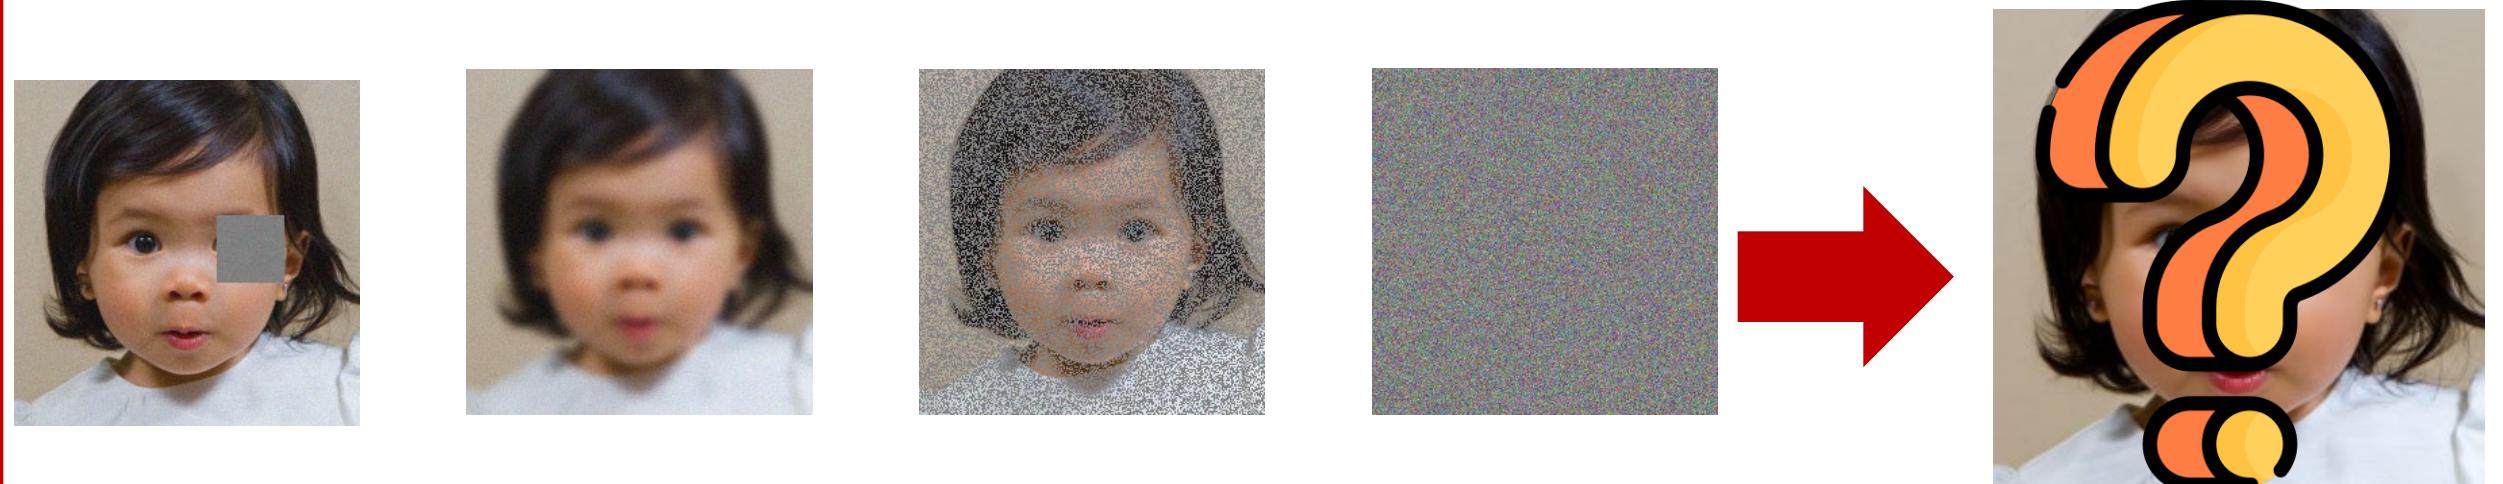


Diffusion Models for Image Generation and Inverse Problems

Yue Qi
Stanford Online

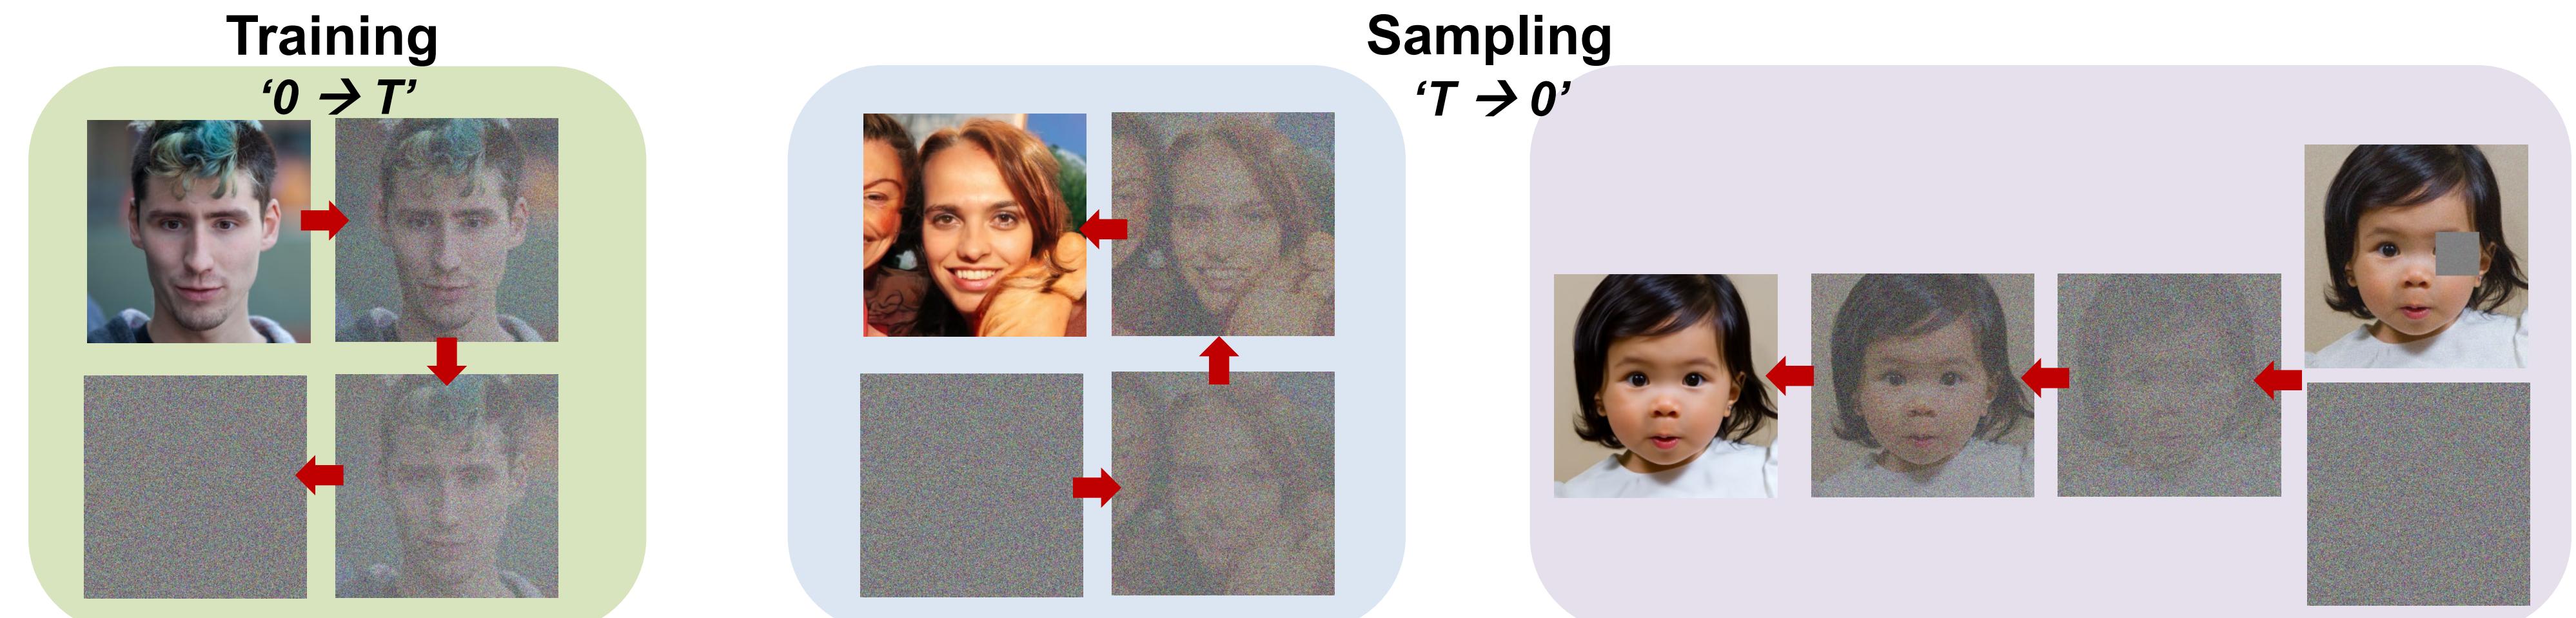
Motivation



Diffusion model for image restoration and generation

- The straightforward structure and efficient training process of diffusion models have made them a popular choice for generative modeling.
- In this work, we implemented several diffusion model approaches utilizing a pre-trained score predictor for image denoising, unconditional image generation, and tackling inverse problems, including inpainting and deconvolution.

Methods



Pre-trained score
($s_\theta(x_t, t)$) predictor

Reverse Diffusion: DDPM

$$\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t + (1 - \bar{\alpha}_t) s_\theta(x_t, t))$$
$$x_{t-1}' = \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \hat{x}_0 + \sqrt{1 - \alpha_t} z$$
$$x_{t-1} = x_{t-1}' - \frac{1}{\sigma^2 + \gamma_t^2} \nabla_{x_t} \|\mathcal{A}(x_t) - y\|^2$$

Posterior Sampling: ScoreALD, DPS

$$\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t + (1 - \bar{\alpha}_t) s_\theta(x_t, t))$$
$$x_{t-1}' = \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \hat{x}_0 + \sqrt{1 - \alpha_t} z$$
$$x_{t-1} = x_{t-1}' - \frac{1}{\sigma^2 + \gamma_t^2} \nabla_{x_t} \|\mathcal{A}(x_t) - y\|^2$$

Related Work



References

- [1] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.Y. Zhu, S. Ermon, "SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations", ICLR 2022
- [2] J. Ho, et al, "Denoising Diffusion Probabilistic Models", NeurIPS 2020
- [3] H. Chung, et al, "Diffusion posterior sampling for general noisy inverse problems", ICLR 2023

Experimental Results

