Principled Probabilistic Imaging using Diffusion Models as Plug-and-Play Priors

NeurIPS 2024 Poster Presentation

Zihui (Ray) Wu
Computational Cameras Group, Caltech

We aim to design a rigorous Bayesian **posterior sampler** using **diffusion models** for solving **inverse problems**

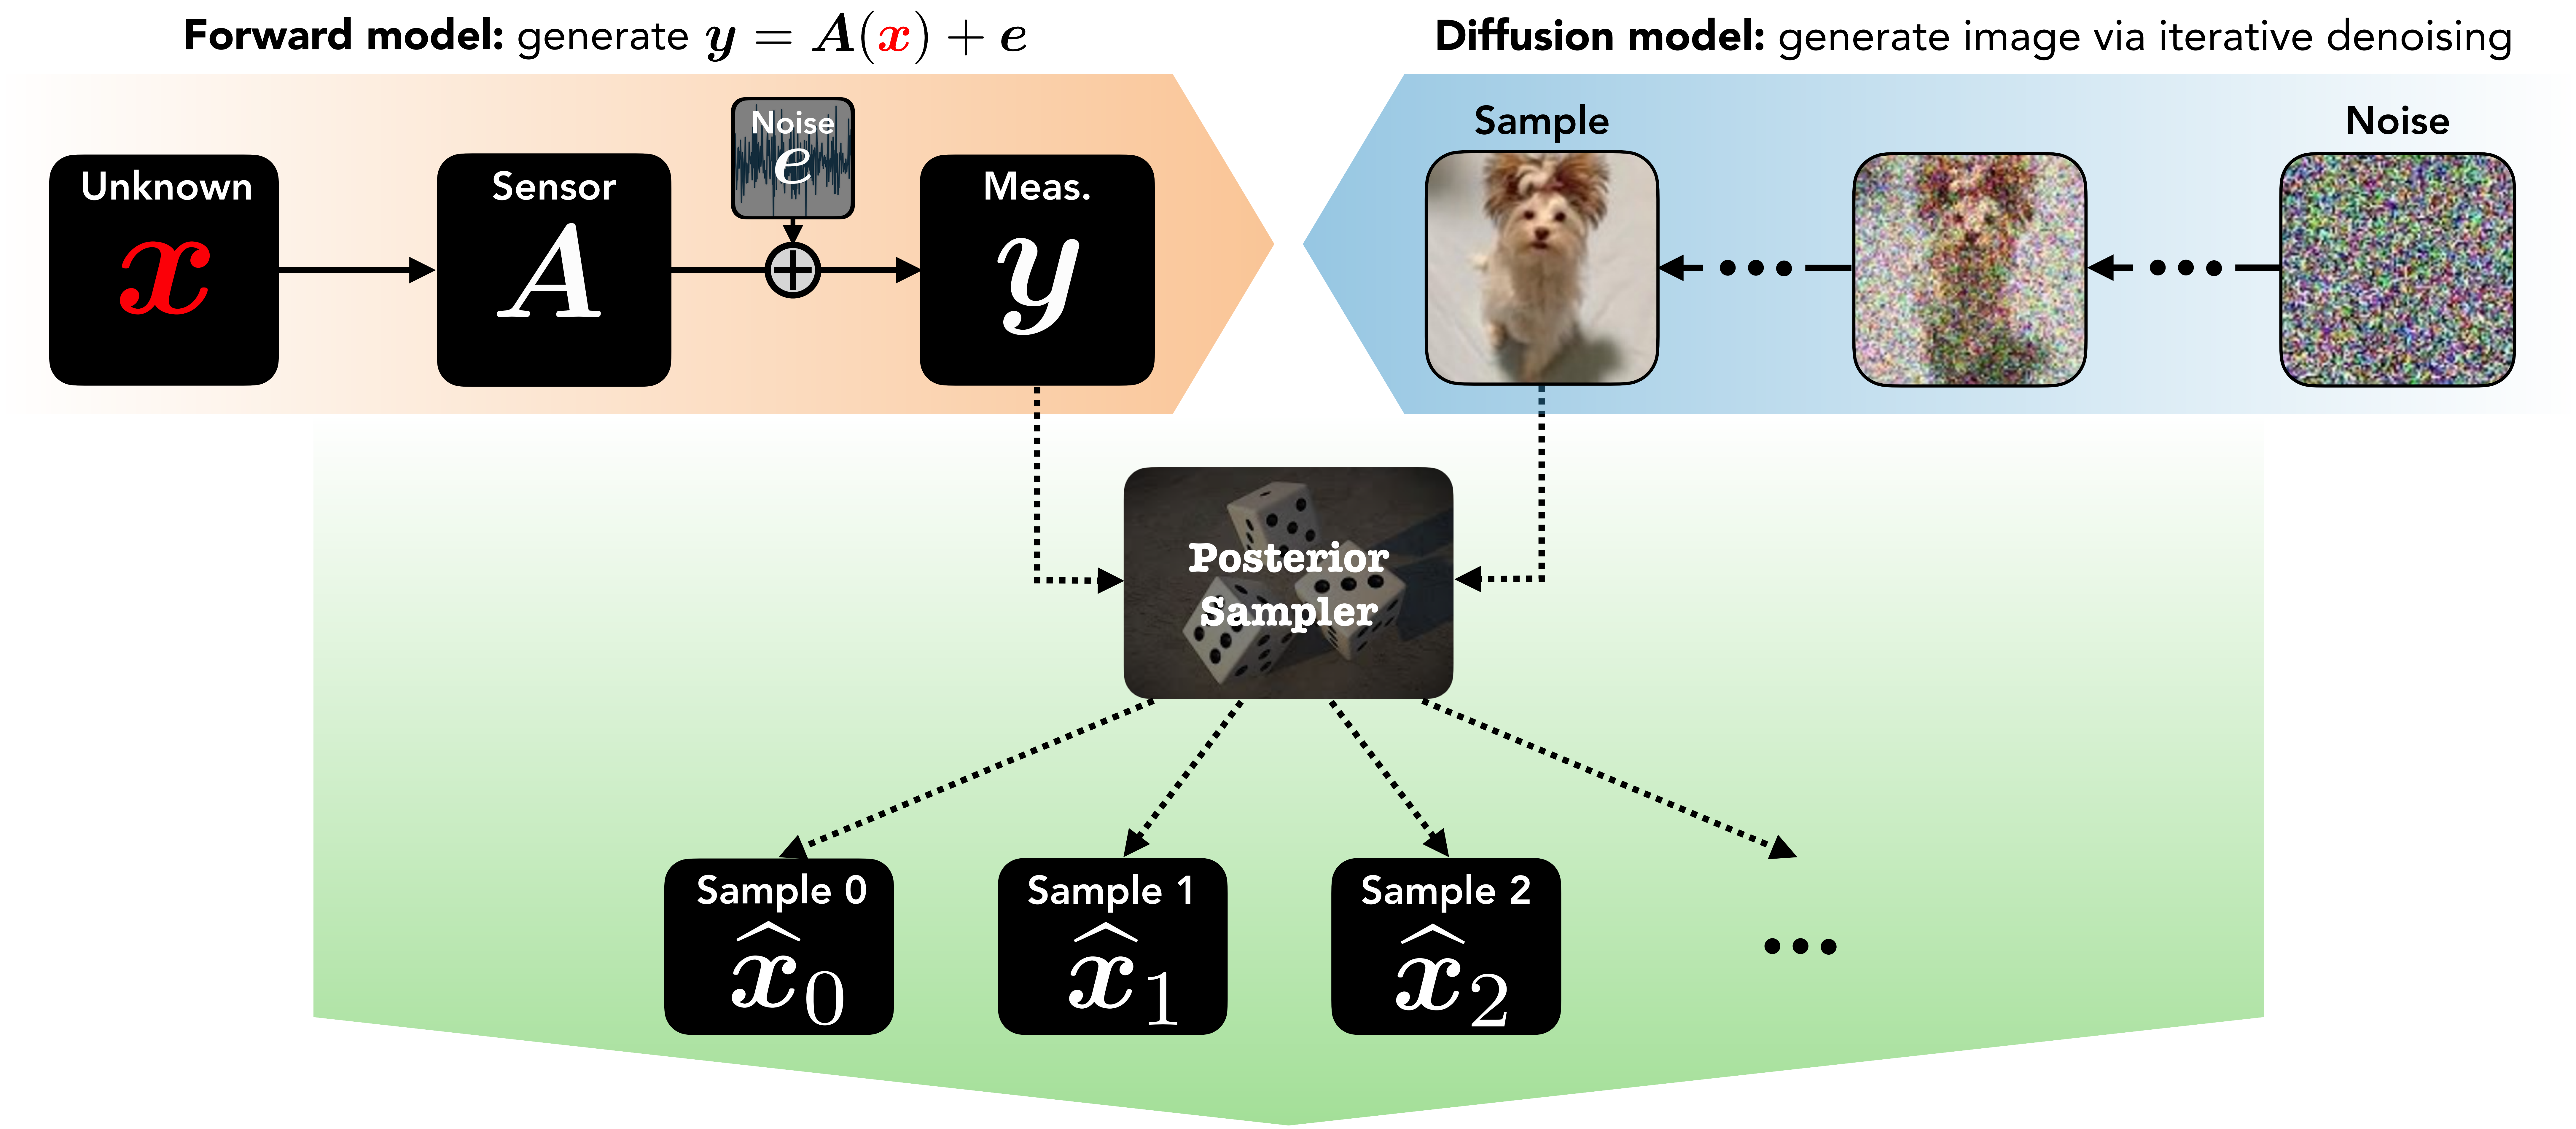

Bayesian approach: sample from the posterior $p(\boldsymbol{x}|\boldsymbol{y})$

We adopt the **Split Gibbs Sampler** formulation to sample from the posterior distribution

Posterior **Likelihood** **Prior**

$$p(\boldsymbol{x}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{x}) \; p(\boldsymbol{x})$$

Augmented distribution

$$\overset{\rho \to 0}{\approx} \underbrace{p(\boldsymbol{y}|\boldsymbol{z}) \;\; p(\boldsymbol{x})}_{\text{Variable splitting}} \; \underbrace{\exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}\|_2^2\right)}_{\text{Coupling term}}$$

Augmented variable

Split Gibbs Sampler (SGS) [1]

1. **Likelihood step:** fix $\boldsymbol{x}$ and sample $\boldsymbol{z}$: $\quad \boldsymbol{z}^{(k)} \sim p(\boldsymbol{y}|\boldsymbol{z}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x}^{(k)} - \boldsymbol{z}\|_2^2\right)$

2. **Prior step:** fix $\boldsymbol{z}$ and sample $\boldsymbol{x}$: $\quad \boldsymbol{x}^{(k+1)} \sim p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right)$

[1]: Vono et al., Split-and-augmented Gibbs sampler - Application to large-scale inference problems (2018).

We identify a key connection between **the prior step** of Split Gibbs Sampler and **the EDM framework**

1. Likelihood step: fix $\boldsymbol{x}$ and sample $\boldsymbol{z}$:

$$\boldsymbol{z}^{(k)} \sim p(\boldsymbol{y}|\boldsymbol{z}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x}^{(k)} - \boldsymbol{z}\|_2^2\right)$$

2. Prior step: fix $\boldsymbol{z}$ and sample $\boldsymbol{x}$:

$$\boldsymbol{x}^{(k+1)} \sim p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right)$$

Observation

$$p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right) \propto \underbrace{p(\boldsymbol{x})}_{\text{Prior}} \underbrace{\mathcal{N}(\boldsymbol{x}; \boldsymbol{z}^{(k)}, \rho^2 \boldsymbol{I})}_{\text{Likelihood of Gaussian denoising}}$$

Noise level = 0

$$\mathrm{d}\boldsymbol{x}_t = -\dot{\sigma}(t)\sigma(t)\nabla \log p(\boldsymbol{x}_t; \sigma(t))\mathrm{d}t$$

Noise level = ρ

EDM reverse diffusion [2]

[2]: Karras et al., Elucidating the Design Space of Diffusion-Based Generative Models (2022).

We propose **PnP-DM** as a principled method to leverage **diffusion models** for solving imaging inverse problems

1. Likelihood step: fix $\boldsymbol{x}$ and sample $\boldsymbol{z}$:

$$\boldsymbol{z}^{(k)} \sim p(\boldsymbol{y}|\boldsymbol{z}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x}^{(k)} - \boldsymbol{z}\|_2^2\right)$$

2. Prior step: fix $\boldsymbol{z}$ and sample $\boldsymbol{x}$:

$$\boldsymbol{x}^{(k+1)} \sim p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right)$$

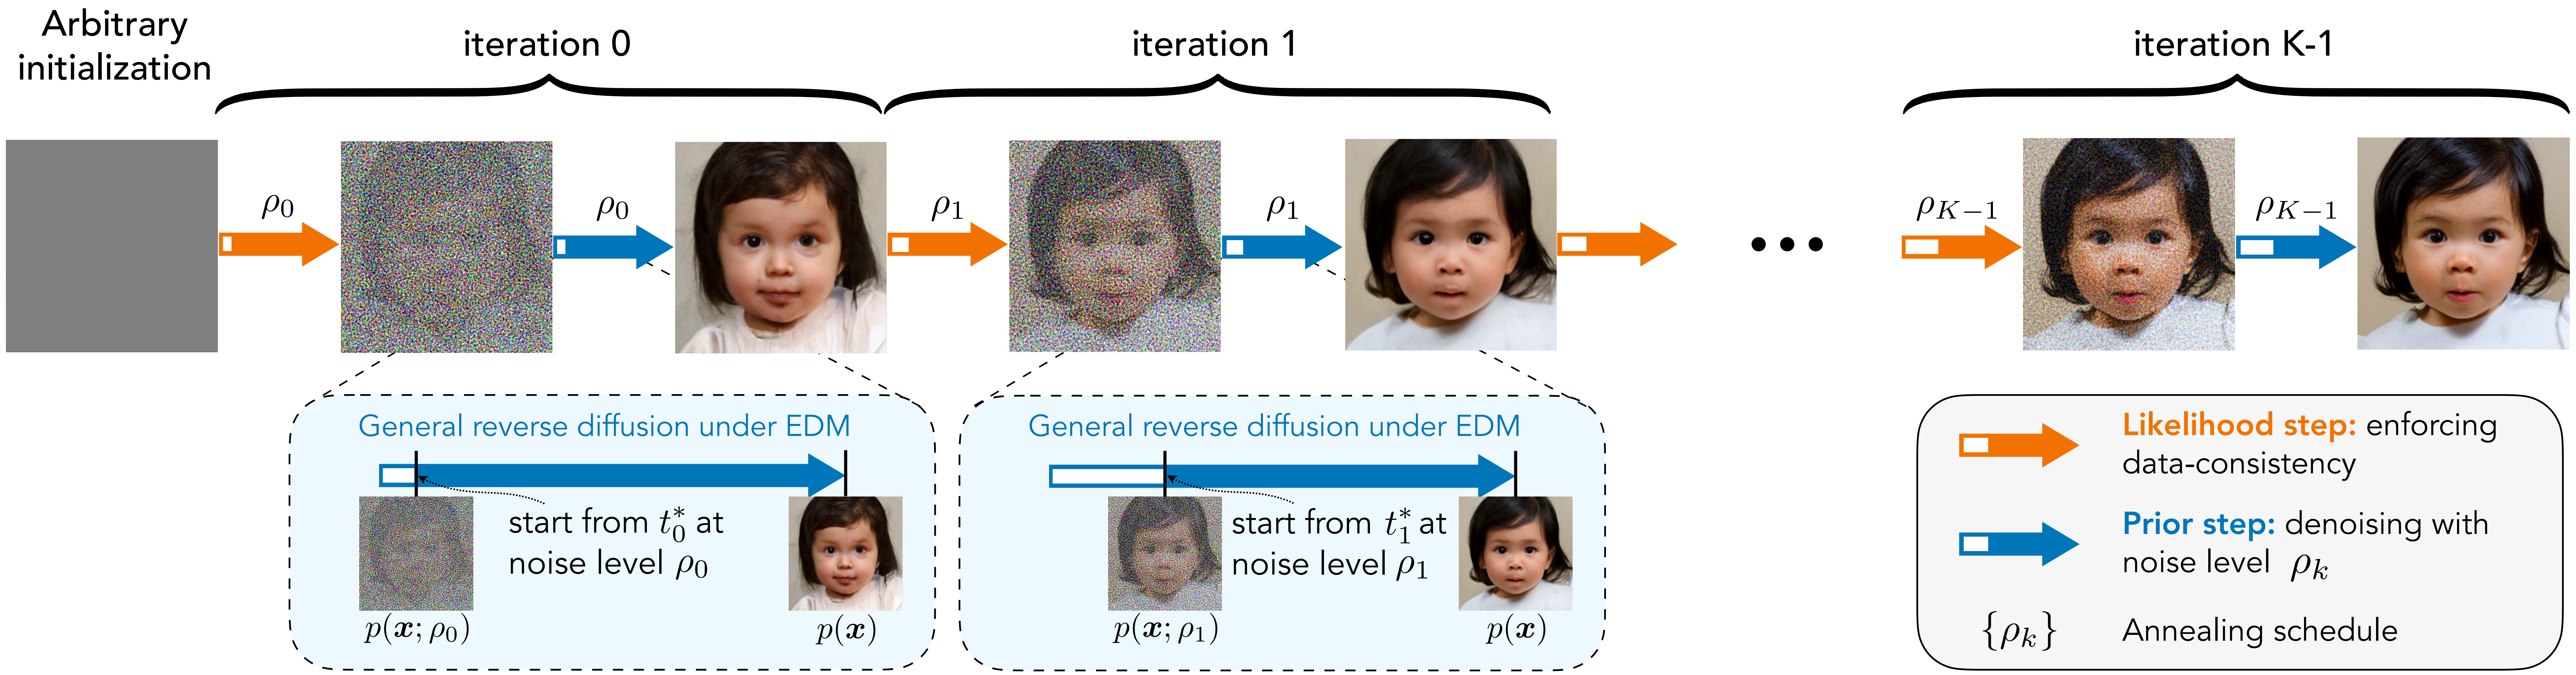

We propose **PnP-DM** as a principled method to leverage **diffusion models** for solving imaging inverse problems

1. Likelihood step: fix $\boldsymbol{x}$ and sample $\boldsymbol{z}$:

$$\boldsymbol{z}^{(k)} \sim p(\boldsymbol{y}|\boldsymbol{z}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x}^{(k)} - \boldsymbol{z}\|_2^2\right)$$

$\boldsymbol{z}$ $\boldsymbol{x}$

$$\boldsymbol{x}^{(k+1)} \sim p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right)$$

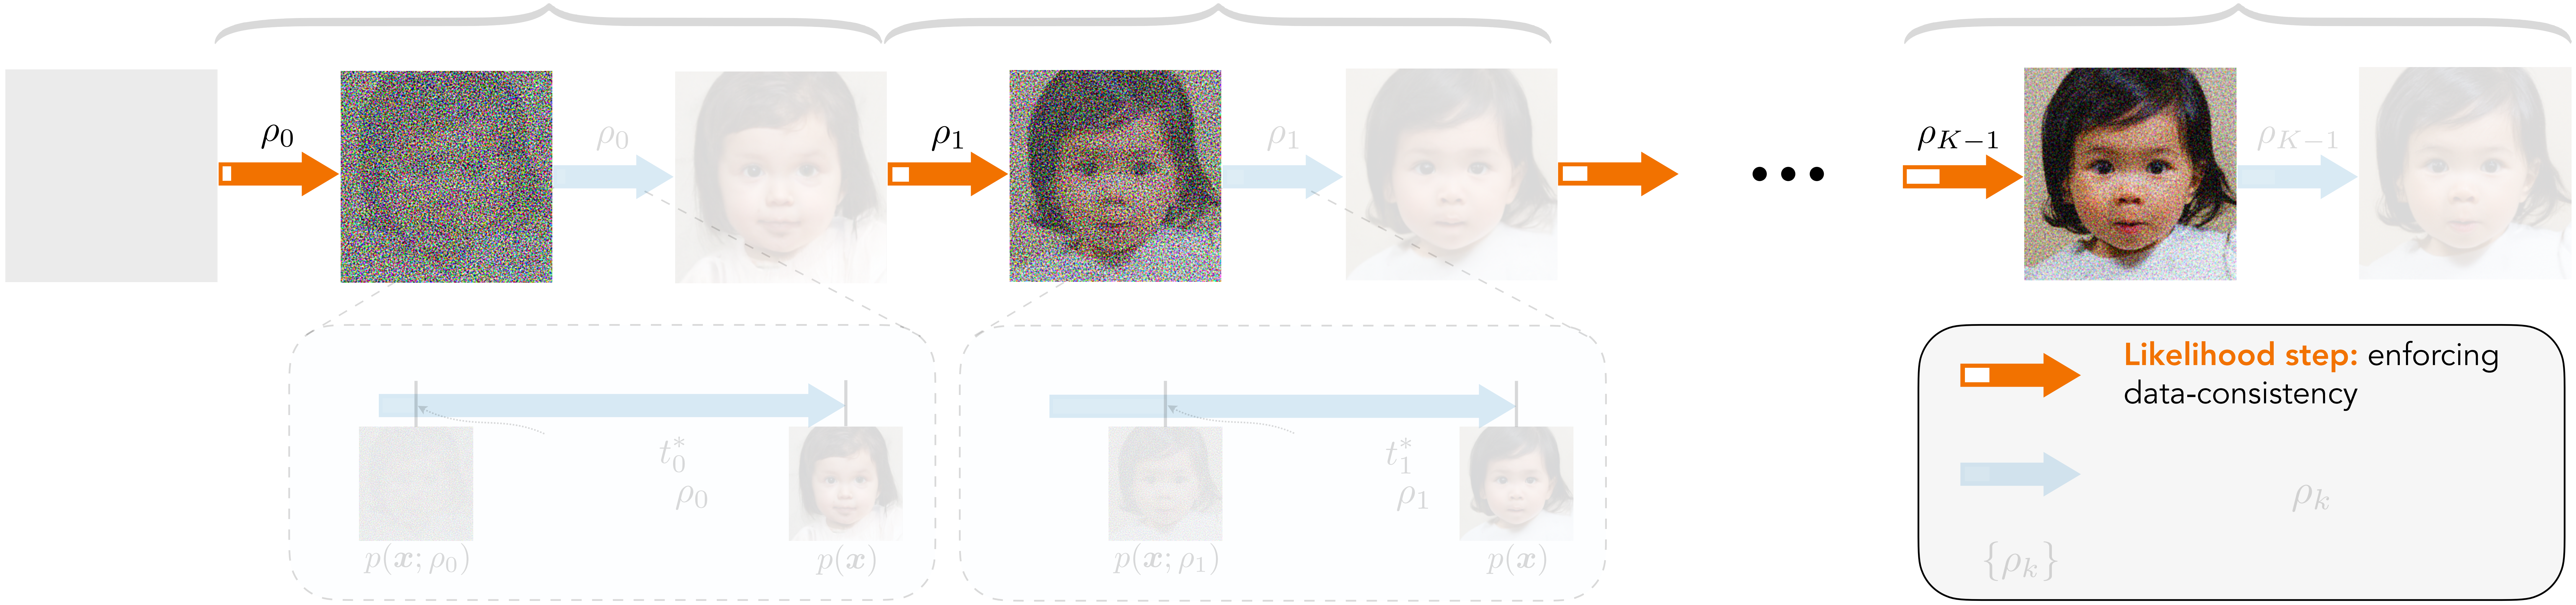

We propose **PnP-DM** as a principled method to leverage **diffusion models** for solving imaging inverse problems

$\boldsymbol{x}$ $\boldsymbol{z}$ $\boldsymbol{z}^{(k)} \sim p(\boldsymbol{y}|\boldsymbol{z}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x}^{(k)} - \boldsymbol{z}\|_2^2\right)$

2. Prior step: fix $\boldsymbol{z}$ and sample $\boldsymbol{x}$:

$$\boldsymbol{x}^{(k+1)} \sim p(\boldsymbol{x}) \exp\left(-\frac{1}{2\rho^2}\|\boldsymbol{x} - \boldsymbol{z}^{(k)}\|_2^2\right)$$

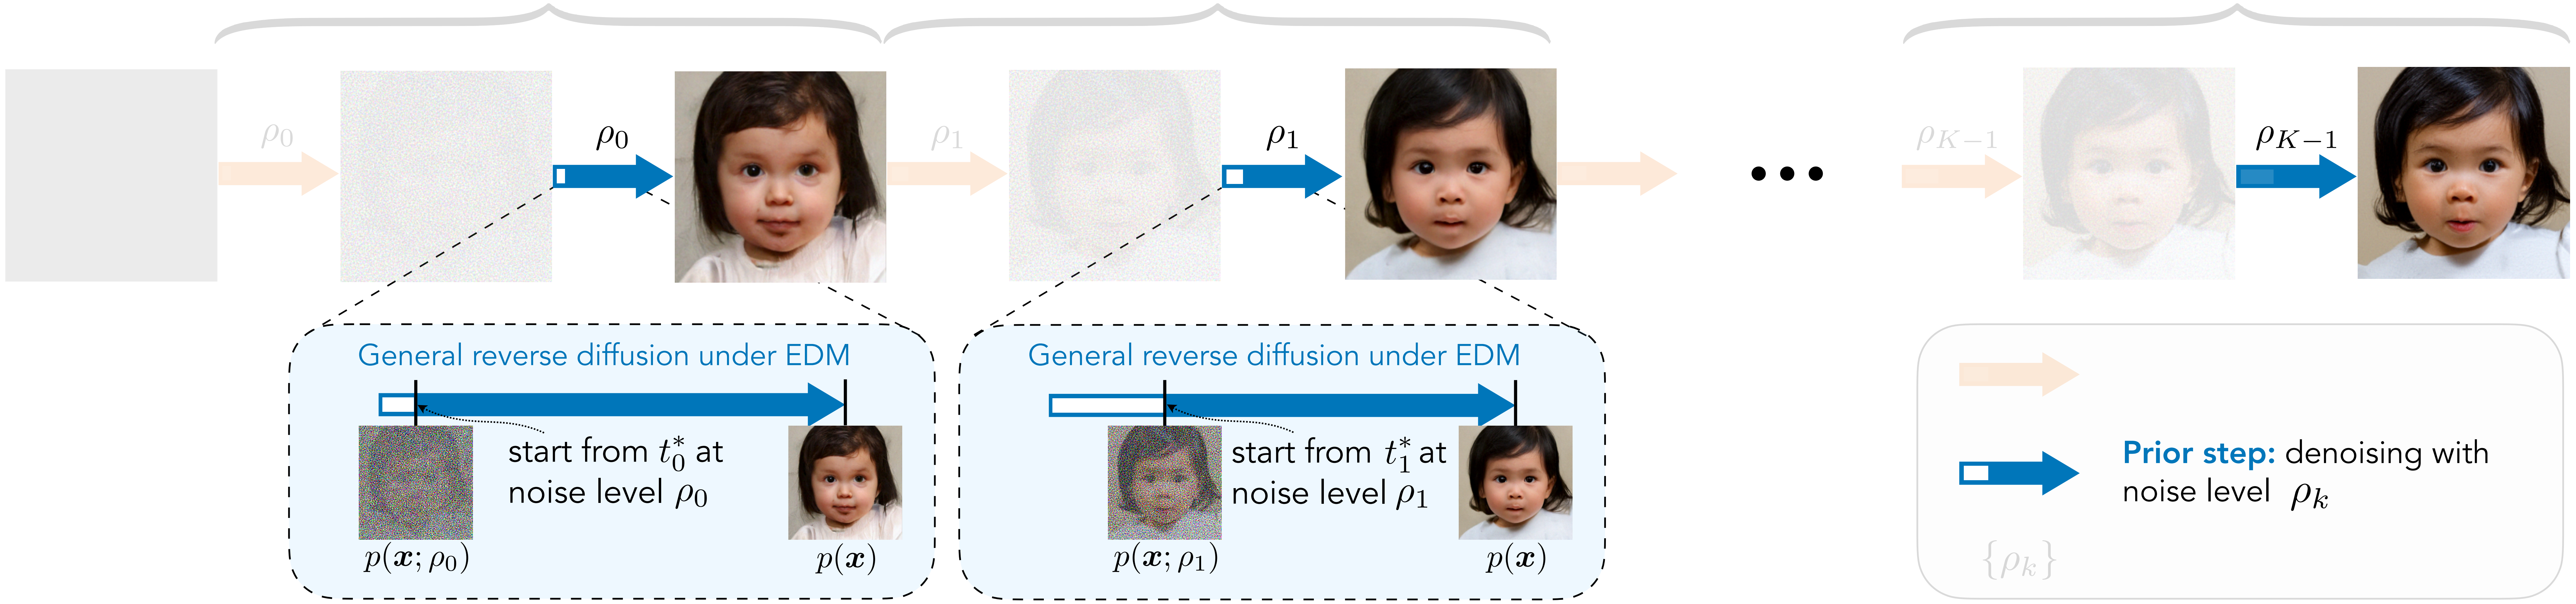

PnP-DM outperforms existing methods on **image restoration problems**

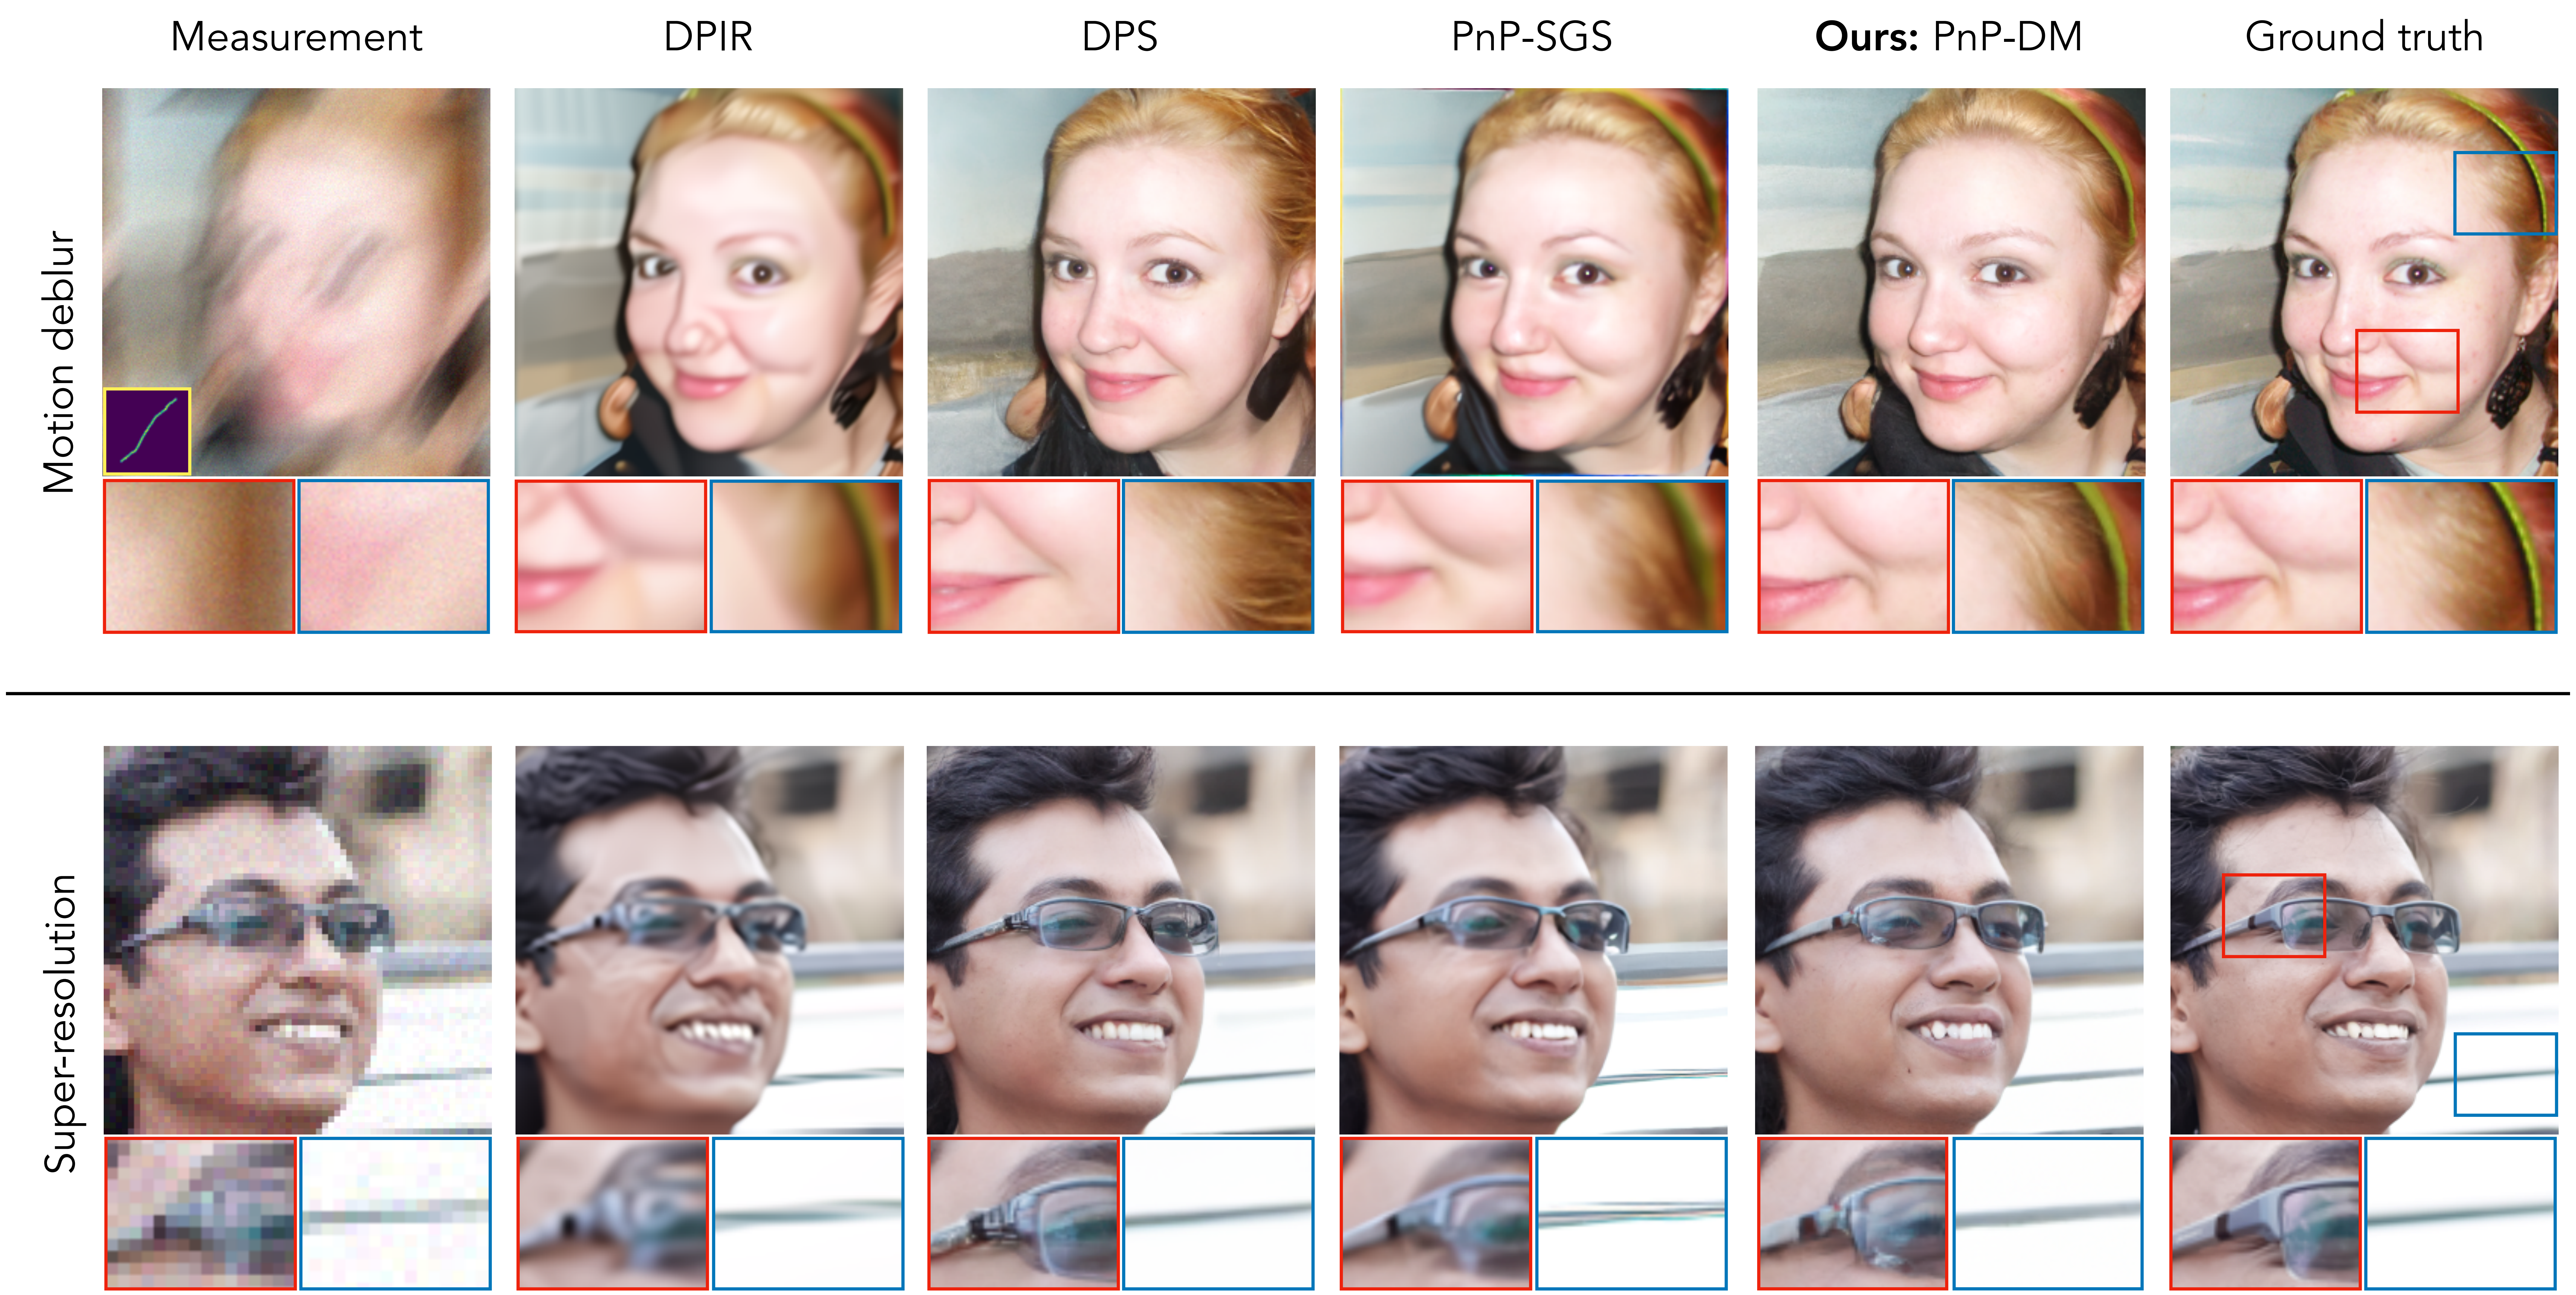

PnP-DM successfully recovers **the M87 black hole** using **the real measurement data** from Event Horizon Telescope (EHT)

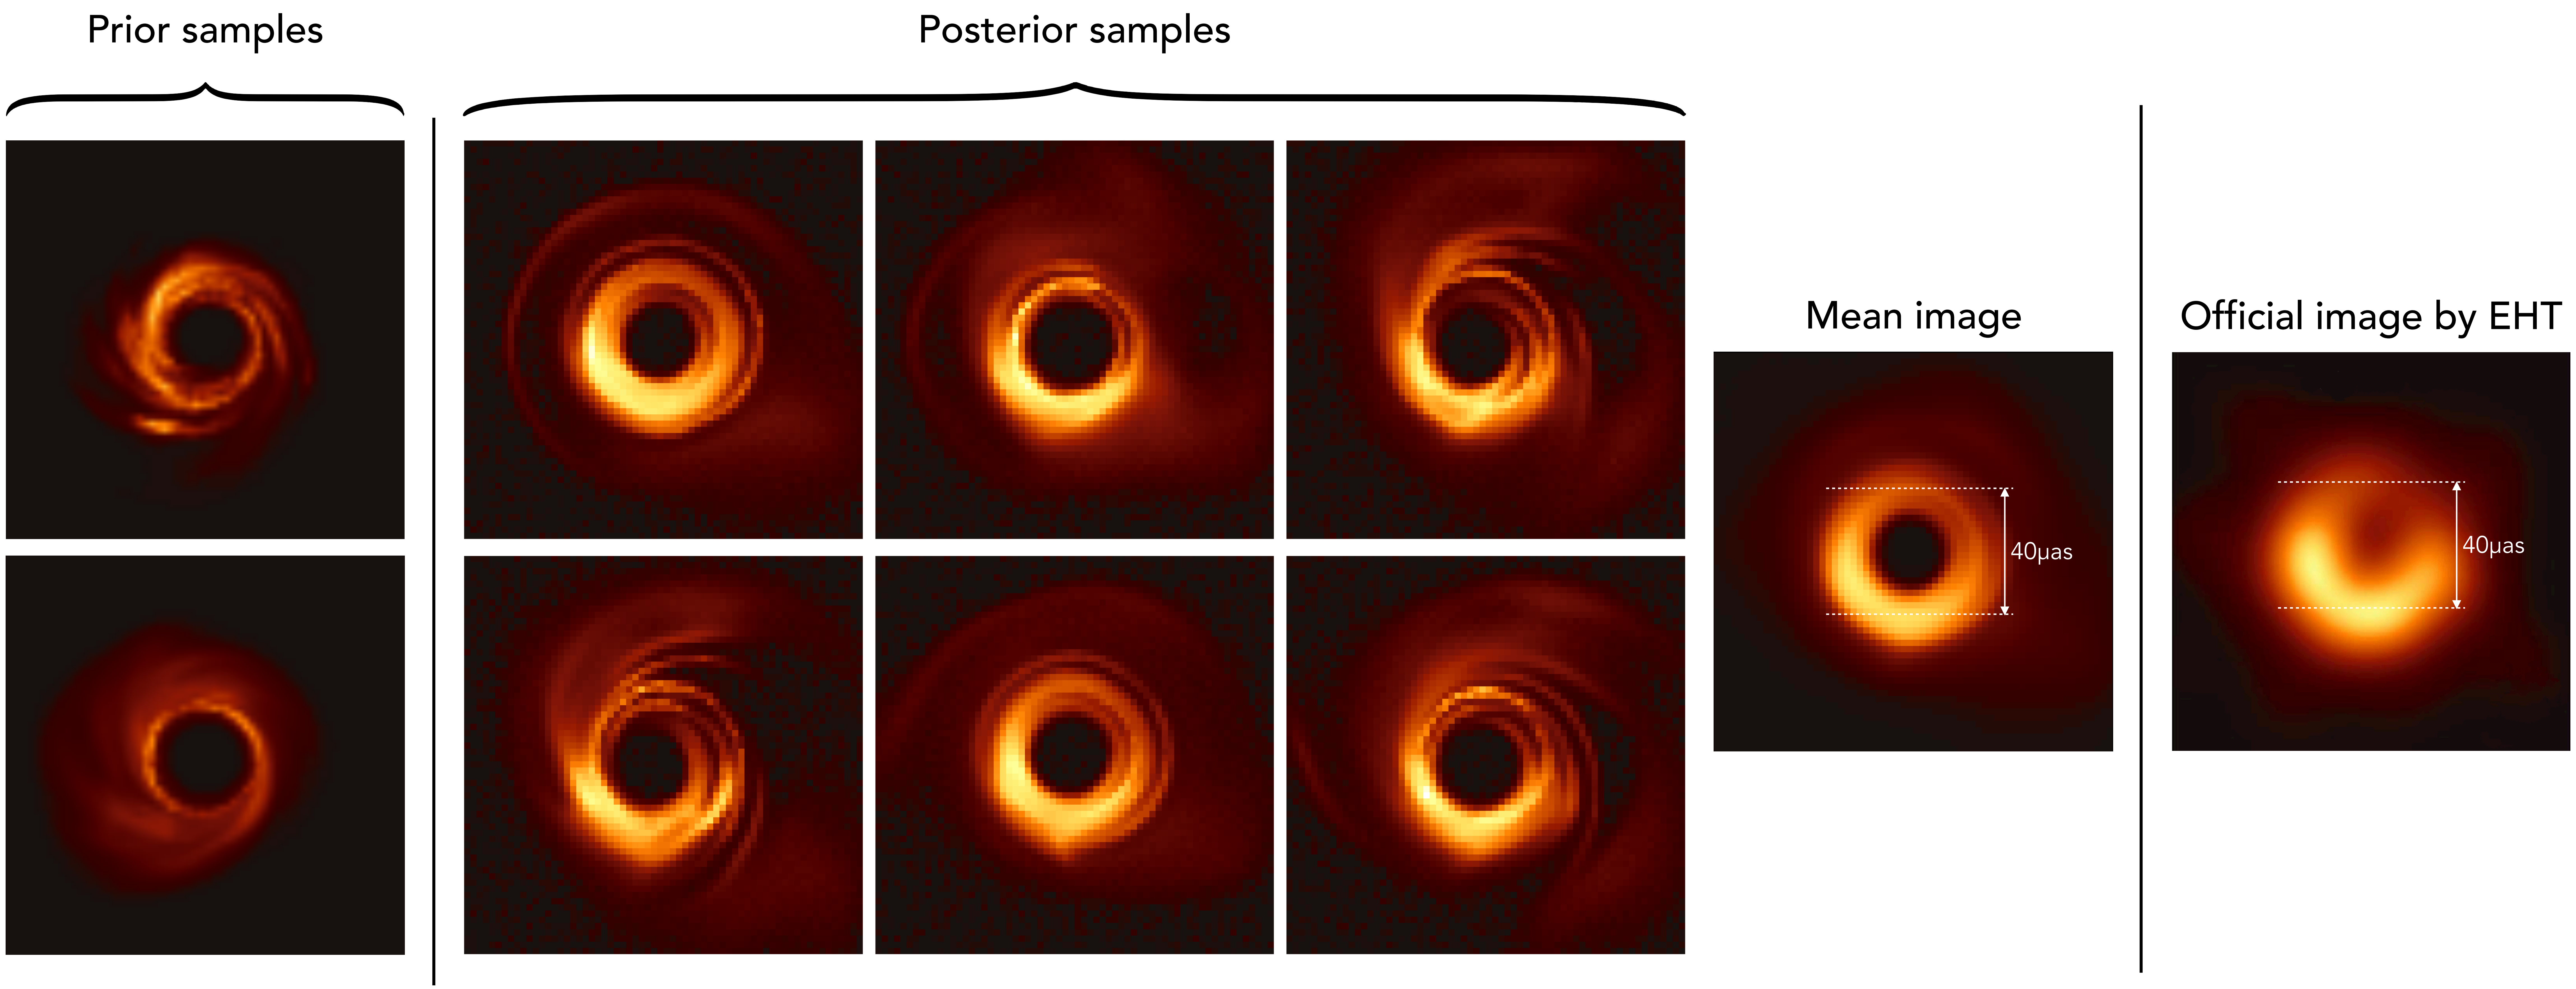

** Experiment was performed with real data for the M87 black hole with non-convex constraints*

Principled Probabilistic Imaging using Diffusion Models as Plug-and-Play Priors

NeurIPS 2024 Poster

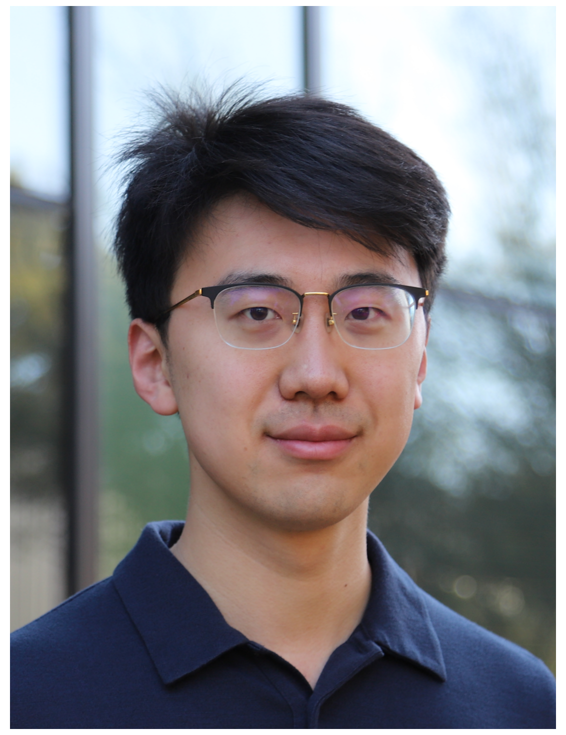

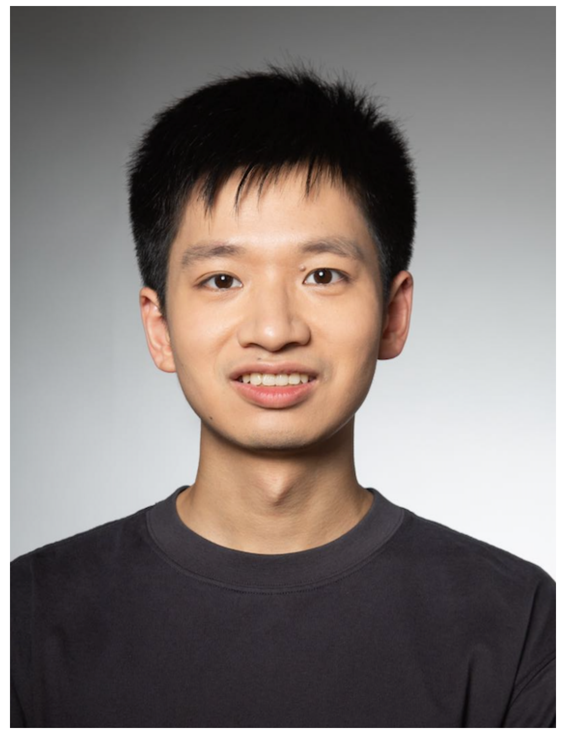
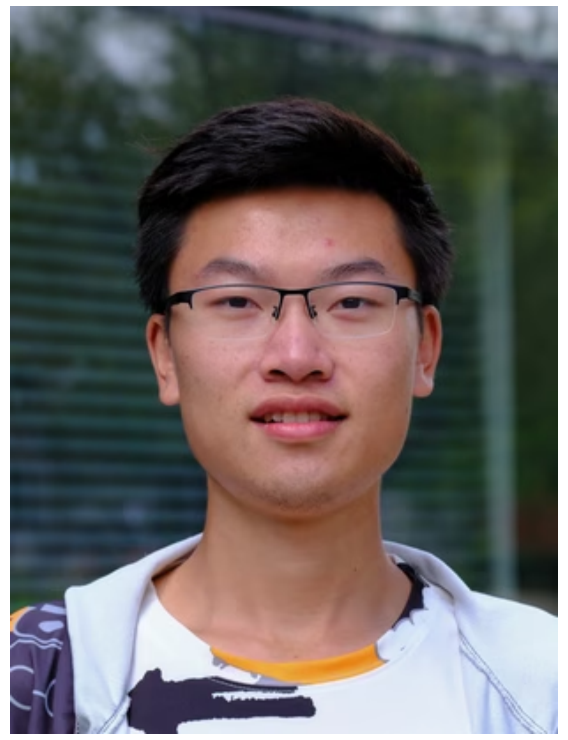

Zihui Wu | Yu Sun | Yifan Chen | Bingliang Zhang | Yisong Yue | Katherine L. Bouman

---

- **Project page:** http://imaging.cms.caltech.edu/pnpdm/
- **Code:** https://github.com/zihuiwu/PnP-DM-public/
- **Scan the QR code for more information!**