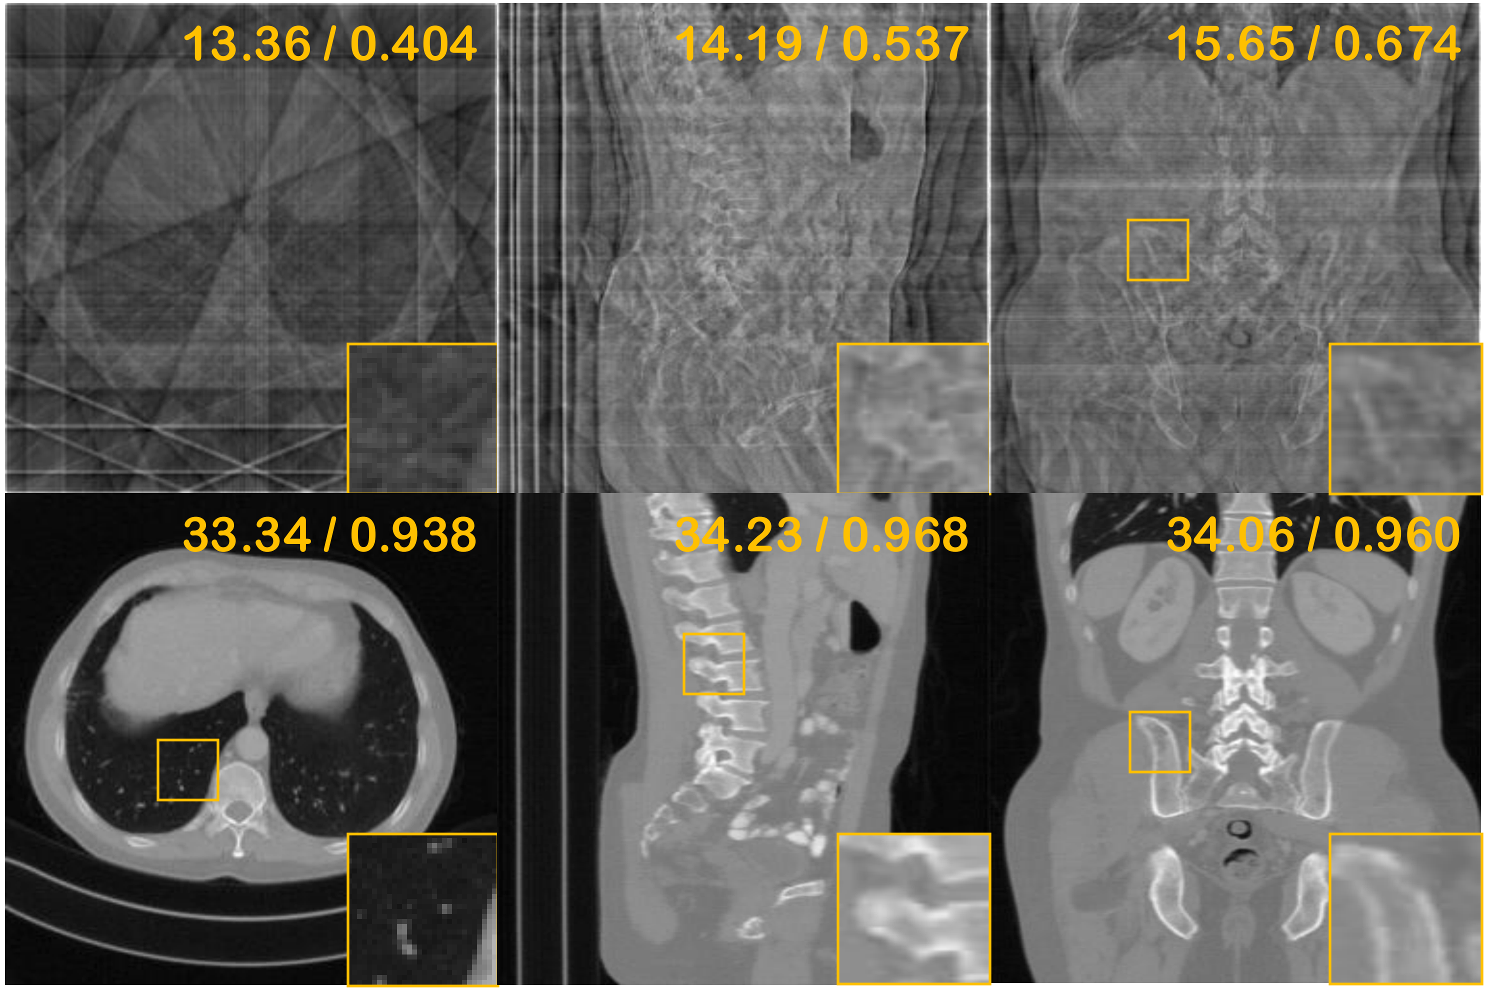

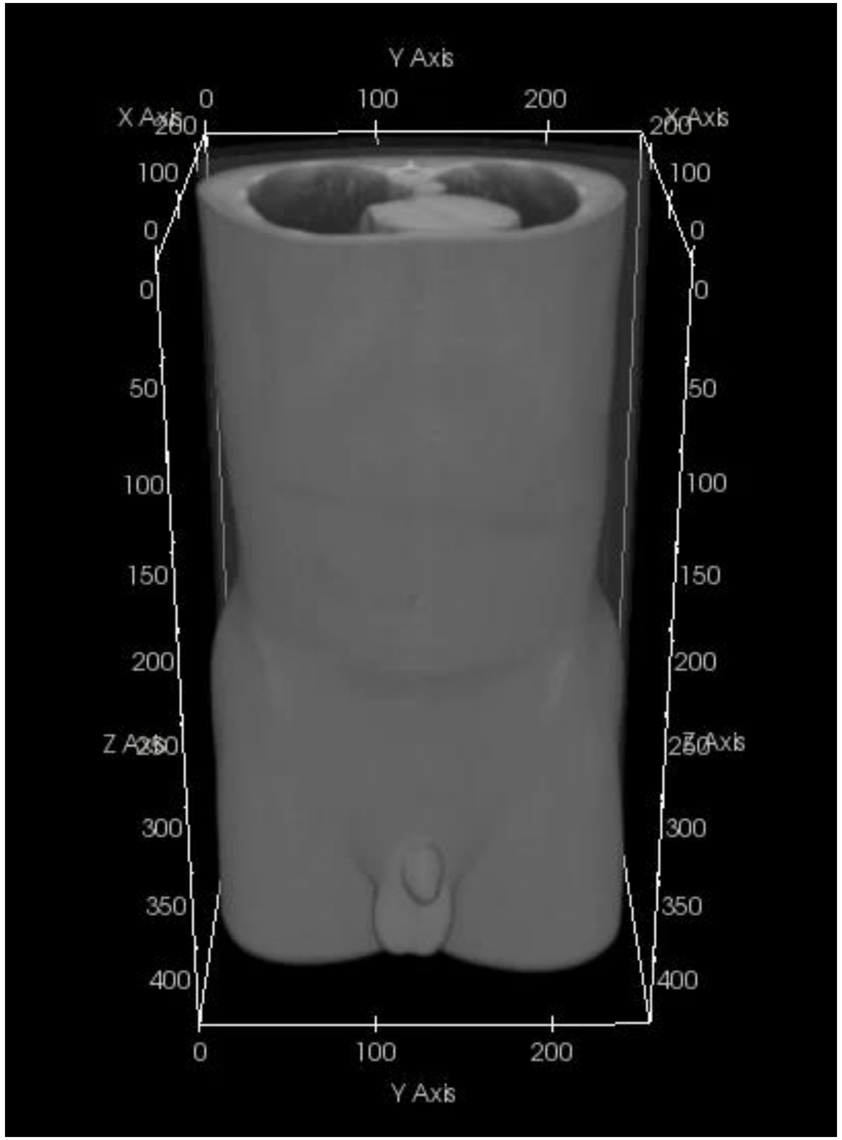

Chung et al., “*Solving 3D Inverse Problems using Pre-trained 2D Diffusion Models*”, CVPR 2023

Factored Diffusion Prior for Both Planes

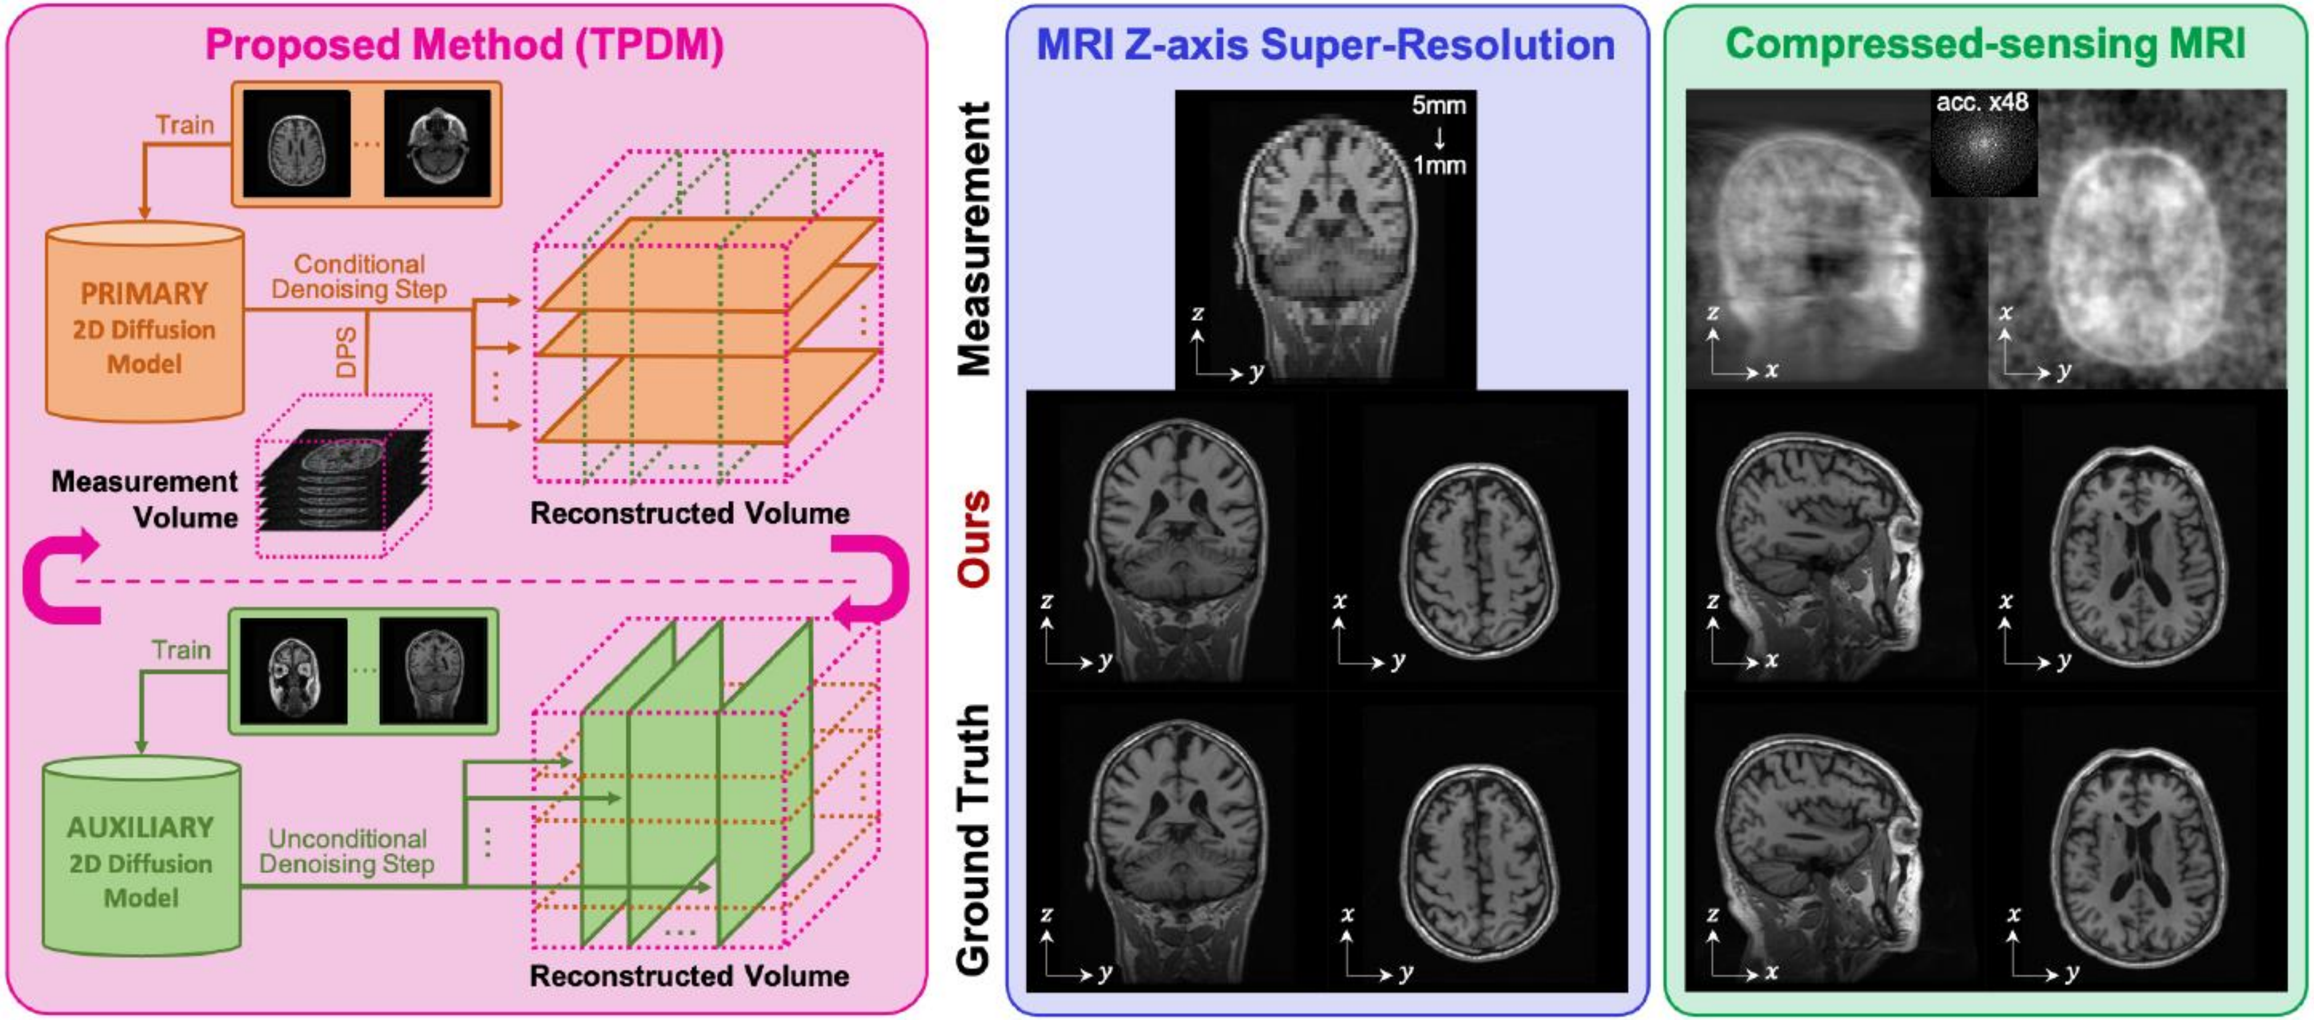

Lee & Chung et al., "*Improving 3D imaging with pre-trained perpendicular 2D diffusion models*", ICCV 2023

Factored Diffusion Prior for Both Planes

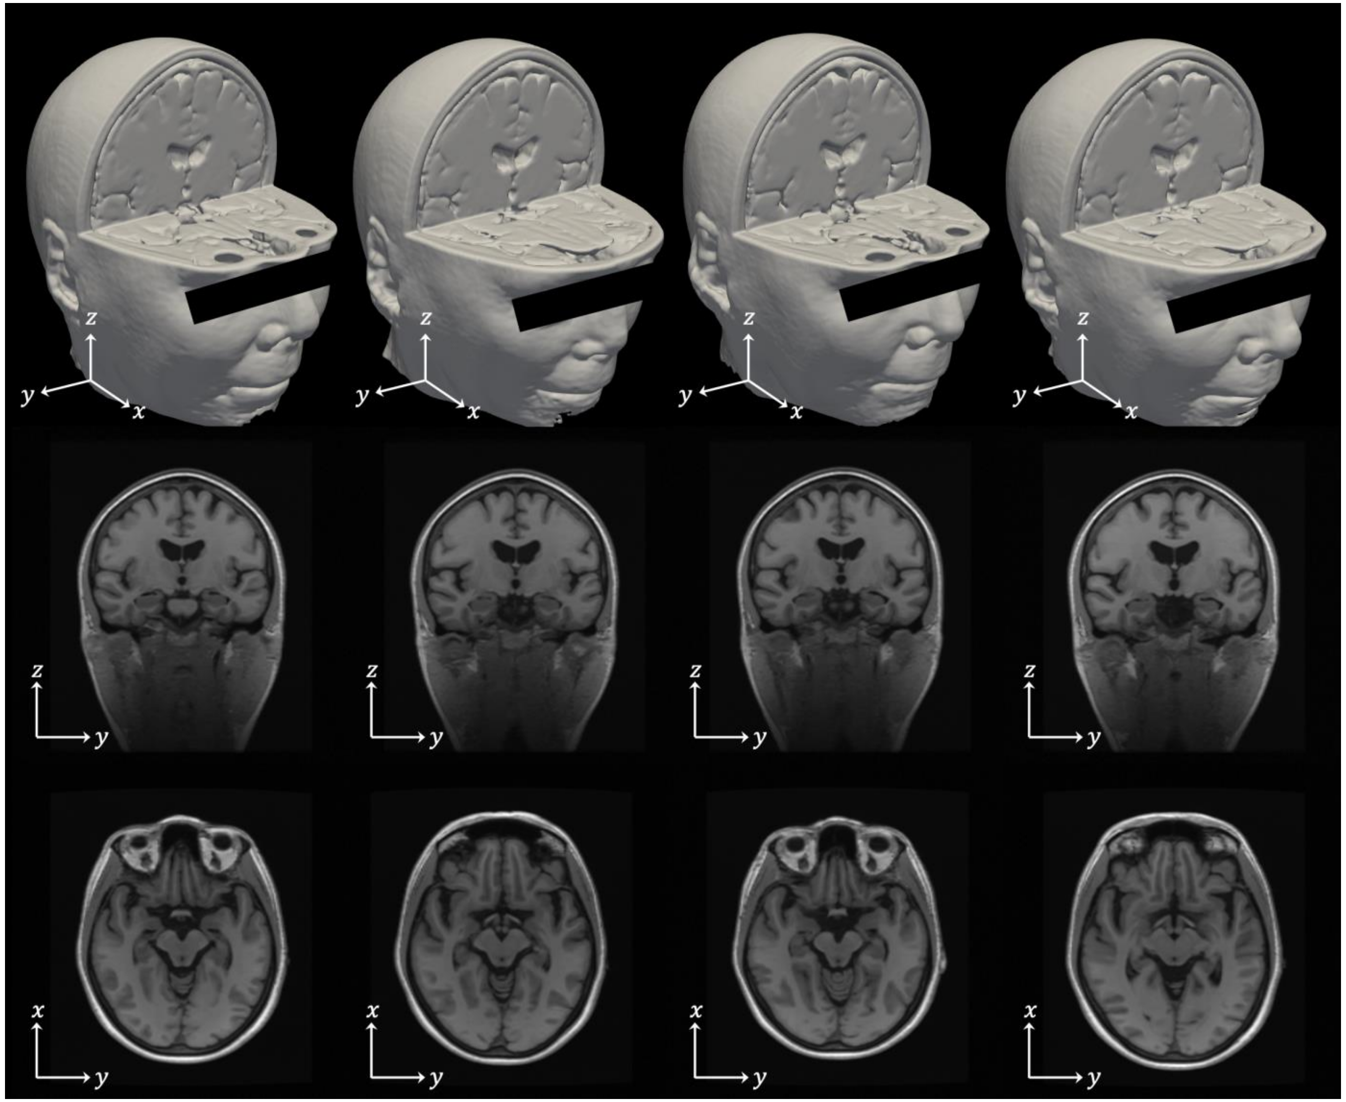

Lee & Chung et al., "*Improving 3D imaging with pre-trained perpendicular 2D diffusion models*", ICCV 2023

Other Approximations

$$\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{y}|\boldsymbol{X}_t = \boldsymbol{x}_t)$$

Score ALD:
$$\propto -A^\top (\boldsymbol{y} - A\boldsymbol{x}_t)$$

Score-SDE:
$$\propto -A^\top (\boldsymbol{y}_t - A\boldsymbol{x}_t)$$

ILVR:
$$\propto -(A^\top A)^{-1}A^\top (\boldsymbol{y}_t - A\boldsymbol{x}_t)$$

DPS:
$$\propto \left(I + \nabla^2_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\right)^\top A^\top \cdot \Big(\boldsymbol{y} - A\mathbb{E}[\boldsymbol{X}_0|\boldsymbol{X}_t = \boldsymbol{x}_t]\Big)$$

ΠGDM:
$$\propto -\frac{\partial \mathbb{E}[\boldsymbol{X}_0|\boldsymbol{X}_t = \boldsymbol{x}_t]}{\partial \boldsymbol{x}_t}(r_t^2 AA^\top + \sigma_{\boldsymbol{y}}^2 I)^{-1}A^\top \cdot (\boldsymbol{y} - A\mathbb{E}[\boldsymbol{X}_0|\boldsymbol{X}_t = \boldsymbol{x}_t])$$

Moment Matching:
$$\propto -\nabla_{\boldsymbol{x}_t}\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t]^\top A^\top(\sigma_{\boldsymbol{y}}^2 I + A\sigma_t^2 \nabla_{\boldsymbol{x}_t}\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t]A^\top)^{-1} \cdot (\boldsymbol{y} - A\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t])$$

SNIPS:
$$\propto -\Sigma^\top \left|\sigma_{\boldsymbol{y}}^2 I_m - \sigma_t^2 \Sigma\Sigma^\top\right|^\dagger (\bar{\boldsymbol{y}} - \Sigma\bar{\boldsymbol{x}}_t)$$

DDRM:
$$\propto -\Sigma^\top \left|\sigma_{\boldsymbol{y}}^2 I_m - \sigma_t^2 \Sigma\Sigma^\top\right|^\dagger (\bar{\boldsymbol{y}} - \Sigma\bar{\boldsymbol{x}}_{0|t})$$

DDNM:
$$\propto \Sigma_t A^\dagger \Big(\boldsymbol{y} - A\mathbb{E}[\boldsymbol{X}_0|\boldsymbol{X}_t = \boldsymbol{x}_t, \boldsymbol{y}]\Big)$$

$$\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{y}|\boldsymbol{X}_t = \boldsymbol{x}_t) \approx -\frac{\mathcal{L}_t \quad \mathcal{M}_t}{\mathcal{G}_t}.$$

Daras et al., “*A Survey on Diffusion Models for Inverse Problems*”

Out-of-distribution Test Data

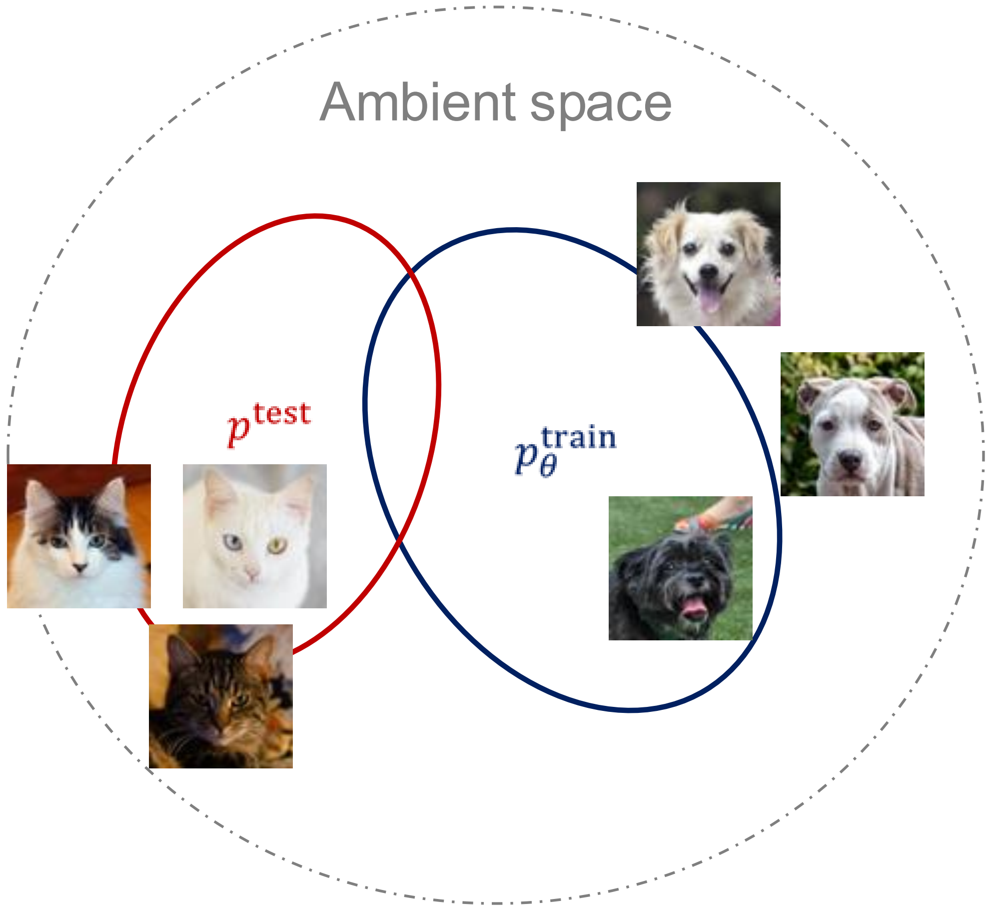

Black Hole Imaging: Unknown Prior

Sparse Fourier phase retrieval

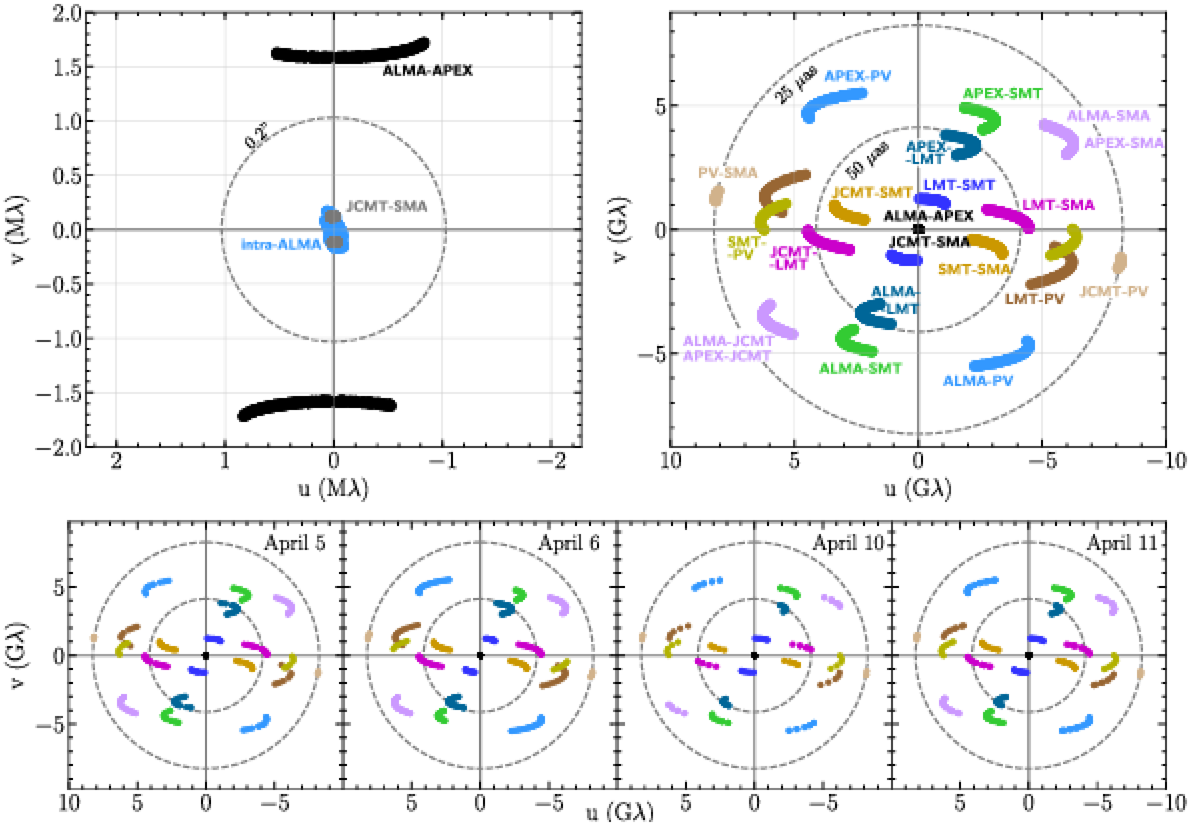

Prior models used: unknown

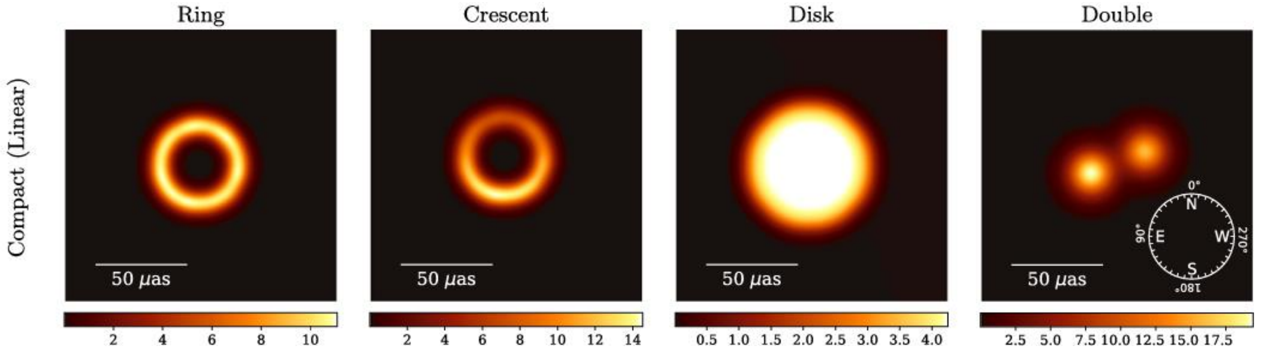

If we don't have a dataset to train a generative model

We always have prior mismatch (OOD)

The best we can hope for

- Learning from the measurement
- Train on phantoms, then deal with OOD

Akiyama et al., "*First M87 event horizon telescope results. IV. Imaging the central supermassive black hole*", The Astrophysical Journal Letters, 2019

Consequence: Medical Imaging

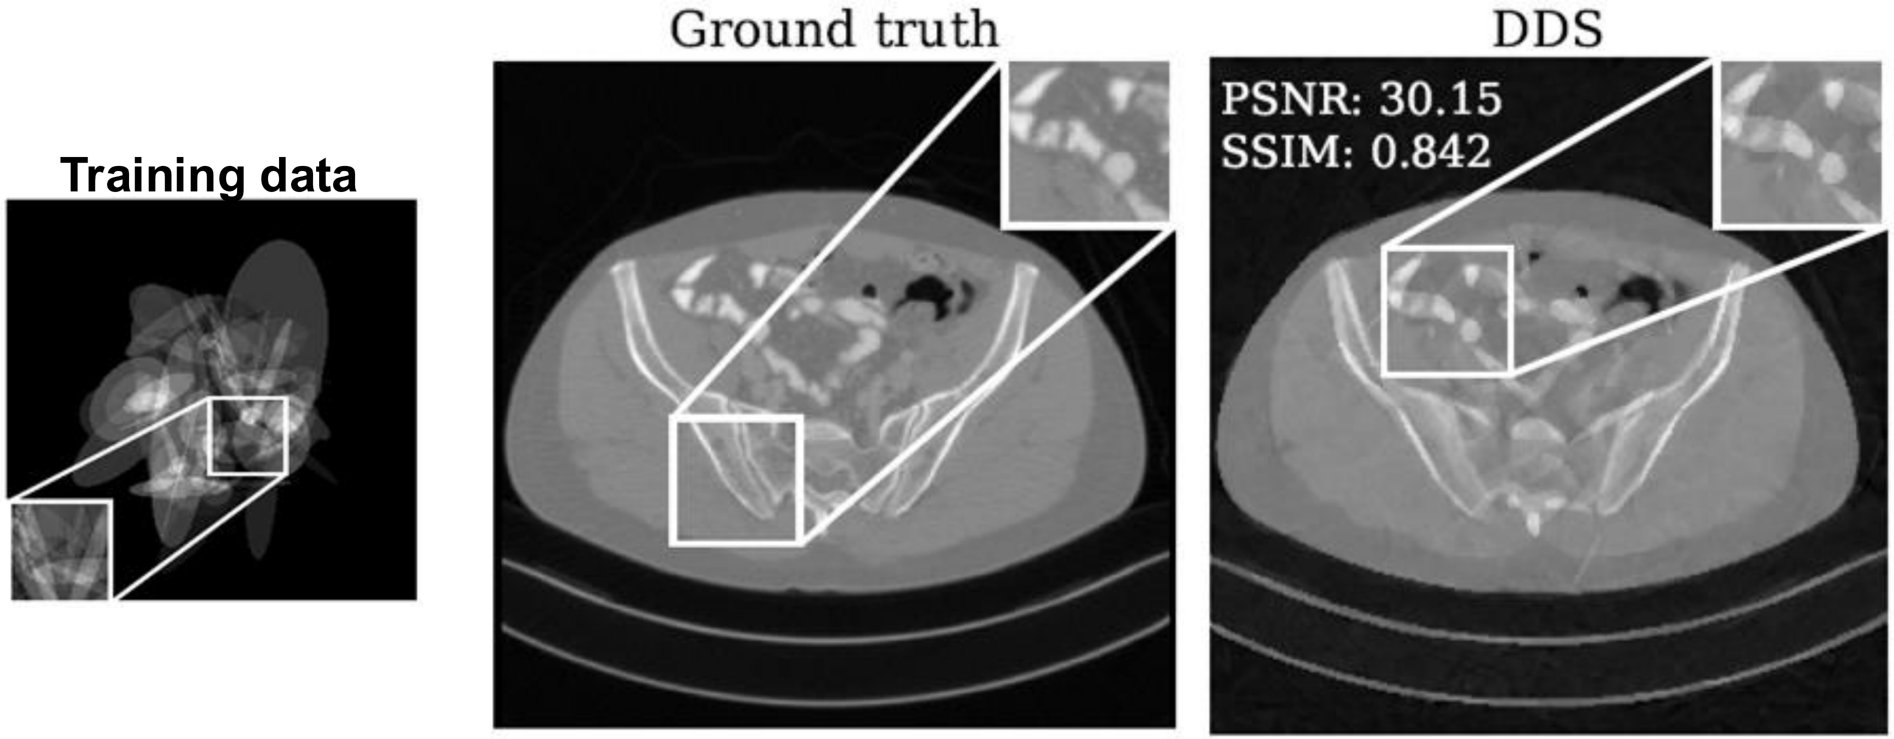

Revisiting Deep Image Prior (DIP)

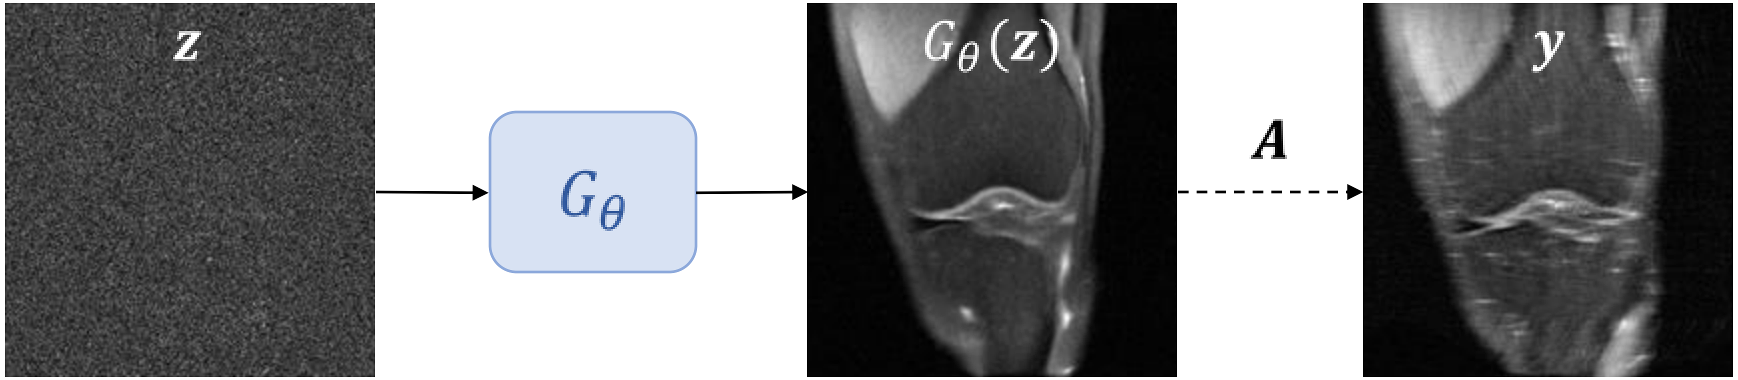

$$\min_{\theta} \|y - AG_{\theta}(z)\|_2^2$$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Deep Diffusion Image Prior (DDIP)

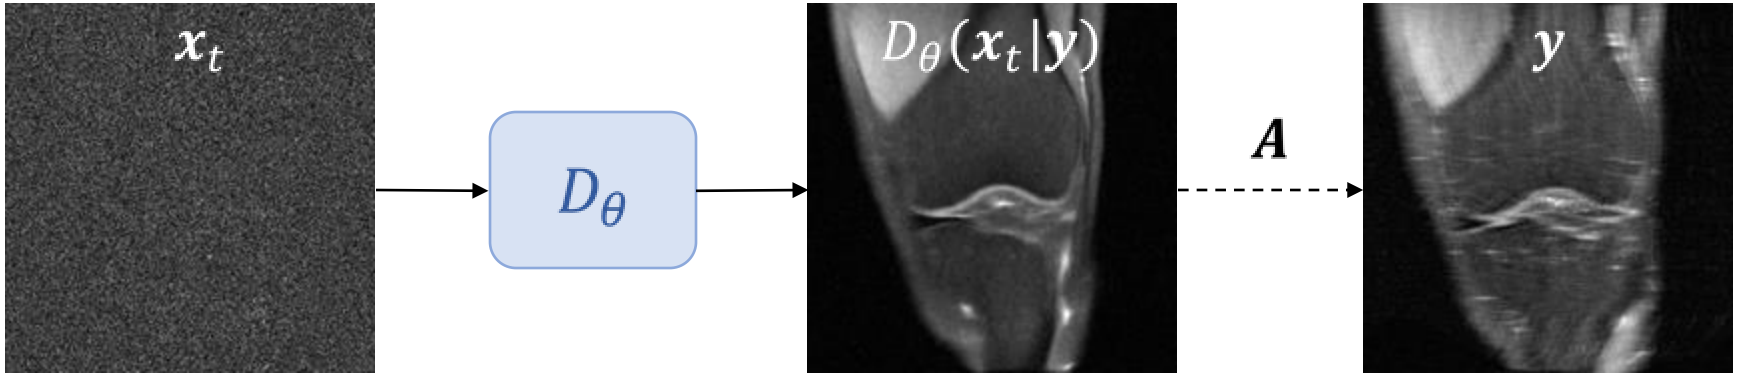

$$\min_{\theta} \|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}_t|\boldsymbol{y})\|_2^2$$

$$\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}] \approx D_\theta(\boldsymbol{x}_t|\boldsymbol{y}) := D_\theta(\boldsymbol{x}_t) + \frac{\sigma_t^2}{2\sigma_y^2}\nabla_{\boldsymbol{x}_t}\|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}_t)\|_2^2$$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Deep Diffusion Image Prior (DDIP)

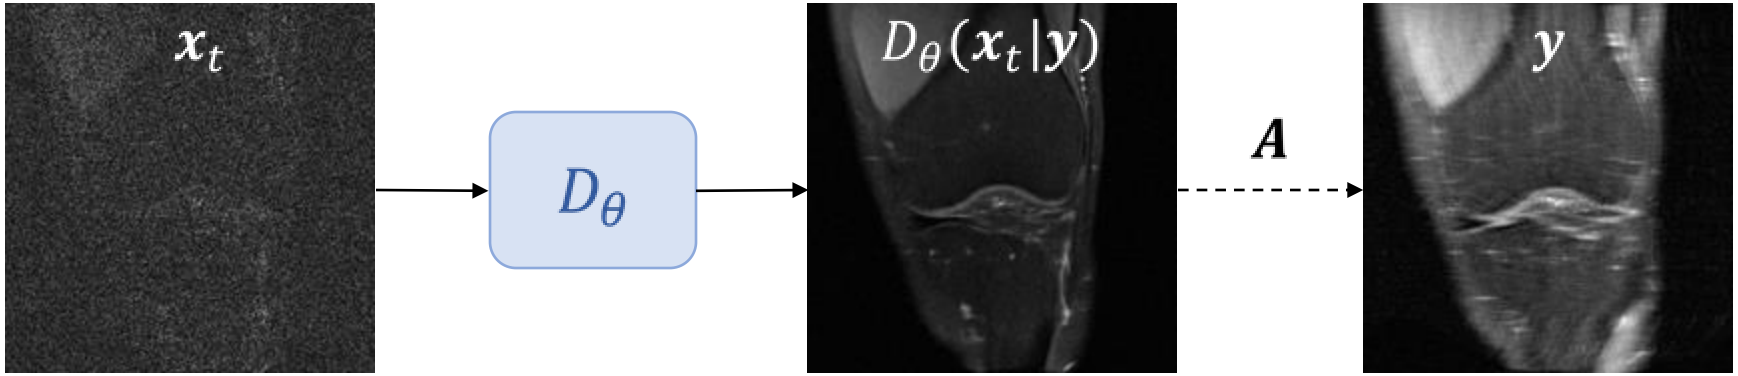

$$\min_{\theta} \|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}_t|\boldsymbol{y})\|_2^2$$

$$\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}] \approx D_\theta(\boldsymbol{x}_t|\boldsymbol{y}) := D_\theta(\boldsymbol{x}_t) + \frac{\sigma_t^2}{2\sigma_y^2}\nabla_{\boldsymbol{x}_t}\|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}_t)\|_2^2$$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Consistency matters with DDIP

$$\text{for } t = T, \ldots, 1 : \theta_{t-1} \leftarrow \arg\min_{\theta_t} \|\boldsymbol{y} - \boldsymbol{A} D_{\theta_t}(\boldsymbol{x}_t|\boldsymbol{y})\|_2^2,$$
$$\boldsymbol{x}_{t-1} \leftarrow \mathrm{DDIM}_{\theta_{t-1}}(D_{\theta_{t-1}}(\boldsymbol{x}_t|\boldsymbol{y}), \eta)$$

Solver	DDS	DDNM	DPS	DDIP							
$\mathbb{E}[\boldsymbol{x}_0\|\boldsymbol{x}_t, \boldsymbol{y}]$	DDS	DDNM	DPS	DDS			DDNM				DPS
K	-	-	-	1 (SCD [3])	3	**5 (ours)**	1	5	10	**20 (ours)**	**1 (ours)**
PSNR	31.37	30.69	23.49	35.43	35.83	**36.31**	34.92	35.50	35.98	**36.07**	**31.32**
SSIM	0.887	0.870	0.608	0.911	0.921	**0.922**	0.902	0.911	0.920	**0.920**	**0.841**

Table 1: Sparse-view CT reconstruction results on ELLIPSES → AAPM OOD reconstruction by varying the solver, $\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t, \boldsymbol{y}]$ approximation method, and the number of CG iterations K. Standard non-adapted DIS , SCD [3] , Proposed DDIP method .

Barbano et al., "*Steerable Conditional Diffusion for Out-of-Distribution Adaptation in Medical Image Reconstruction*", IEEE TMI 2025

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

DIP → DDIP

Deep Image Prior (DIP)

$$\min_{\theta} \|\mathbf{y} - AG_{\theta}(\mathbf{z})\|_2^2$$

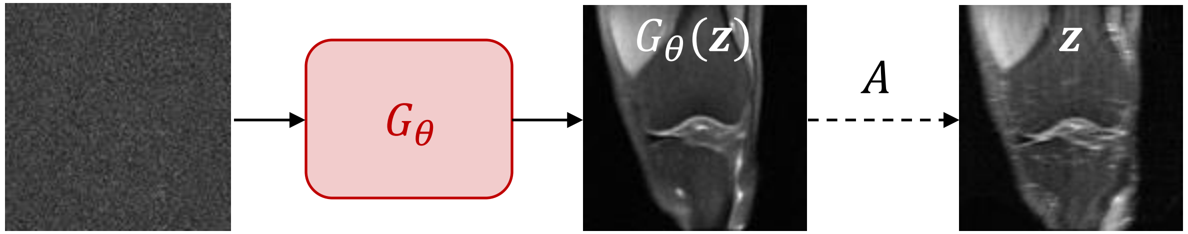

- Training possible due to implicit prior of CNNs

- Slow optimization
- Wide search space (a lot of bad local minima)
- Volatile, often fails

Deep Diffusion Image Prior (DDIP)

$$\min_{\theta} \|\mathbf{y} - A\mathbb{E}^{(\theta)}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}]\|_2^2$$

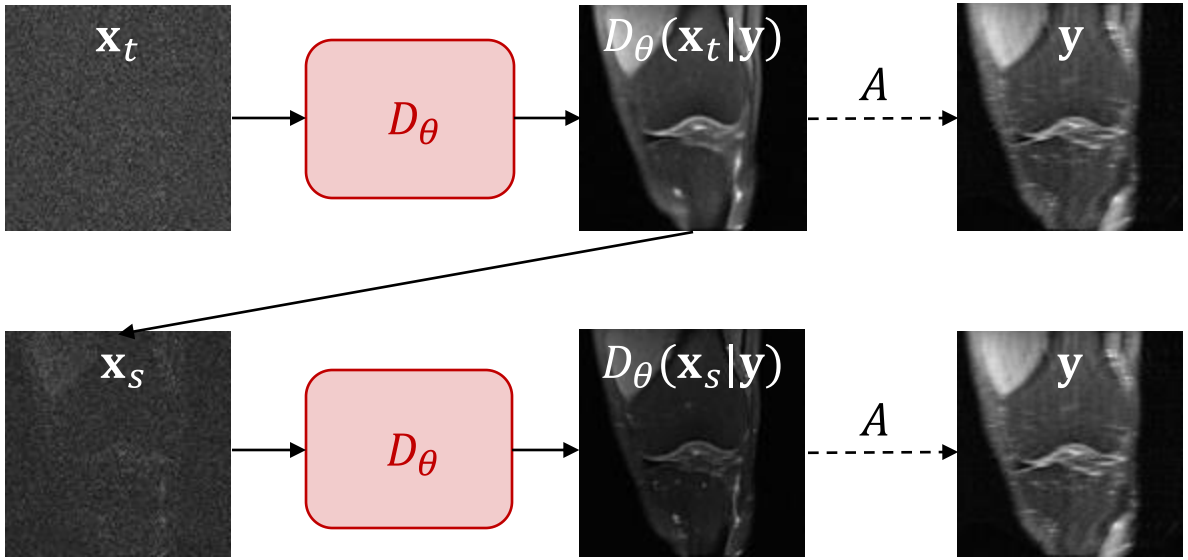

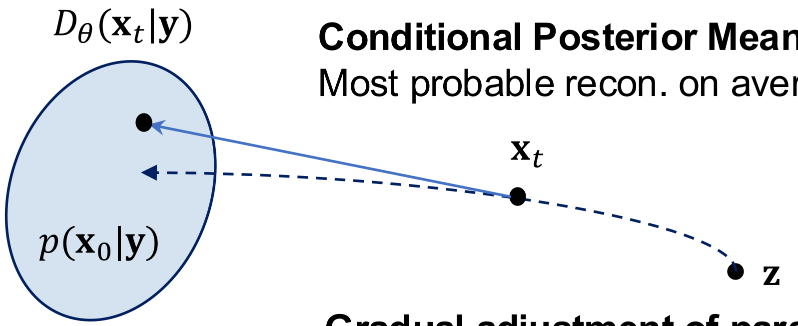

Conditional Posterior Mean:
Most probable recon. on average

Gradual adjustment of parameters:

- Fast optimization
- Constrained search space
- Multi-grid DIP with good initialization

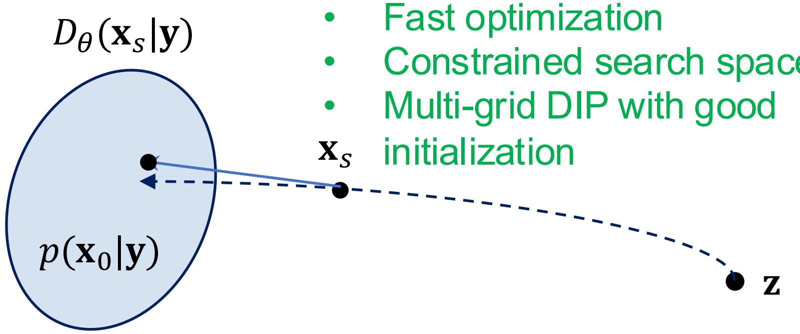

DDS approx.: $\mathbb{E}^{(\theta)}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}] \approx D_\theta(\mathbf{x}_t|\mathbf{y})$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

DDIP with LoRA

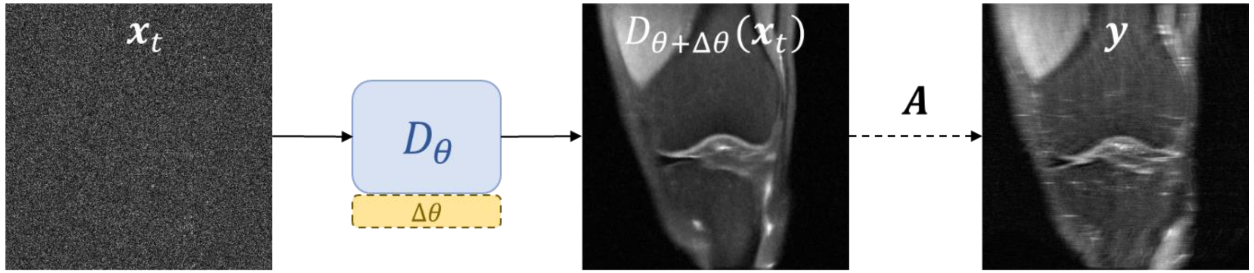

$$\min_{\Delta\theta} \|\boldsymbol{y} - \boldsymbol{A}D_{\theta+\Delta\theta}(\boldsymbol{x}_t|\boldsymbol{y})\|_2^2$$

0.6% new parameters

$$\text{for } t = T, \ldots, 1: \theta_{t-1} \leftarrow \arg\min_{\theta_t} \|\boldsymbol{y} - \boldsymbol{A}D_{\theta_t}(\boldsymbol{x}_t|\boldsymbol{y})\|_2^2,$$

$$\boldsymbol{x}_{t-1} \leftarrow \mathrm{DDIM}_{\theta_{t-1}}(D_{\theta_{t-1}}(\boldsymbol{x}_t|\boldsymbol{y}), \eta)$$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Results

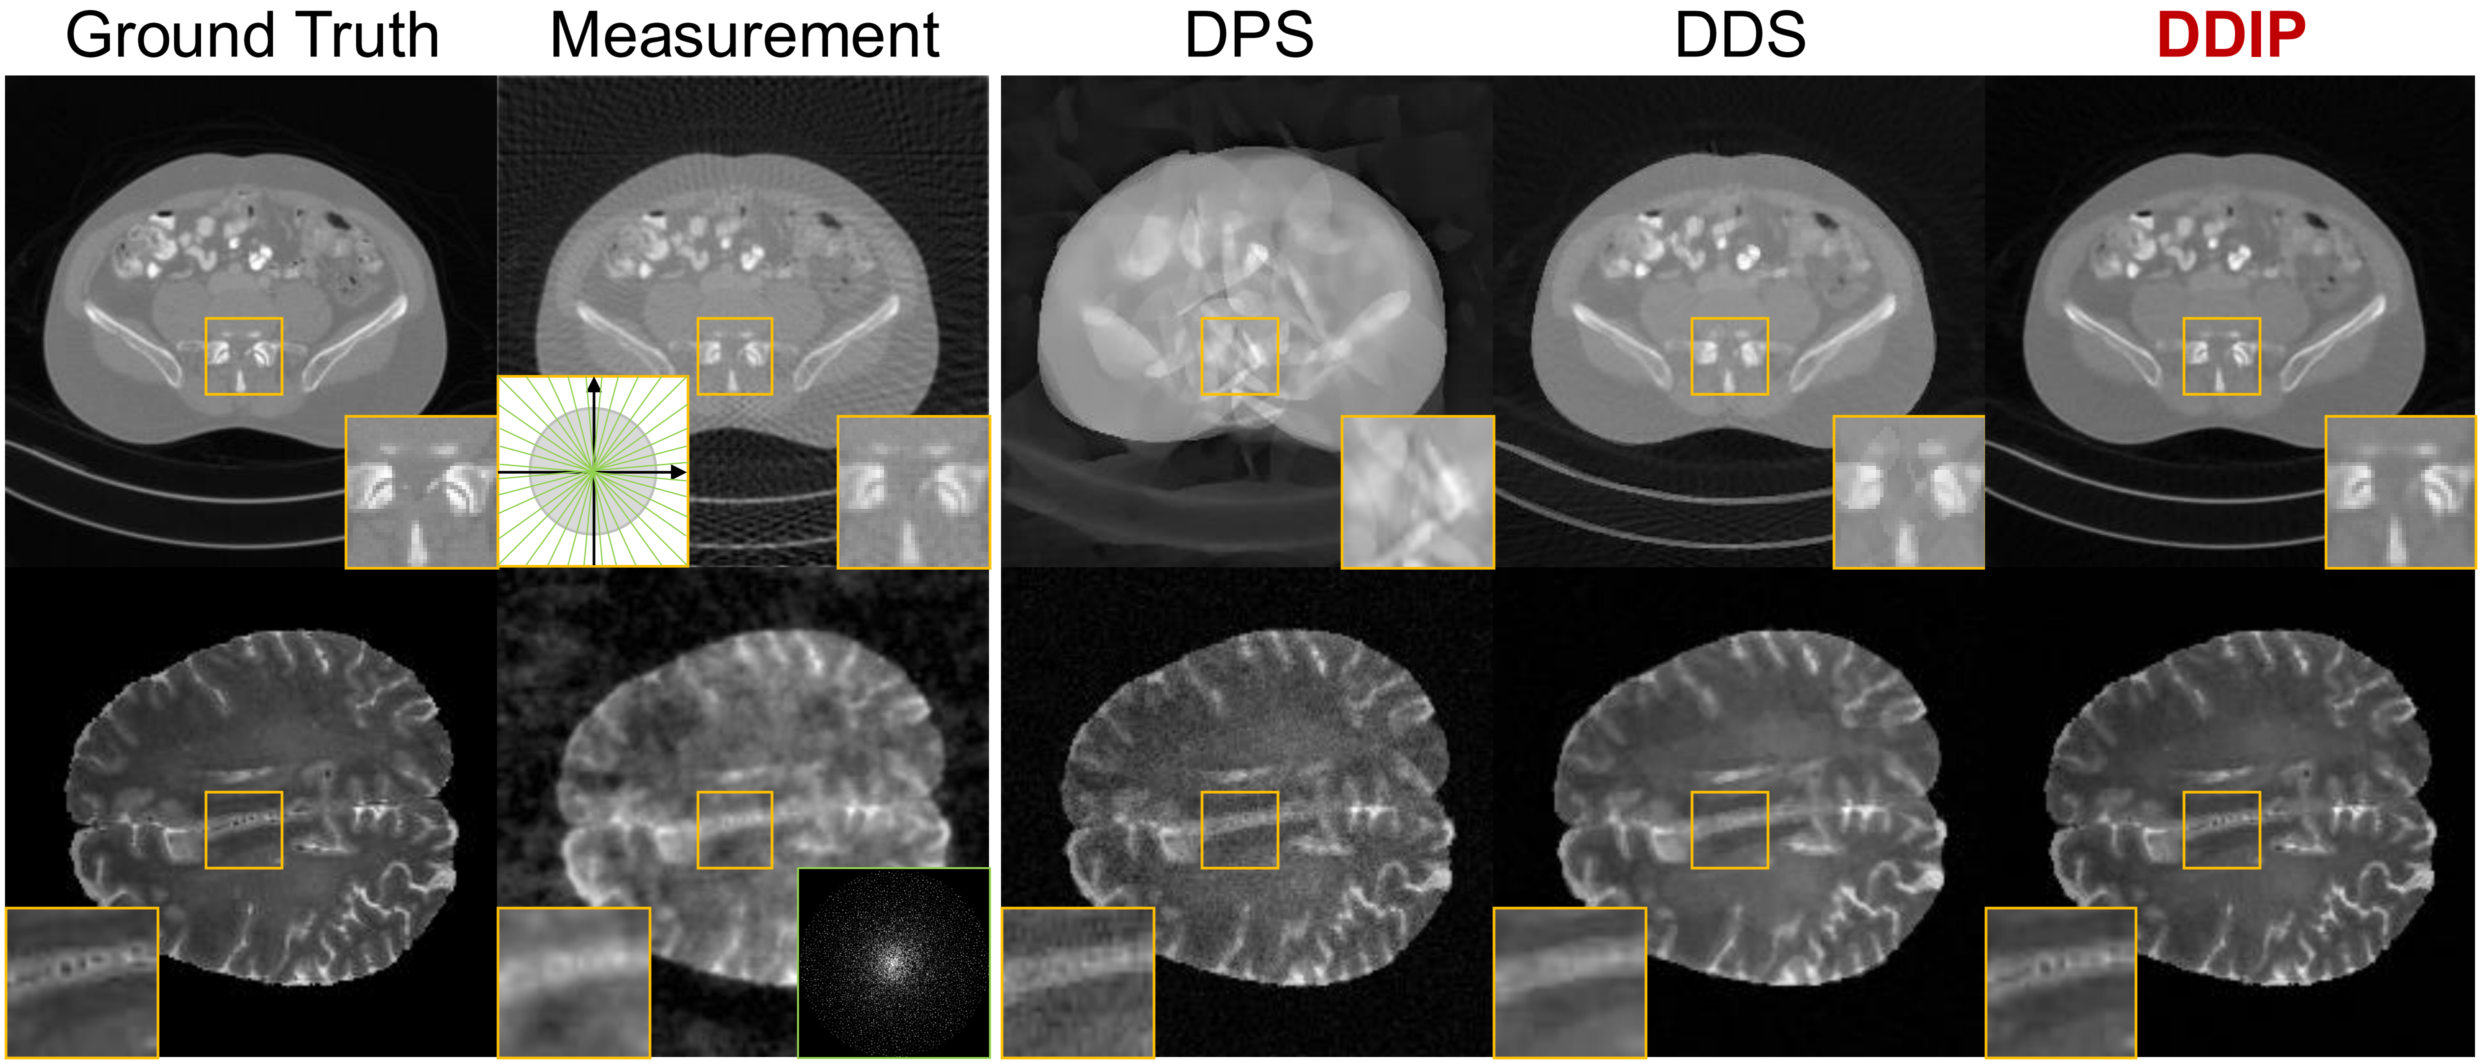

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Extending DDIP to D3IP (3D)

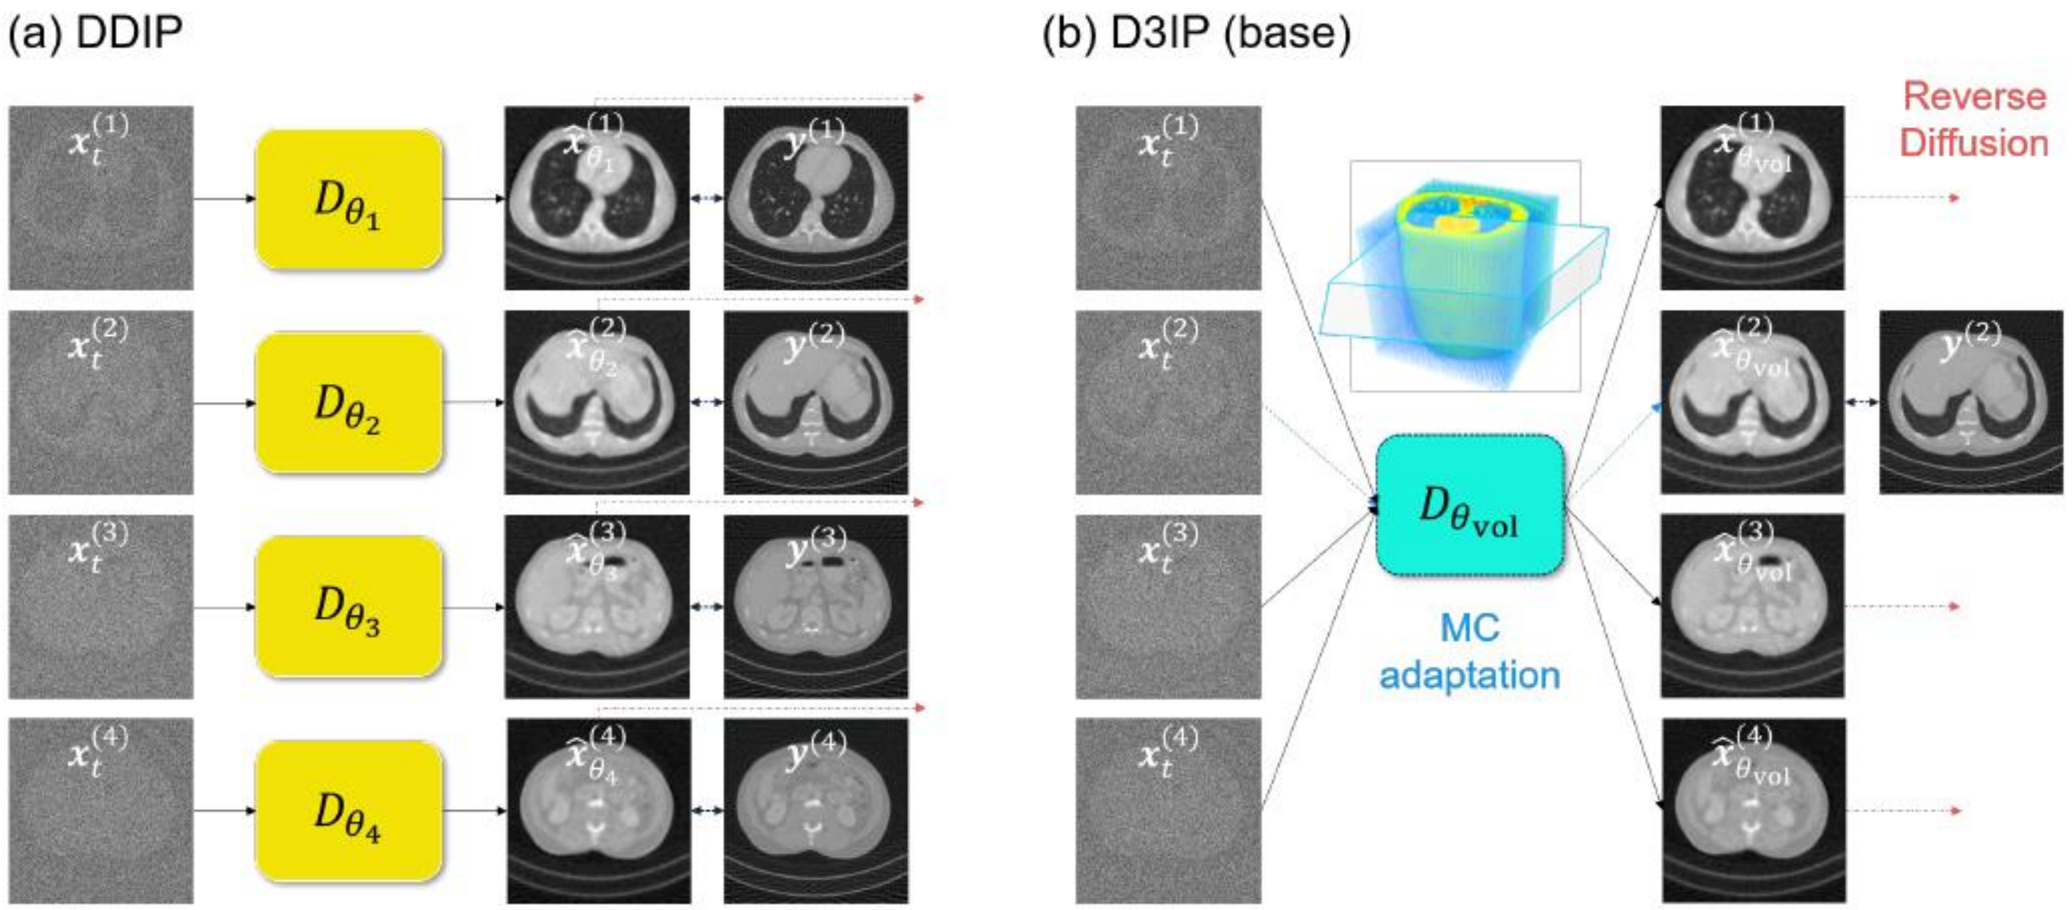

$$\text{for } t = T, \ldots, 1 : \min_{\theta} \mathbb{E}_i \left[\|\boldsymbol{y}^i - \boldsymbol{A} D_\theta^t(\boldsymbol{x}_t^i|\boldsymbol{y}^i)\|_2^2 \right] \approx \frac{1}{K} \sum_{i=1}^{K} \|\boldsymbol{y}^i - \boldsymbol{A} D_\theta^t(\boldsymbol{x}_t^i|\boldsymbol{y}^i)\|_2^2,$$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Extending DDIP to D3IP (3D)

Algorithm 1 DDIP

Require: $\theta, D^t, N, T', \eta, \boldsymbol{y}$

$\mathcal{L}(\boldsymbol{x}, \boldsymbol{y}, D_\theta) := \|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}|\boldsymbol{y})\|_2^2$

for $i = 1$ **to** N **do**

  $\theta_{T'}^{(i)} \leftarrow \theta$

  $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$

  $\boldsymbol{x}_T^{(i)} \leftarrow \sqrt{\bar{\alpha}_{T'}}\boldsymbol{A}^\dagger \boldsymbol{y} + \sqrt{1 - \bar{\alpha}_{T'}}\boldsymbol{\epsilon}$

  for $t = T'$ **to** 1 **do**

    $\theta_{t-1}^i \leftarrow \arg\min_{\theta_t^i} \mathcal{L}(\boldsymbol{x}_t^i, \boldsymbol{y}^i, D_{\theta_t^i}^t)$

    $\hat{\boldsymbol{x}}_{0|t}^i \leftarrow D_{\theta_{t-1}^i}^t(\boldsymbol{x}_t^i|\boldsymbol{y}^i)$

    $\boldsymbol{x}_{t-1}^i \leftarrow \mathrm{DDIM}(\hat{\boldsymbol{x}}_{0|t}^i, \eta)$

  end for

  return $\boldsymbol{x}_0^i$

end for

Algorithm 2 D3IP

Require: $\theta, D^t, N, T', K, \eta, \boldsymbol{Y}$

$\mathcal{L}(\boldsymbol{x}, \boldsymbol{y}, D_\theta) := \|\boldsymbol{y} - \boldsymbol{A}D_\theta(\boldsymbol{x}|\boldsymbol{y})\|_2^2$

$\theta_{T'} \leftarrow \theta$

$\boldsymbol{\epsilon}_1, \boldsymbol{\epsilon}_N \sim \mathcal{N}(0, \boldsymbol{I})$

$\boldsymbol{\epsilon} \leftarrow \mathrm{slerp}(\boldsymbol{\epsilon}_1, \boldsymbol{\epsilon}_N, \frac{i}{N})$

$\boldsymbol{X}_{T'} \leftarrow \sqrt{\bar{\alpha}_{T'}}\boldsymbol{A}^\dagger \boldsymbol{Y} + \sqrt{1 - \bar{\alpha}_{T'}}\boldsymbol{\epsilon}$

for $t = T'$ **to** 1 **do**

  $\boldsymbol{x}_t^{\{i\}}, \boldsymbol{y}^{\{i\}} \sim \mathrm{MC}((\boldsymbol{X}_t, \boldsymbol{Y}), K)$

  $\theta_{t-1} \leftarrow \arg\min_{\theta_t} \mathcal{L}(\boldsymbol{x}_t^{\{i\}}, \boldsymbol{y}^{\{i\}}, D_{\theta_t}^t)$

  $\hat{\boldsymbol{X}}_{0|t} \leftarrow D_{\theta_{t-1}}^t(\boldsymbol{X}_t|\boldsymbol{Y})$

  $\boldsymbol{X}_{t-1} \leftarrow \mathrm{DDIM}(\hat{\boldsymbol{X}}_{0|t}, \eta)$

end for

return $\boldsymbol{X}_0$

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

Results: 3D MRI

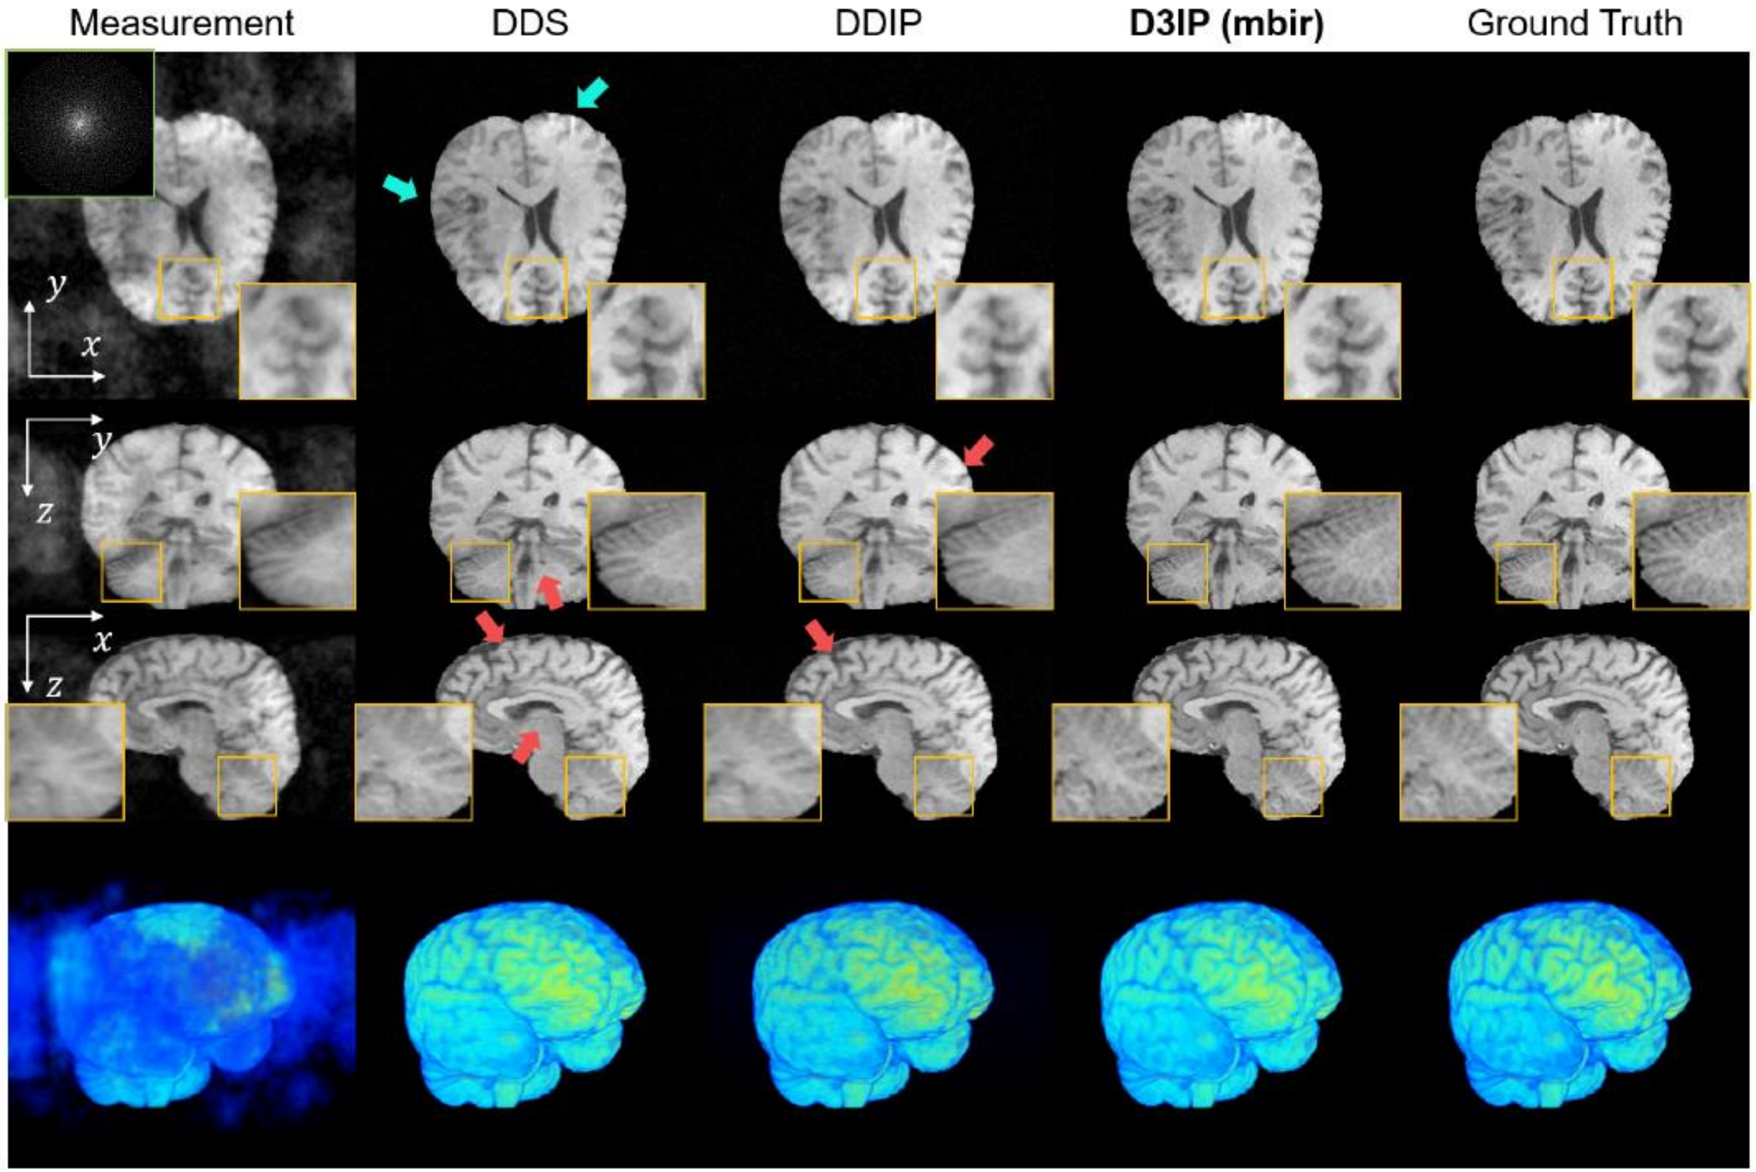

Meta Learning Further Enhances Results

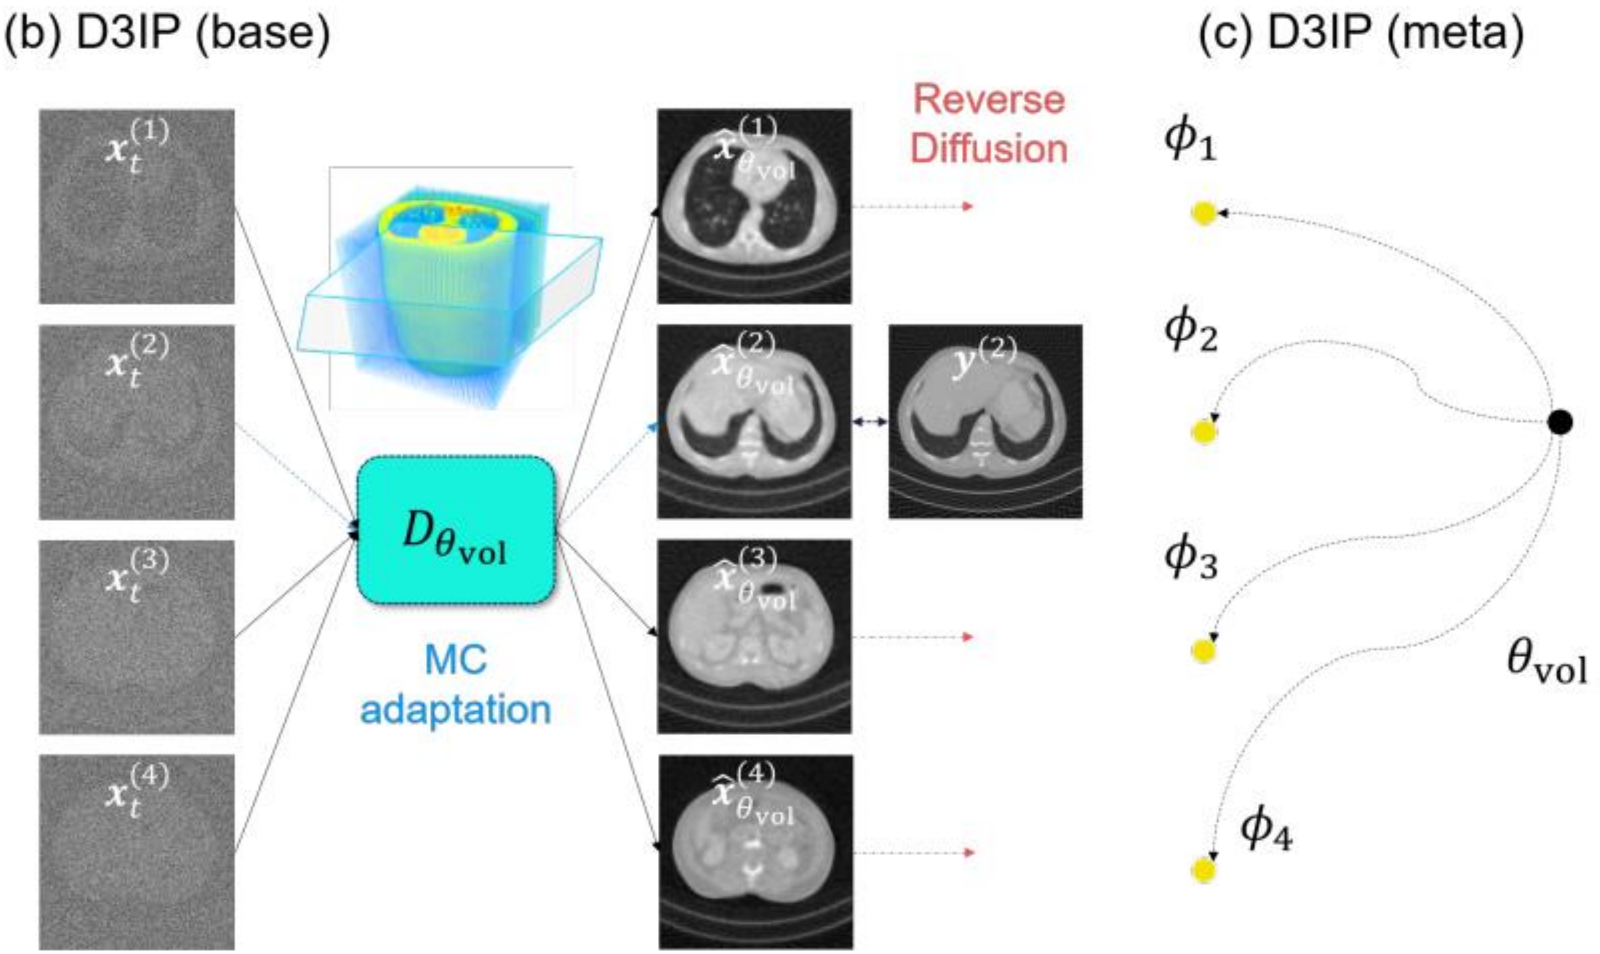

Chung & Ye, "*Deep Diffusion Image Prior for Efficient OOD Adaptation in 3D Inverse Problems*", ECCV 2024

3D Super-resolution from OOD Measurements

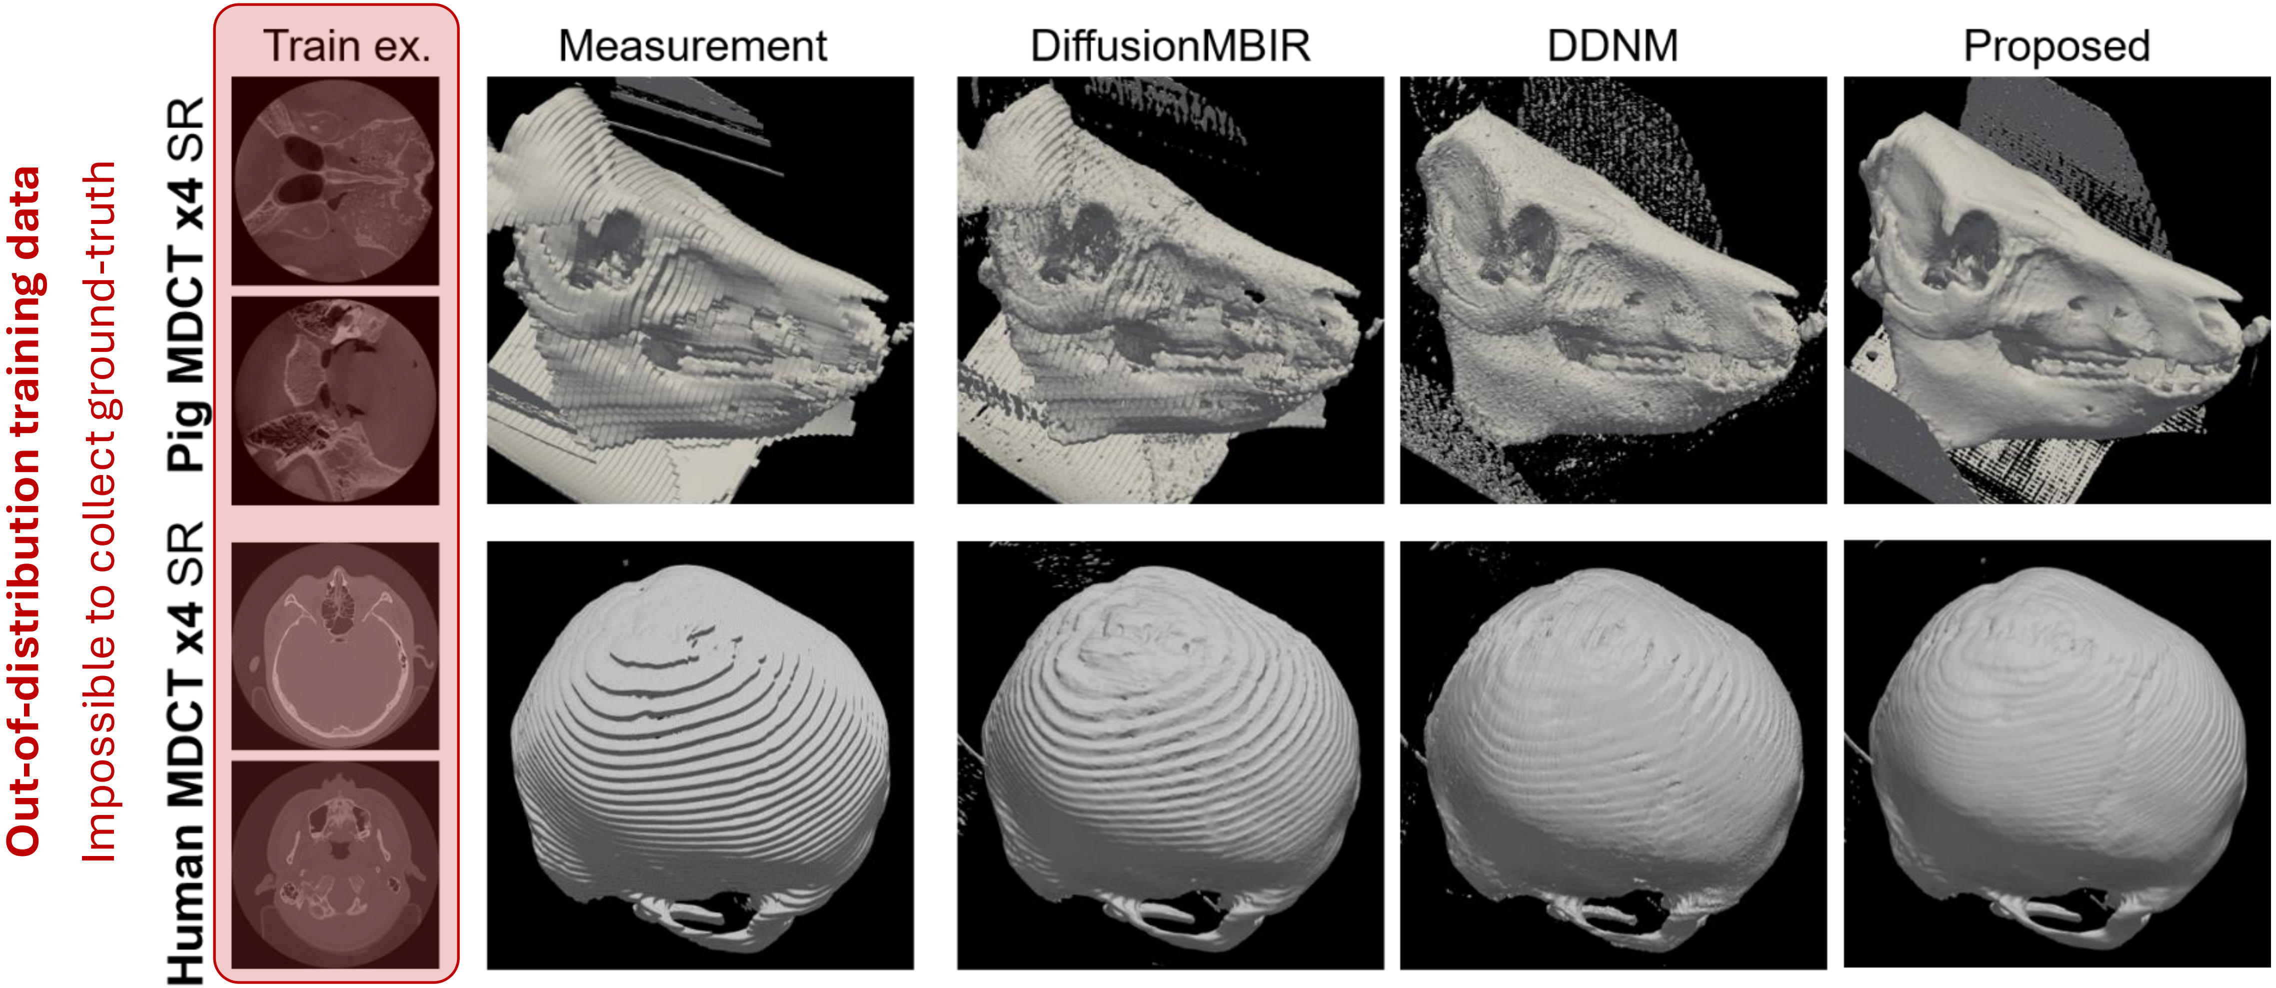

3D Super-resolution from OOD Measurements

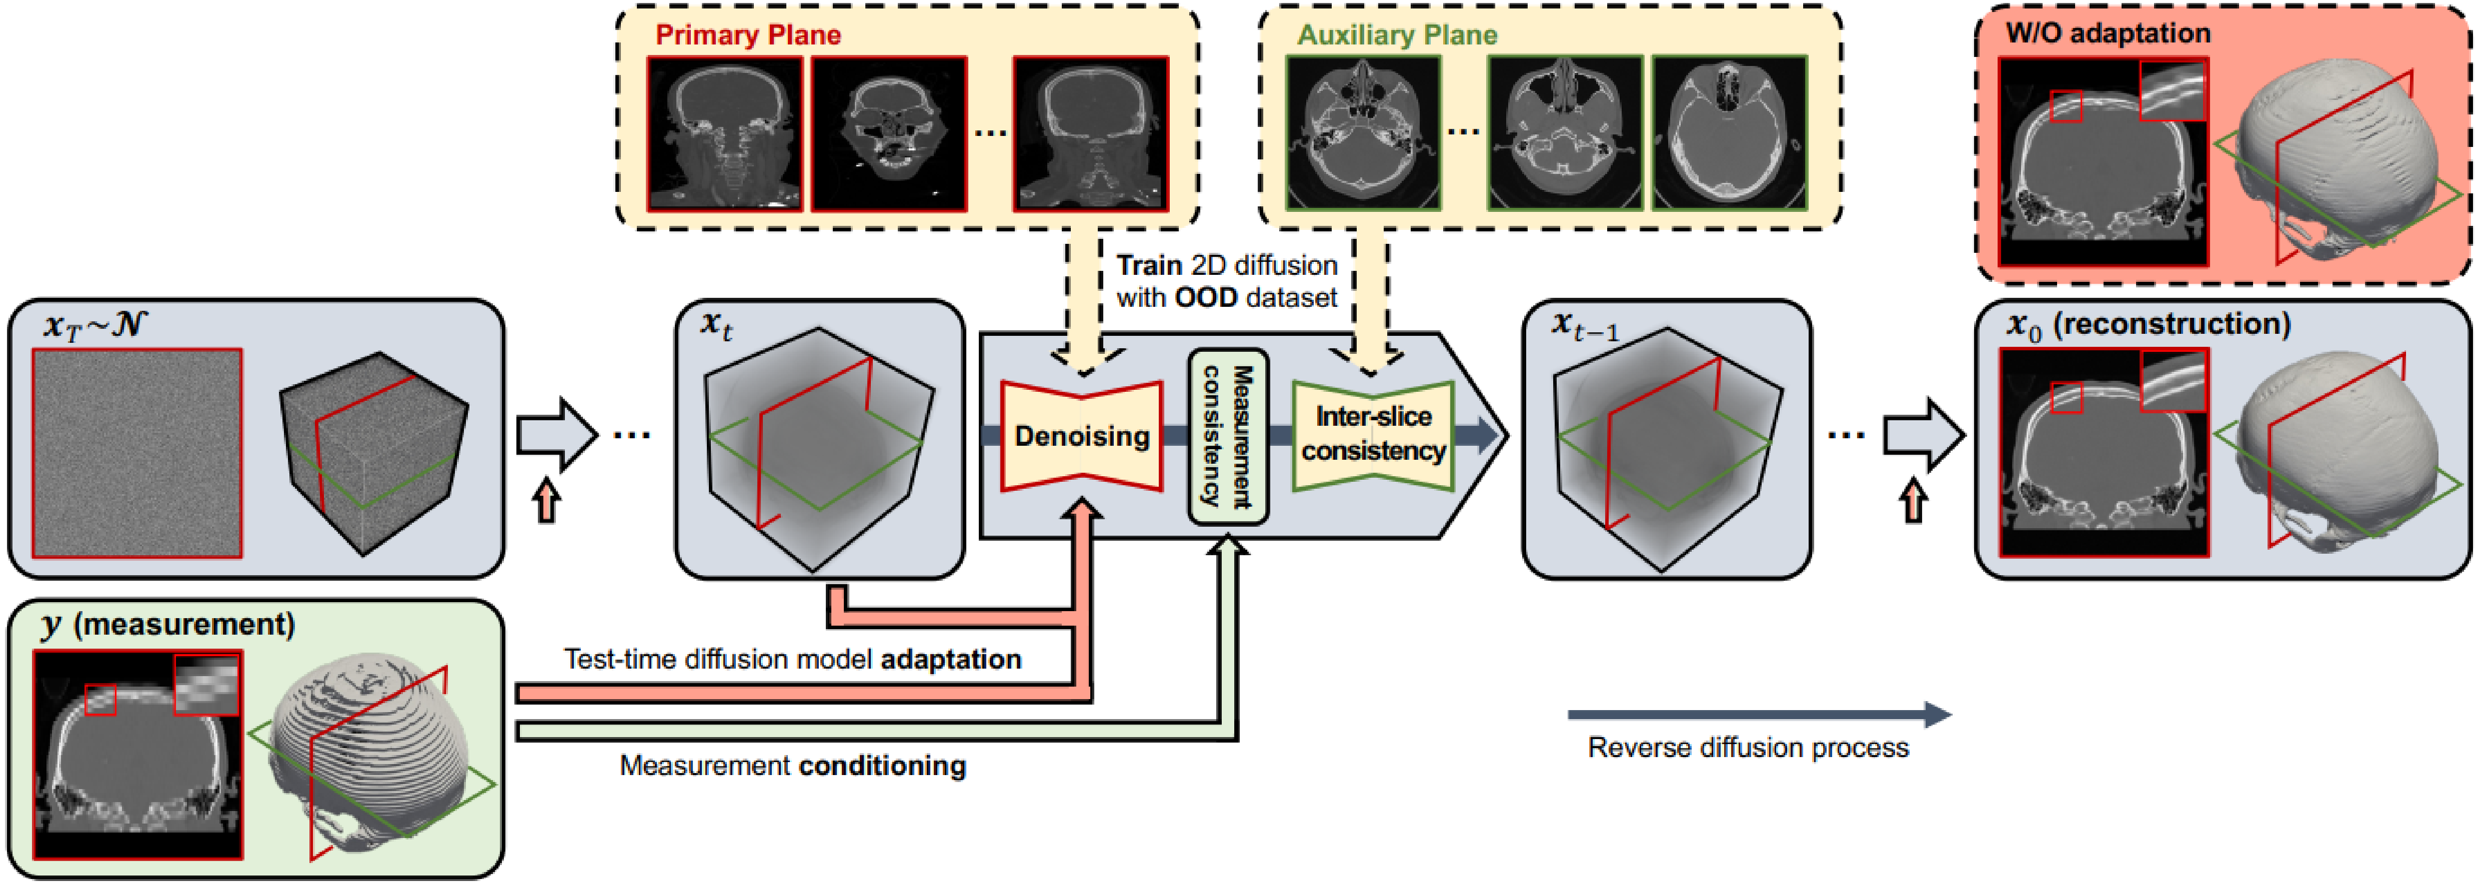

Results

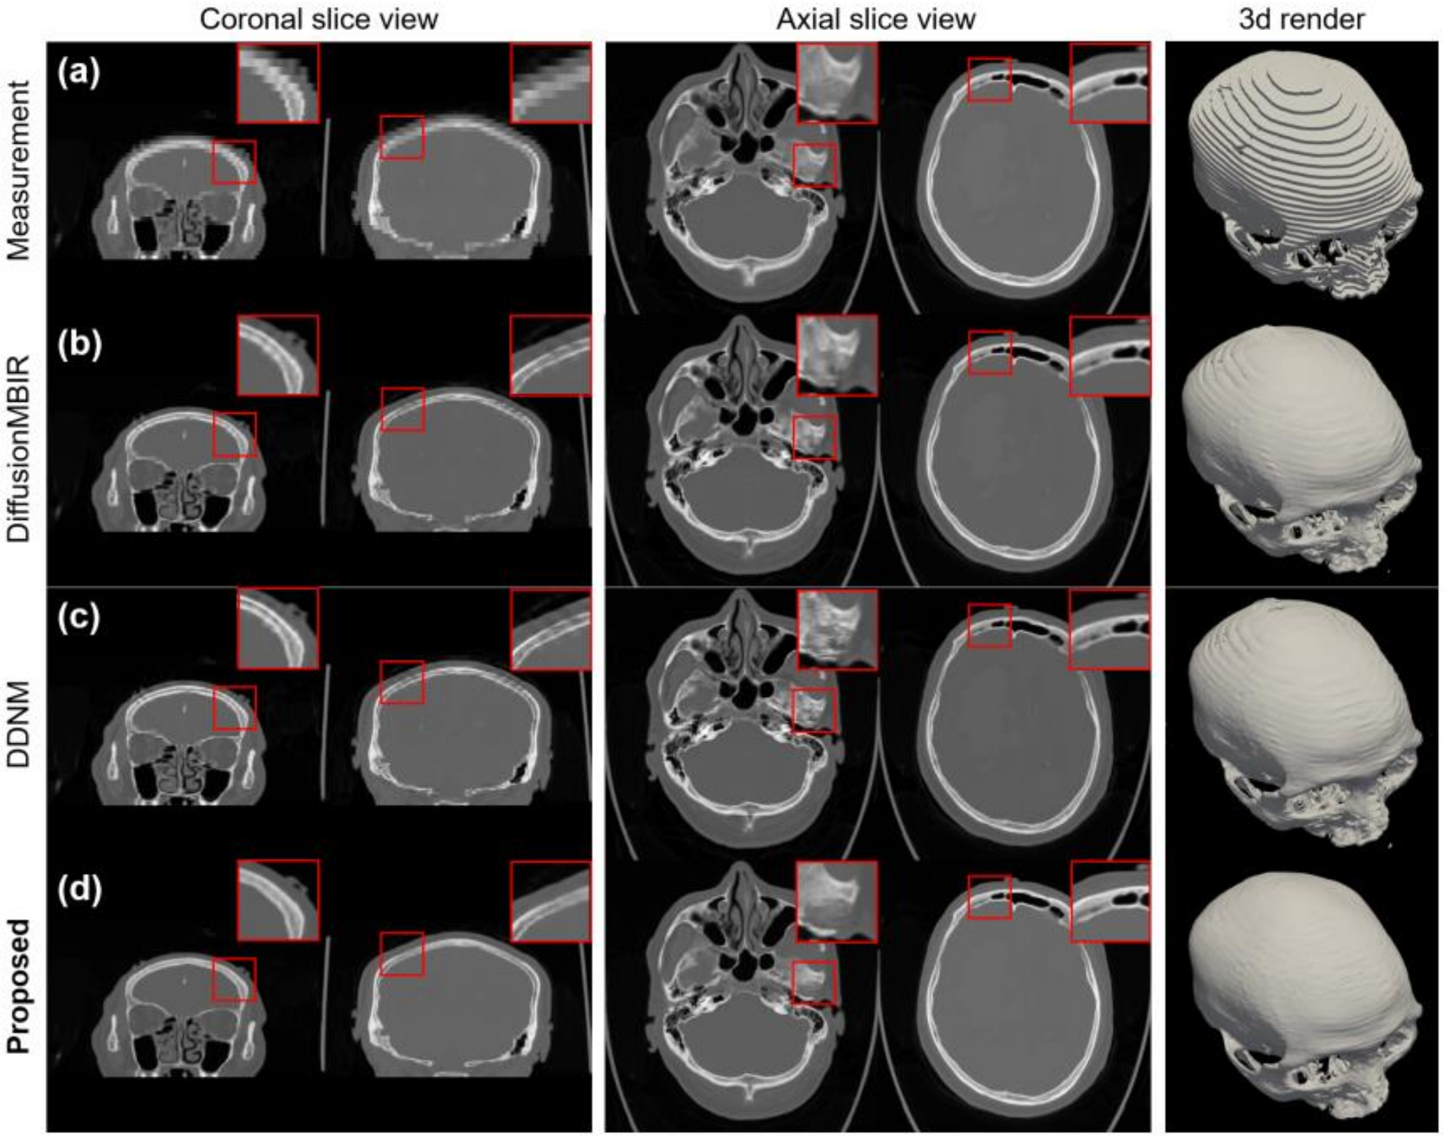

Coronal slice view
Axial slice view
3d render
(a)
Measurement
(b)
DiffusionMBIR
(c)
DDNM
(d)
Proposed

Results

Method	×4 SR				×8 SR			
	PSNR ↑	SSIM(cor) ↑	SSIM(axi) ↑	SSIM(sag) ↑	PSNR ↑	SSIM(cor) ↑	SSIM(axi) ↑	SSIM(sag) ↑
Upsample (nearest)	31.43	0.9048	0.9152	0.9298	27.45	0.8180	0.8299	0.8732
Upsample (bilinear)	33.08	0.9177	0.9248	0.9298	28.59	0.8455	0.8475	0.8692
U-Net [29]	32.62	0.9931	0.9451	0.9412	28.99	0.8647	0.8788	0.8847
SRGAN [30]	30.84	0.8111	0.8275	0.8729	27.62	0.7956	0.8129	0.8503
SRGAN (2.5D) [30]	30.26	0.8999	0.9141	0.9189	28.53	0.8130	0.8321	0.8565
DDNM [28]	31.22	0.8159	0.8412	0.8146	28.15	0.7397	0.7944	0.7369
DiffusionMBIR (DDS) [8]	35.03	0.9196	0.9447	0.9262	31.06	0.8628	0.9106	0.8712
D3IP (base) [13]	31.26	0.8166	0.8419	0.8154	28.20	0.7405	0.7754	0.7378
D3IP (mbir) [13]	35.10	0.9196	0.9445	0.9262	31.11	0.8626	0.9104	0.8711
Ours	**38.91**	**0.9572**	**0.9641**	**0.9585**	**35.31**	**0.9221**	**0.9343**	**0.9232**
Ours-ID (oracle)	40.90	0.9662	0.9716	0.9674	36.22	0.9344	0.9450	0.9356

Training Diffusion Models with Measurements

Goal: Train $p_{\text{data}}(\mathbf{x}) \approx p_\theta(\mathbf{x})$ with measurements $\mathbf{y} \sim p(\mathbf{y}|\mathbf{x})$

SURE

- Train DM with Generalized SURE-like loss
- Kawar et al. TMLR 2024

Ambient diffusion

- Further corrupt, then try to reconstruct
- For noise levels below measurement, regularize with consistency
- Daras et al. NeurIPS 2024, ICML 2024

EM

- Iterate between training DM with current posterior samples and sampling from the posterior
- Rozet et al. NeurIPS 2024

Equivariance

- Regularize with rotation/translation/normalization equivariance
- Can be used even with single A
- Levac et al. arxiv 2025

Q & A