An Expectation-Maximization Algorithm for Training Clean Diffusion Models from Corrupted Observations

Weimin Bai, Yifei Wang, Wenzheng Chen, He Sun

Peking University

Diffusion Models

Clean Dataset for Training: *FFHQ*

What about few clean images for training?

Corrupted Observations are Sufficient

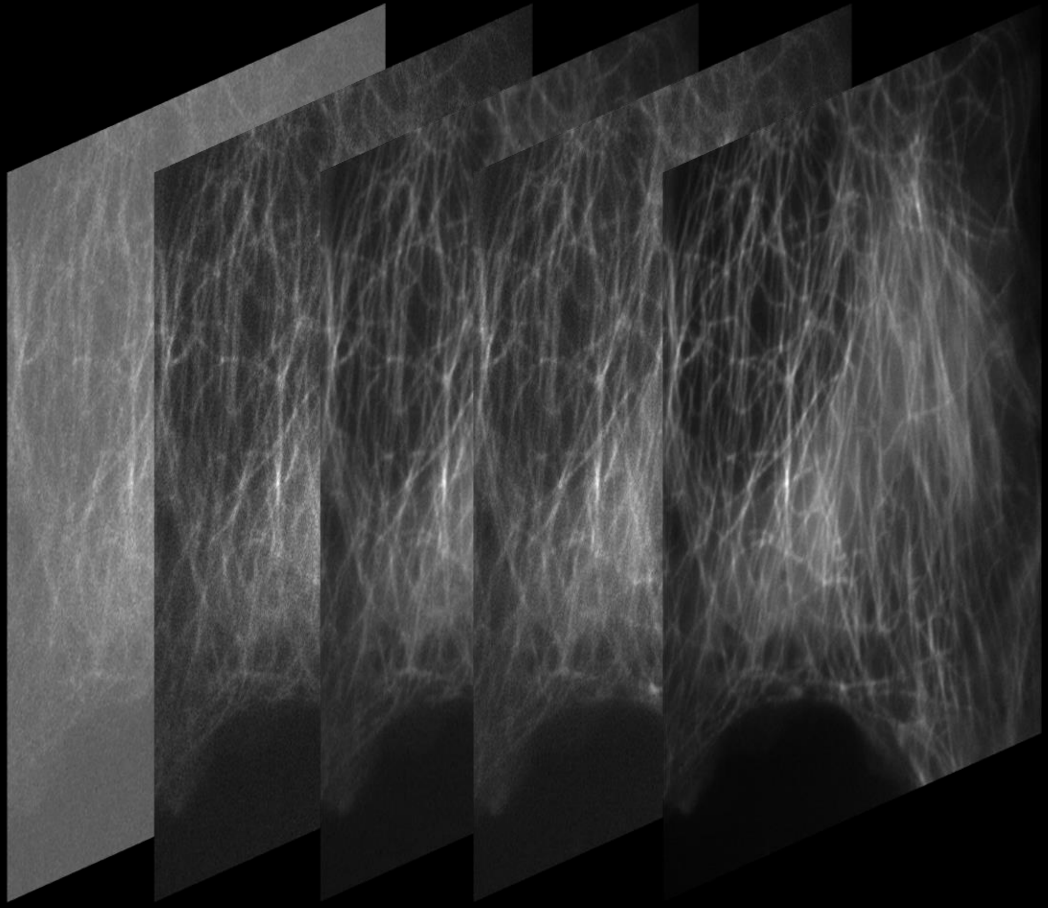

Corrupted Observations

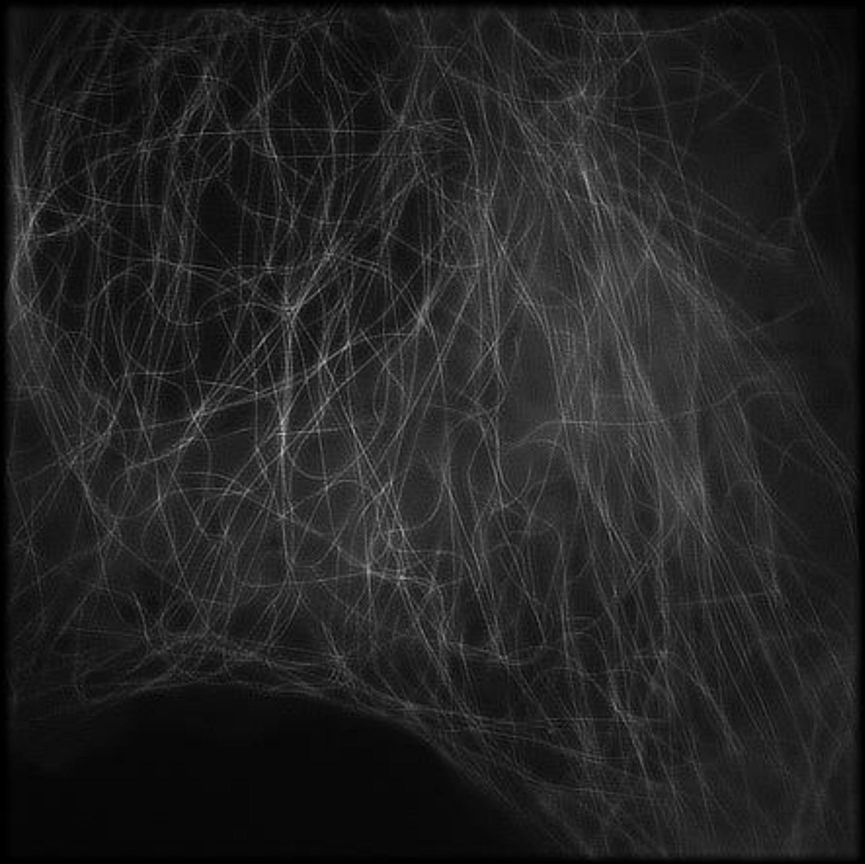

Clean Image

An example in super-resolution fluorescent microscopy

EMDiffusion: Learn Generative Image Priors from Corrupted Observations

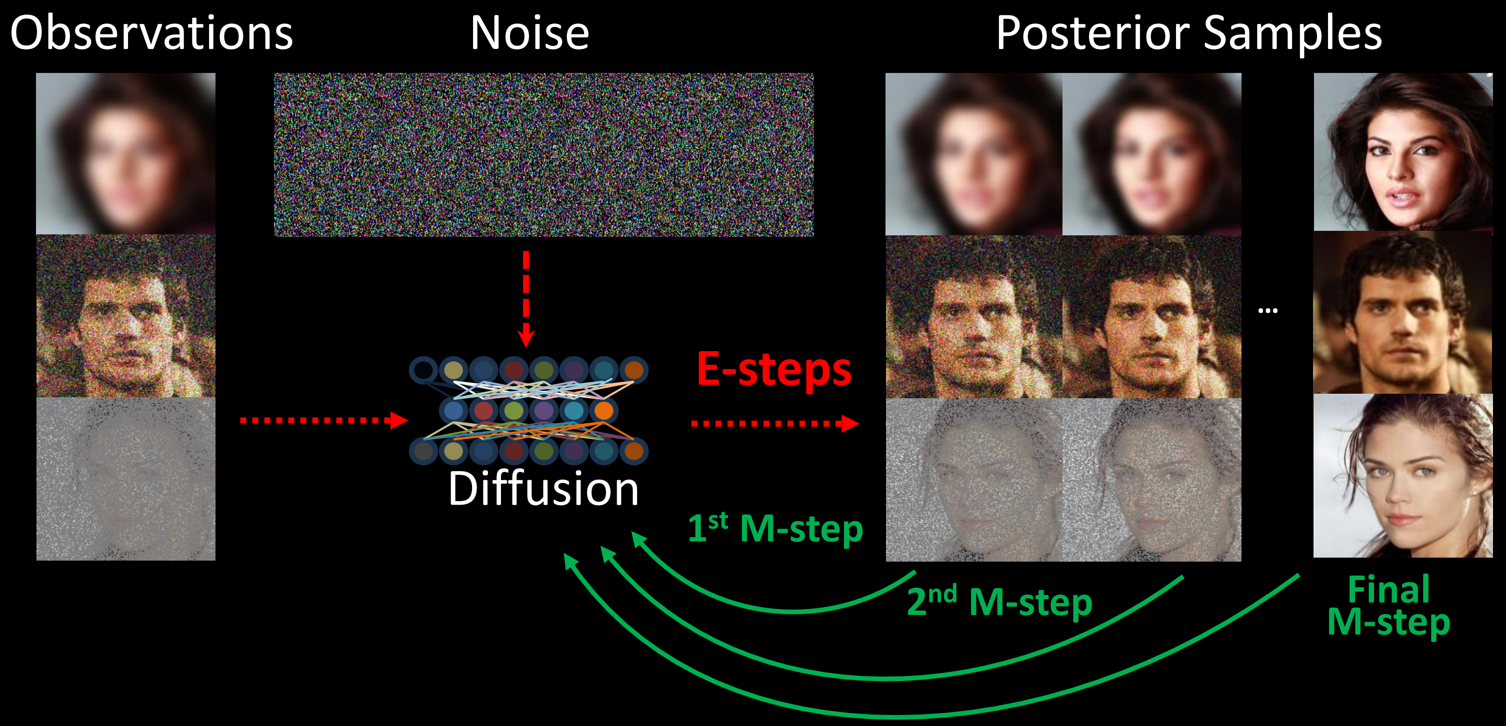

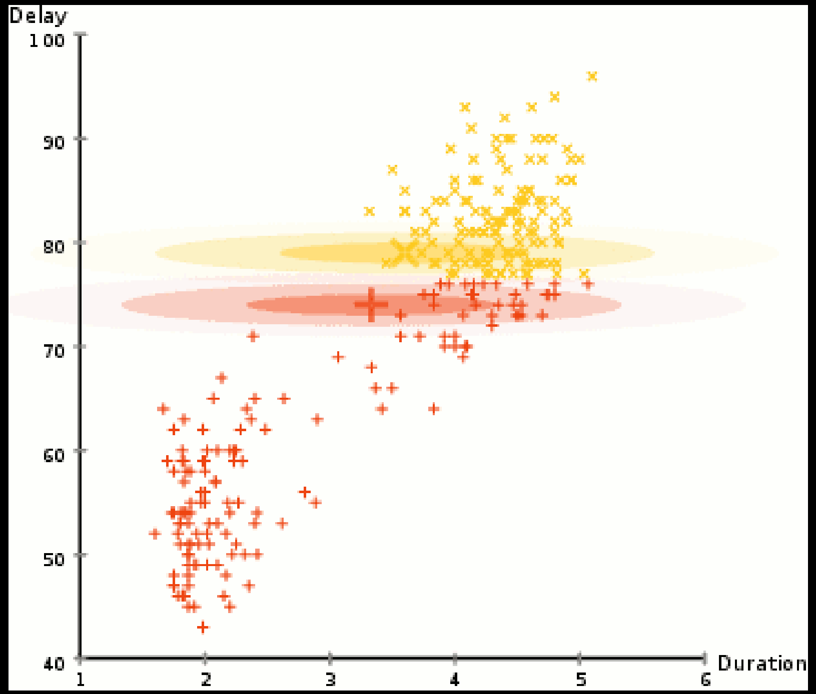

[1] *An Expectation-Maximization Algorithm for Training Clean Diffusion Models from Corrupted Observations.* **Weimin Bai**, Yifei Wang, Wenzheng Chen, He Sun. NeurIPS 2024.

Initialization: Training Diffusion with Limited Clean Data

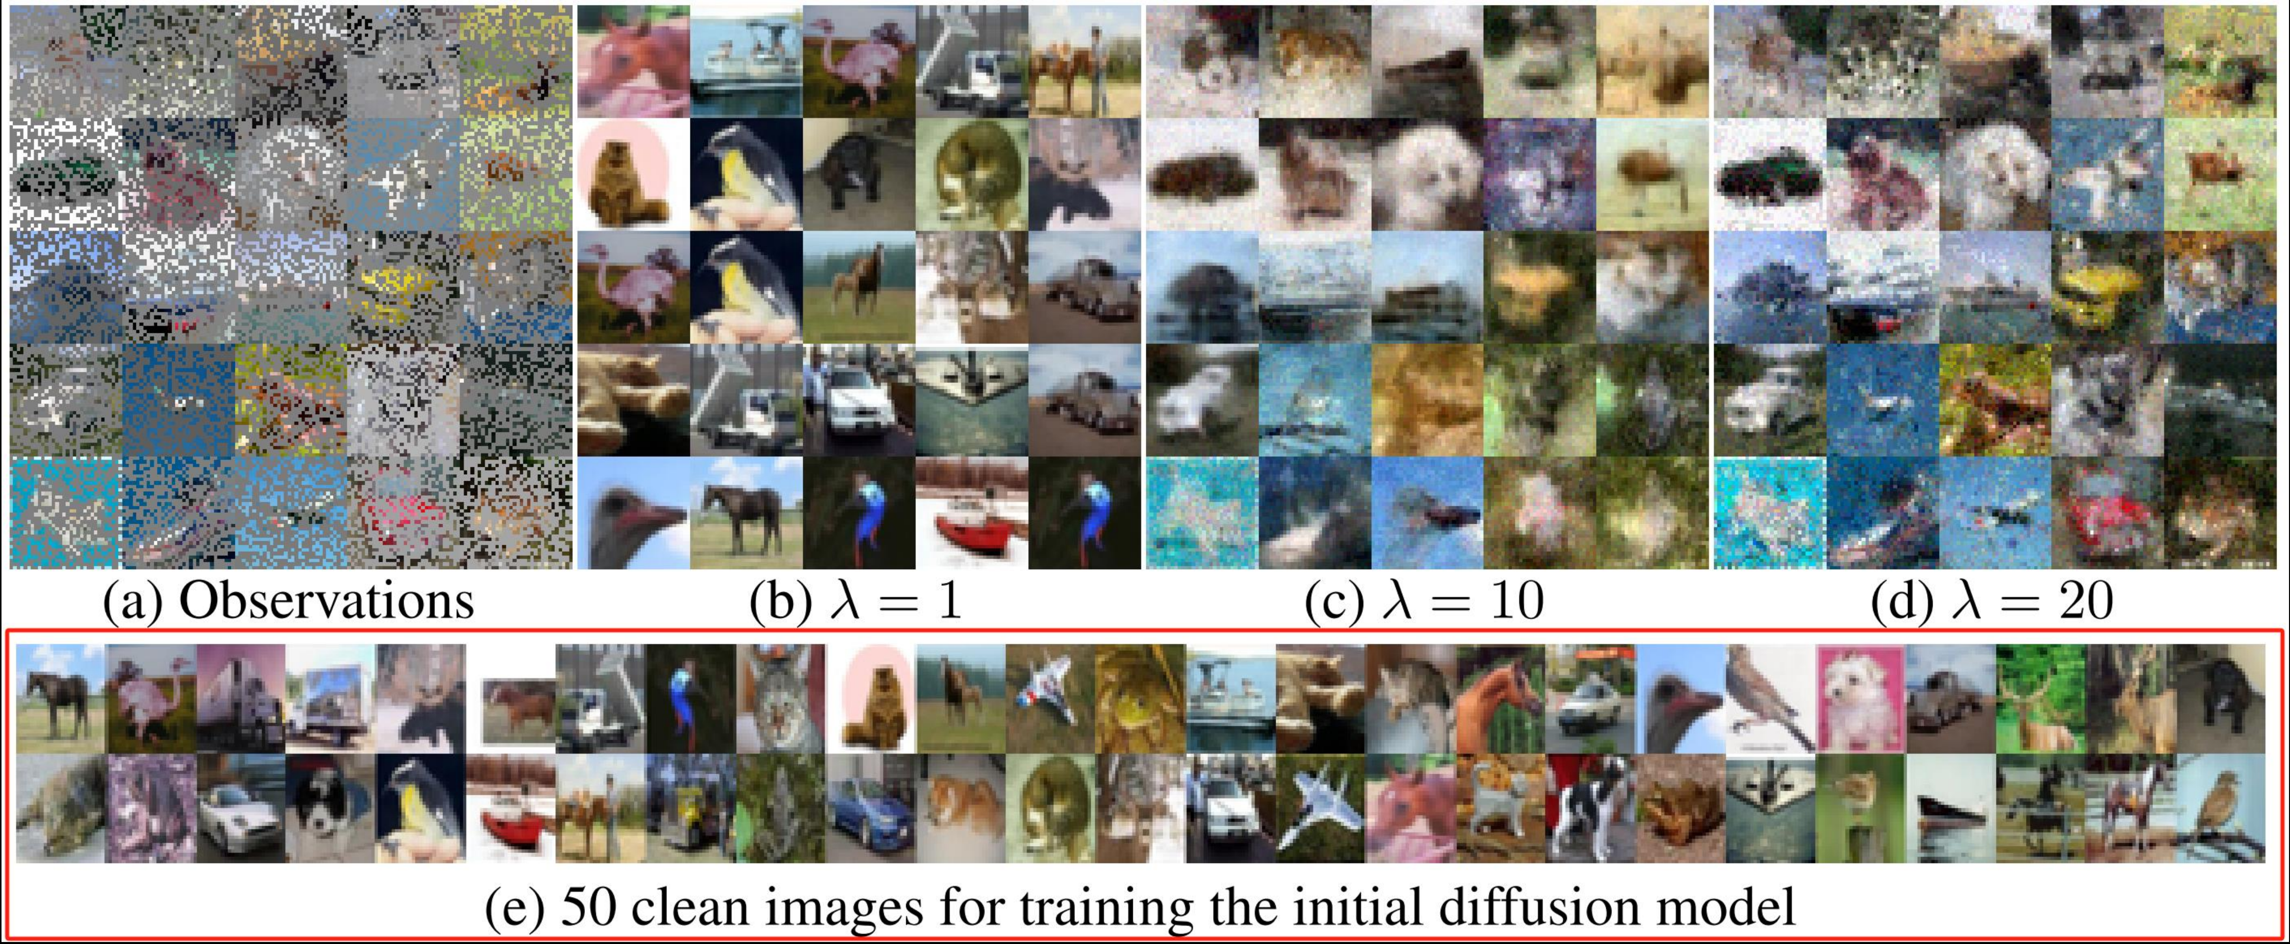

(a) Observations (b) $\lambda = 1$ (c) $\lambda = 10$ (d) $\lambda = 20$

(e) 50 clean images for training the initial diffusion model

E-step: Diffusion Posterior Sampling

$$p(x|y) \propto p(y|x)\ p(x)$$

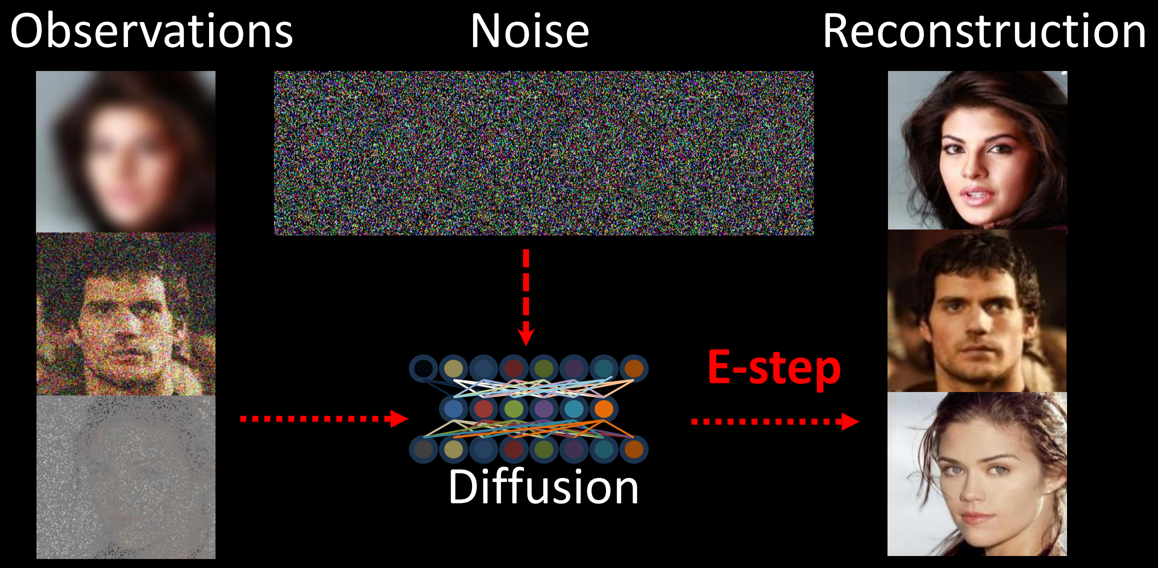

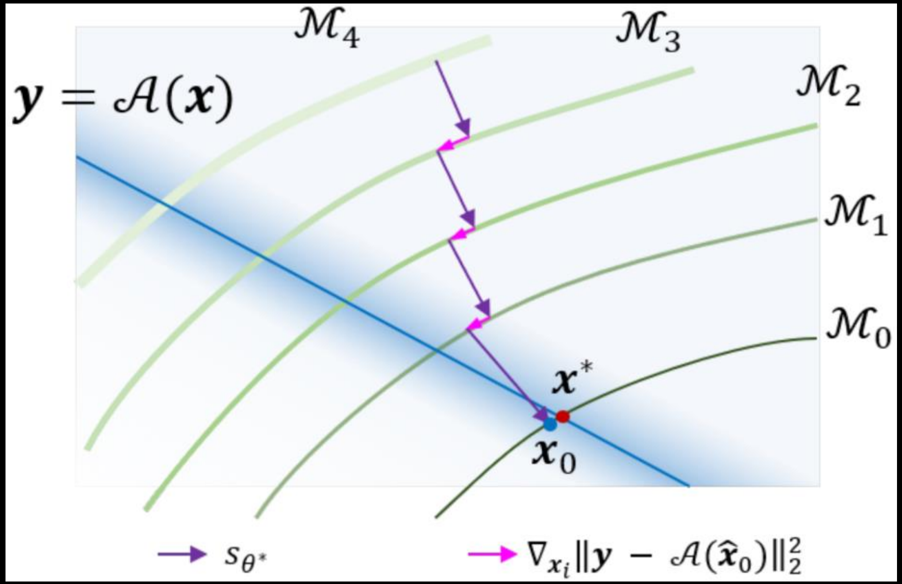

E-step: **Reconstruct images from corrupted observations using current generative image prior**

[1] *Diffusion Posterior Sampling for General Noisy Inverse Problems.* Hyungjin Chung, et al. ICLR 2023.

M-step: Updating Diffusion Model’s Weights

$$p(x|y) \propto p(y|x)\ {\color{red}p(x)}$$

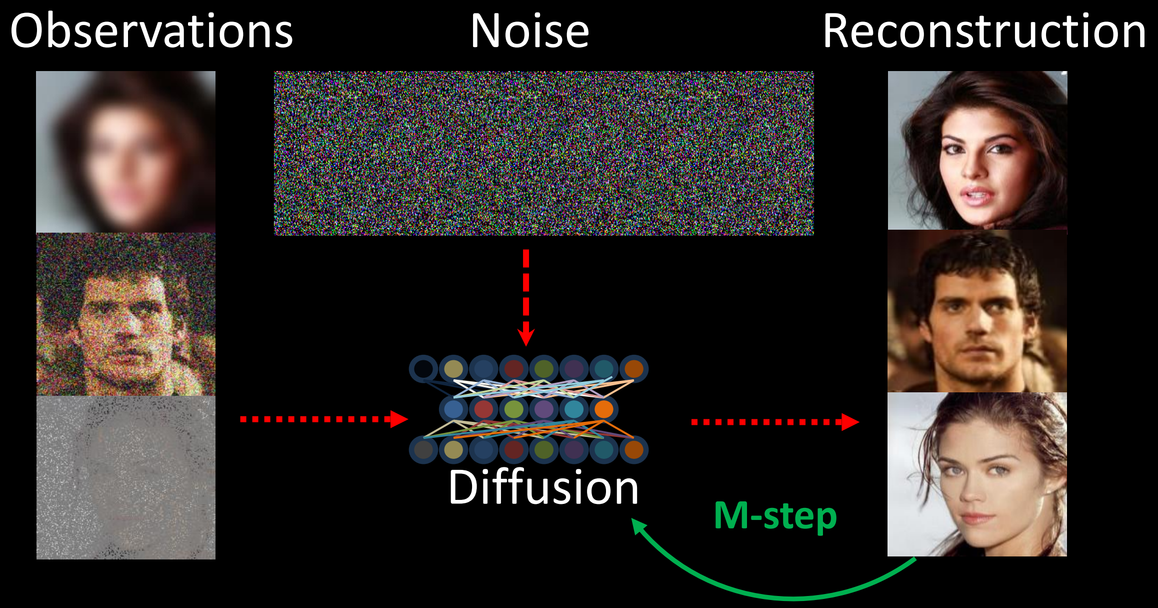

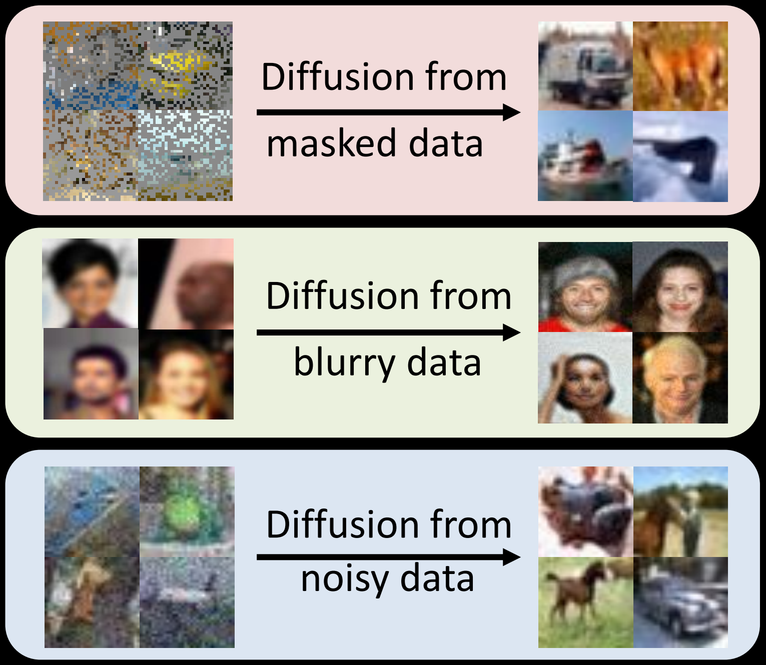

M-step: Refining generative image prior using reconstructions

[1] *An Expectation-Maximization Algorithm for Training Clean Diffusion Models from Corrupted Observations.* **Weimin Bai**, Yifei Wang, Wenzheng Chen, He Sun. NeurIPS 2024.

EMDiffusion: Learn Generative Image Priors from Corrupted Observations

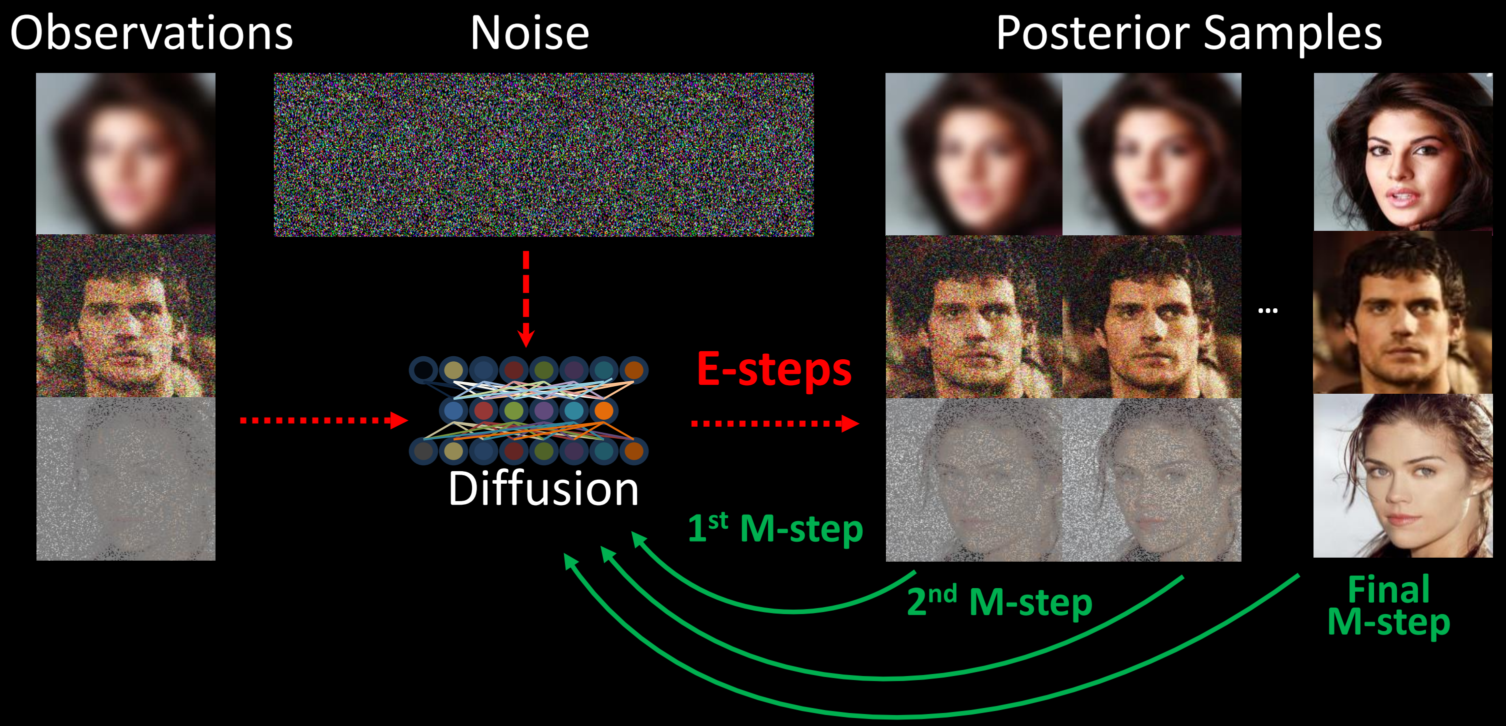

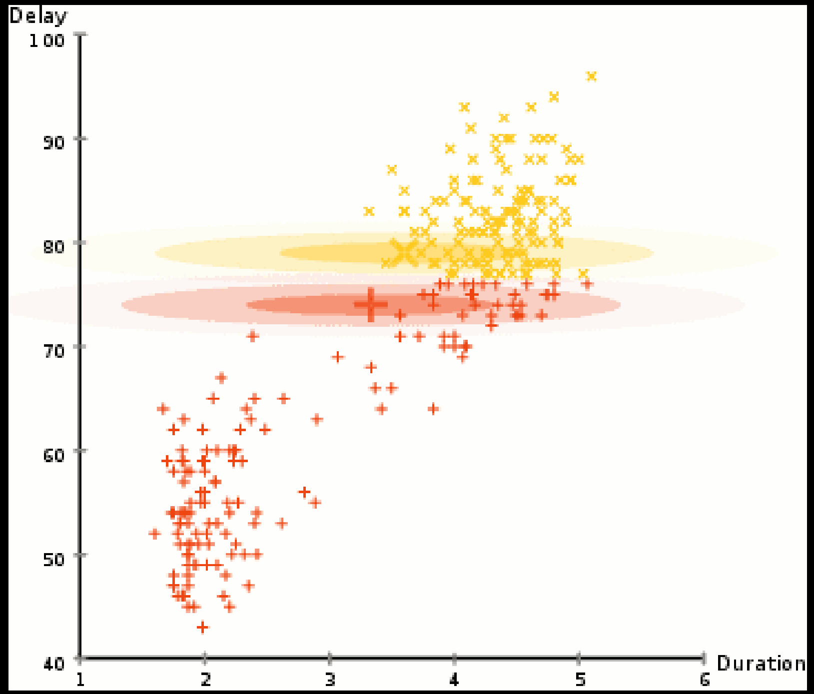

[1] *An Expectation-Maximization Algorithm for Training Clean Diffusion Models from Corrupted Observations.* **Weimin Bai**, Yifei Wang, Wenzheng Chen, He Sun. NeurIPS 2024.

Progressive Learning Process

$$p(x|y) \propto p(y|x)\ p(x)$$

Learned Diffusion Model

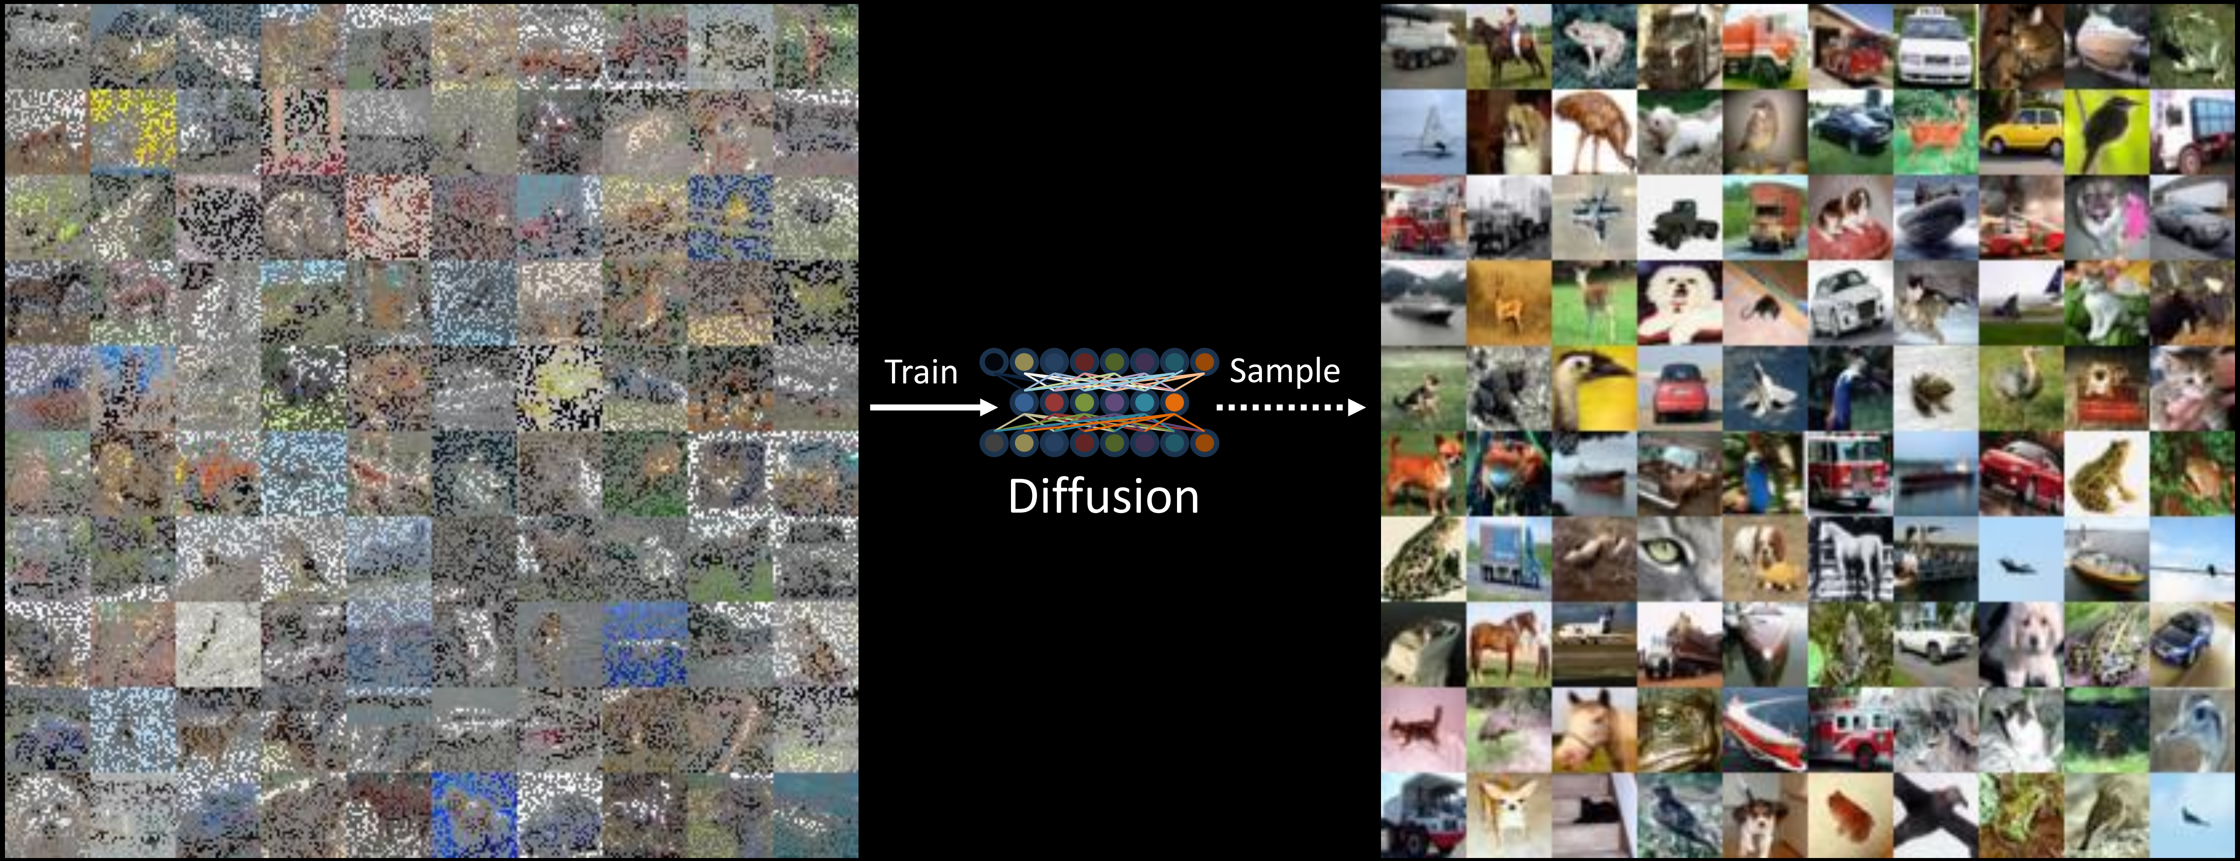

Masked Observations

Generated Samples

Learned Diffusion Model

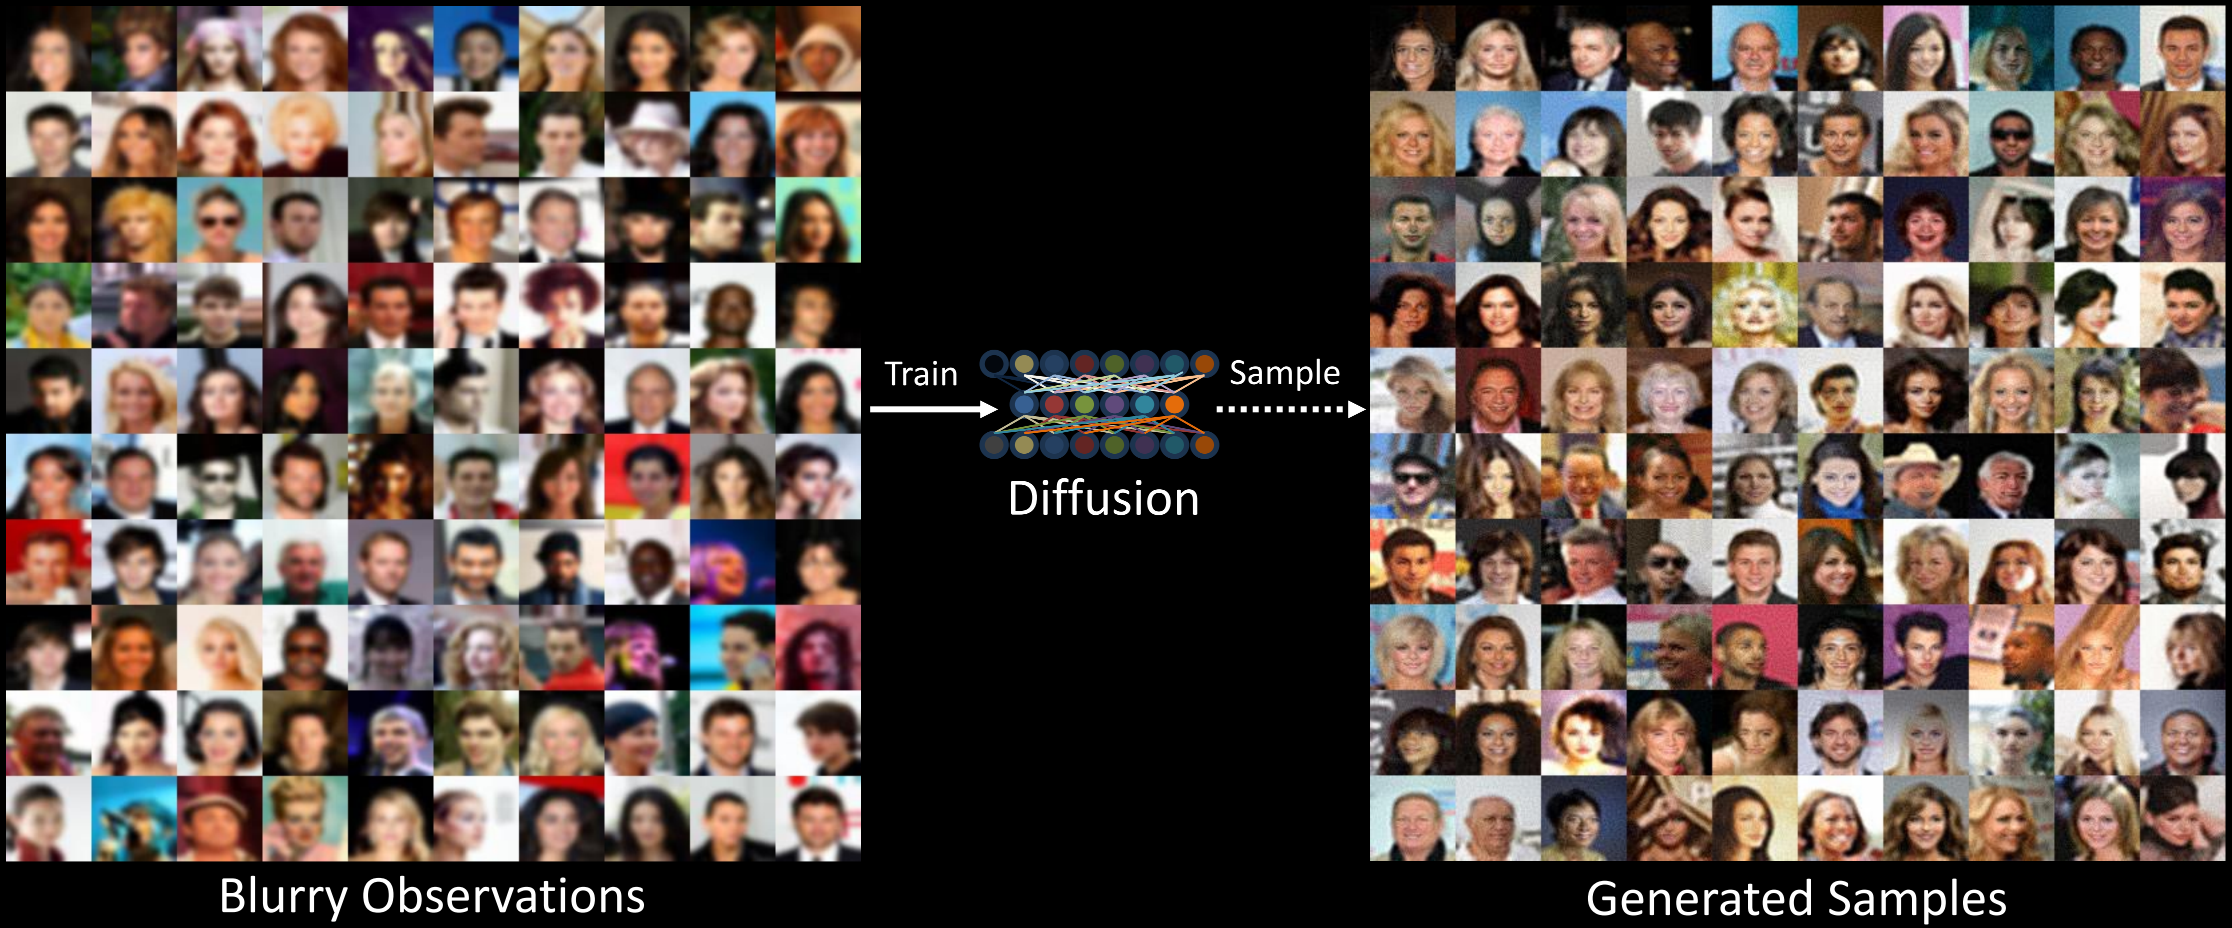

Learned Diffusion Model

Noisy Observations

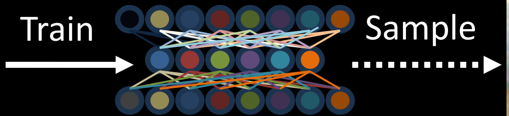

Diffusion

Generated Samples

Quantitative Comparison

Method	CIFAR10-Inpainting			CIFAR10-Denoising			CelebA-Deblurring		
	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓
Observations	13.49	0.295	234.47	18.05	0.047	132.59	22.47	0.365	72.83
DPS w/ clean prior	25.44	0.008	7.08	25.91	0.010	7.08	29.05	0.013	10.24
Noise2Self [3]	-	-	-	21.32	0.227	<u>92.06</u>	-	-	-
SURE-Score [1]	15.75	0.182	220.01	<u>22.42</u>	0.138	132.61	<u>22.07</u>	0.383	191.96
AmbientDiffusion [14]	<u>20.57</u>	<u>0.027</u>	<u>28.88</u>	21.37	<u>0.033</u>	114.13	21.16	<u>0.158</u>	**83.99**
Ours	**24.70**	**0.009**	**21.08**	**23.16**	**0.022**	**86.47**	**23.74**	**0.103**	<u>91.89</u>

[1] *An Expectation-Maximization Algorithm for Training Clean Diffusion Models from Corrupted Observations.* **Weimin Bai**, Yifei Wang, Wenzheng Chen, He Sun. NeurIPS 2024.

Initialization and Annealing of Diffusion Prior

$$p(x|y) \propto p(y|x)\, p^{\lambda}(x)$$

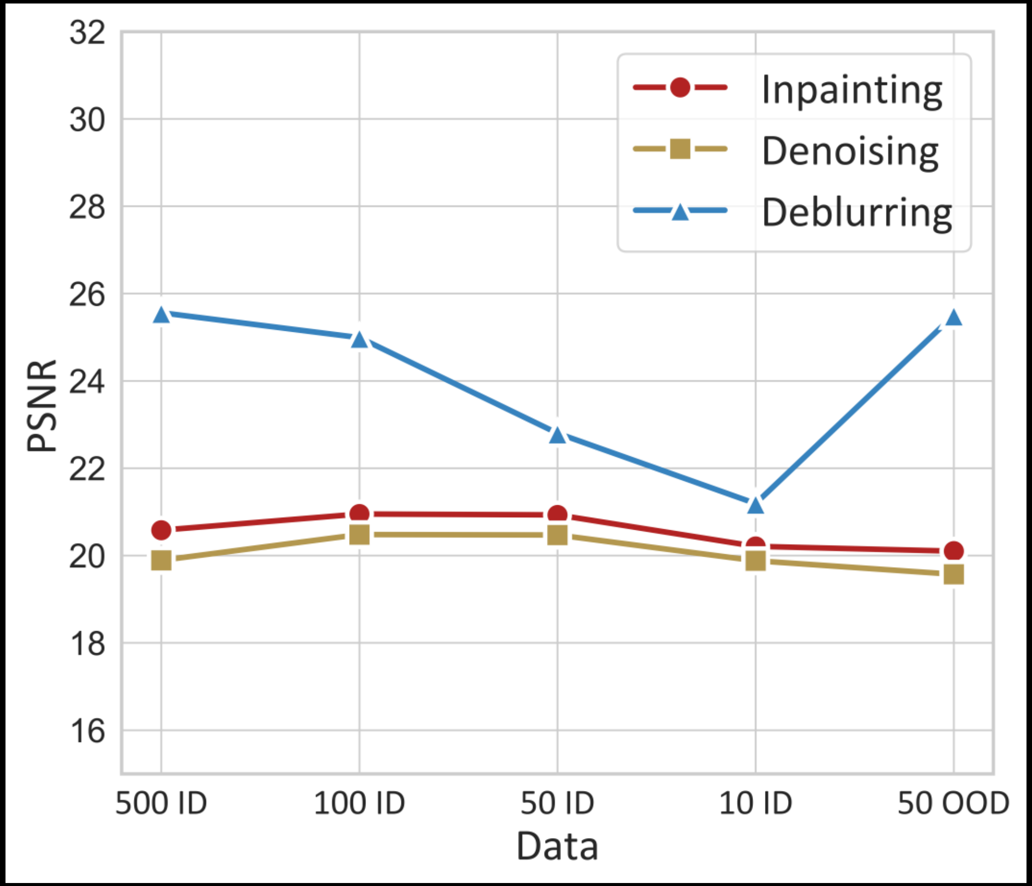

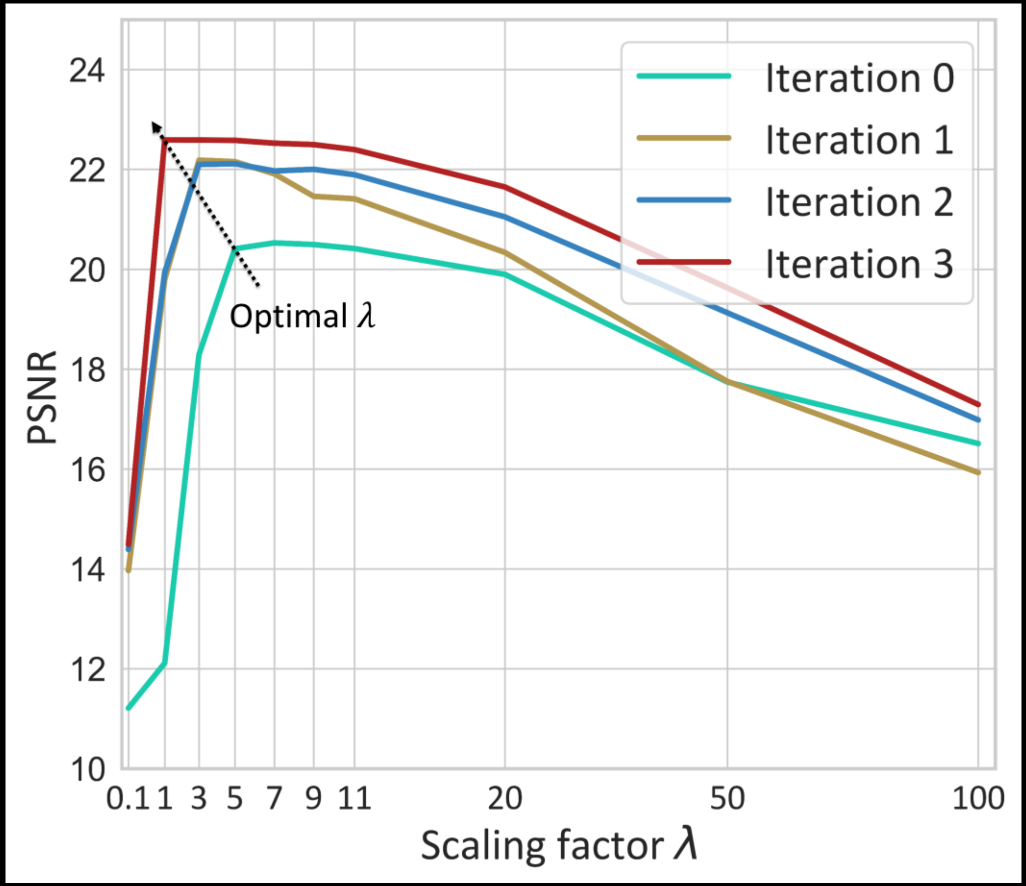

ID: In-distribution clean images
OOD: Out-of-distribution clean images

Scaling factor controls the strength of prior

Reset of Diffusion Model Weights

Conclusion

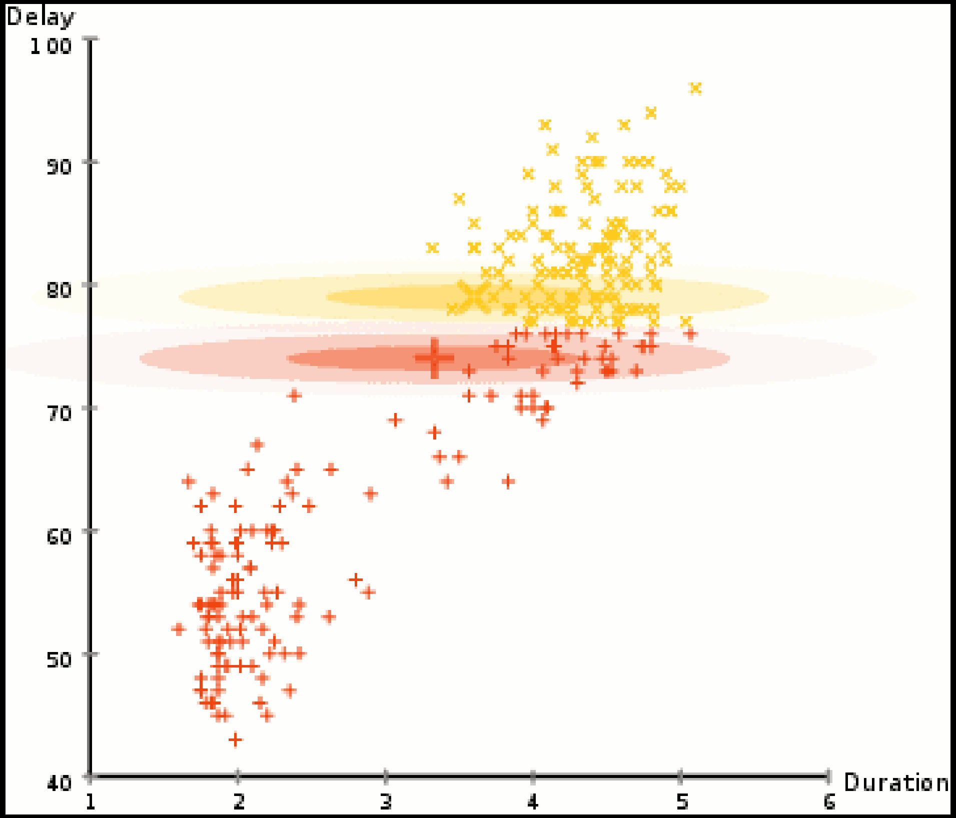

EMDiffusion: Learn from Corruption

Without large-scale clean images, Diffusion models can still be **trained**

- **EM-Diffusion:** E-step (DPS) + M-step (Diffusion Model Training);
- Training data: large-scale corrupted observations
- Future work: eliminate the initialization dependency on clean data.

$$p(x|y) \propto p(y|x)\ p(x)$$