Test-time anchoring for discrete diffusion posterior sampling

Litu Rout, Andreas Lugmayr, Yasamin Jafarian, Srivatsan Varadharajan, Constantine Caramanis, Sanjay Shakkottai, Ira Kemelmacher-Shlizerman

The University of Texas at Austin
Machine Learning Laboratory

The University of Texas at Austin
Chandra Department of Electrical and Computer Engineering
Cockrell School of Engineering

Google

Overview

Use Discrete Diffusion

- High-quality samples (images or text) in significantly fewer steps

Control Generation

- Posterior Sampling
- Tilt categorical distribution for inverse problems

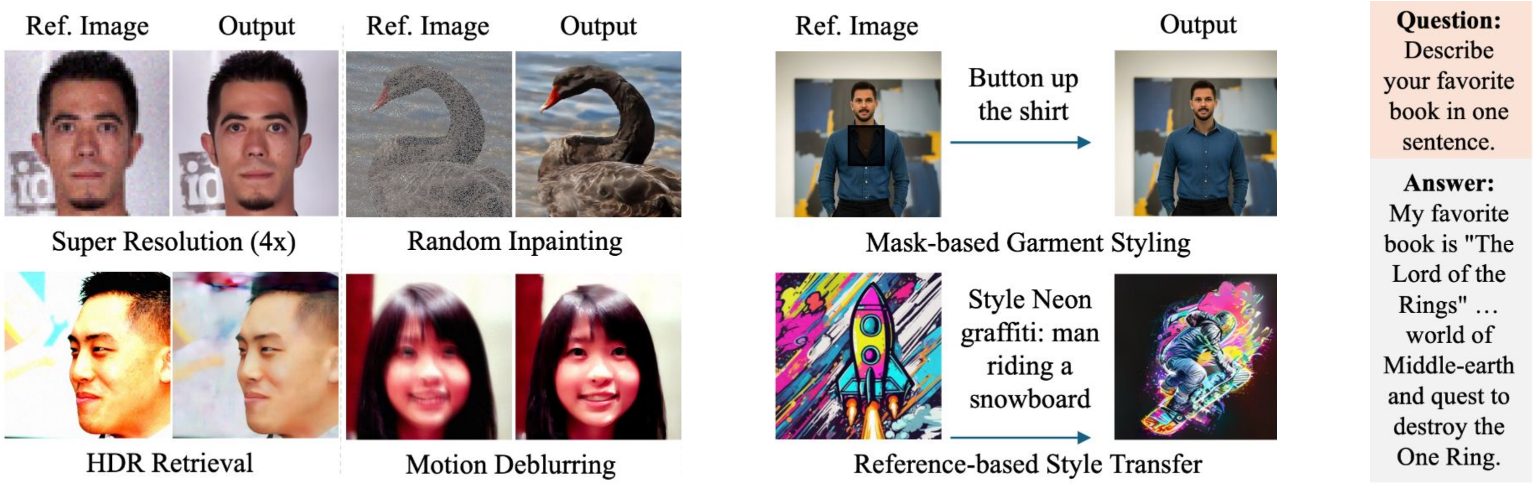

Objective of Anchored Posterior Sampling (APS)

Negative log posterior (Theorem 3.3):

$$-\log p_\varphi(\mathrm{x}|\mathrm{y}) \;\leq \mathcal{L}_{\mathrm{NELBO}}(\mathrm{x};\theta) + \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\left[\sum_{l=1}^{L} \log \mathcal{Z}^l_\varphi(Z_{t(i)}, \mathrm{y})\right]$$

$$+ \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\left[\left[\frac{\alpha_{s(i)} - \alpha_{t(i)}}{1-\alpha_{t(i)}} \sum_{l=1}^{L} \log \frac{\langle \mathrm{x}^l_\theta(Z_{t(i)}), \mathrm{x}^l\rangle}{\langle \mathrm{x}^l_\varphi(Z_{t(i)}), \mathrm{x}^l\rangle}\right] - \mathbb{E}_{q(Z_{s(i)}|Z_{t(i)},\mathrm{x})}\left[\sum_{l=1}^{L} \log q(\mathrm{y}|\mathrm{x}_\varphi(Z_{t(i)}; Z^l_{s(i)}))\right]\right]$$

- Standard masked diffusion loss NELBO
- Normalizing factor via stop-gradient (Sohl-Dickstein et. al., 2015, Murata et. al., 2025)
    - Exclude during gradient computation
    - Normalize transition via softmax
- Log ratio captures adaptation gap
- Measurement likelihood tilt (external reward model)

Approximations

Negative log posterior (Theorem 3.3):

$$-\log p_\varphi(\mathrm{x}|\mathrm{y}) \leq \mathcal{L}_{\mathrm{NELBO}}(\mathrm{x};\theta) + \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\Big[\sum_{l=1}^{L} \log \mathcal{Z}_\varphi^l(Z_{t(i)},\mathrm{y})\Big]$$
$$+ \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\Bigg[\Big[\frac{\alpha_{s(i)}-\alpha_{t(i)}}{1-\alpha_{t(i)}} \sum_{l=1}^{L} \log \frac{\langle \mathrm{x}_\theta^l(Z_{t(i)}),\mathrm{x}^l\rangle}{\langle \mathrm{x}_\varphi^l(Z_{t(i)}),\mathrm{x}^l\rangle}\Big] - \mathbb{E}_{q(Z_{s(i)}|Z_{t(i)},\mathrm{x})}\Big[\sum_{l=1}^{L} \log q(\mathrm{y}|\mathrm{x}_\varphi(Z_{t(i)};Z_{s(i)}^l))\Big]\Bigg]$$

- Approximation 1: At test-time, minimize error at each time step

$$\mathcal{L}_t(\varphi) \coloneqq \frac{\alpha_s-\alpha_t}{1-\alpha_t}\sum_{l=1}^{L}\Big[\log \frac{\langle \mathbf{x}_\theta^l(\mathbf{z}_t),\mathbf{x}^l\rangle}{\langle \mathbf{x}_\varphi^l(\mathbf{z}_t),\mathbf{x}^l\rangle}\,\mathbf{1}_{\{\mathbf{z}_t^l=\mathbf{m}\}}\Big] - \mathbb{E}_{q(Z_s|\mathbf{z}_t,\mathbf{x})}\Big[\sum_{l=1}^{L}\log q(\mathbf{y}|\mathbf{x}_\varphi(\mathbf{z}_t;Z_s^l))\Big]$$

- Approximation 2: Replace original sample with model prediction

$$\hat{\mathcal{L}}_t(\varphi) \coloneqq \frac{\alpha_t-\alpha_s}{1-\alpha_t}\sum_{l=1}^{L}\Big[\log\langle \mathbf{x}_\varphi^l(\mathbf{z}_t),\mathbf{x}_\theta^l(\mathbf{z}_t)\rangle\,\mathbf{1}_{\{\mathbf{z}_t^l=\mathbf{m}\}}\Big] - \mathbb{E}_{q(Z_s|\mathbf{z}_t,\mathbf{x})}\Big[\sum_{l=1}^{L}\log q(\mathbf{y}|\mathbf{x}_\varphi(\mathbf{z}_t;Z_s^l))\Big]$$

Approximations for Likelihood: Quantized Expectation

Negative log posterior (Theorem 3.3):

$$-\log p_{\varphi}(\mathrm{x}|\mathrm{y}) \le \mathcal{L}_{\mathrm{NELBO}}(\mathrm{x};\theta) + \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\left[\sum_{l=1}^{L} \log \mathcal{Z}^{l}_{\varphi}(Z_{t(i)}, \mathrm{y})\right] \qquad \mathcal{O}(K^L)$$

$$+ \sum_{i=1}^{T} \mathbb{E}_{q(Z_{t(i)}|\mathrm{x})}\left[\left[\frac{\alpha_{s(i)} - \alpha_{t(i)}}{1-\alpha_{t(i)}} \sum_{l=1}^{L} \log \frac{\langle \mathrm{x}^{l}_{\theta}(Z_{t(i)}), \mathrm{x}^{l}\rangle}{\langle \mathrm{x}^{l}_{\varphi}(Z_{t(i)}), \mathrm{x}^{l}\rangle}\right] - \mathbb{E}_{q(Z_{s(i)}|Z_{t(i)},\mathrm{x})}\left[\sum_{l=1}^{L} \log q(\mathrm{y}|\mathrm{x}_{\varphi}(Z_{t(i)}; Z^{l}_{s(i)}))\right]\right]$$

- Approximation 3: Our final surrogate objective

$$\bar{\mathrm{x}}^{l} = \sum_{k} c_k \mathrm{x}^{l}_{\varphi}(\mathrm{z}_t)[k] \qquad \mathbf{x}^{l} = \mathcal{Q}_{\mathrm{lfq}}(\bar{\mathbf{x}}^{l})$$

$$\tilde{\mathrm{x}}_{\varphi}(\mathrm{z}_t) = \bar{\mathrm{x}} + [\mathrm{x} - \bar{\mathrm{x}}]_{\mathrm{sg}}$$

$$\hat{\mathcal{L}}_t(\varphi) := \frac{\alpha_t - \alpha_s}{1-\alpha_t} \sum_{l=1}^{L}\left[\log\langle \mathbf{x}^{l}_{\varphi}(\mathbf{z}_t), \mathbf{x}^{l}_{\theta}(\mathbf{z}_t)\rangle\, \mathbf{1}_{\{\mathbf{z}^{l}_{t}=\mathbf{m}\}}\right] - L\cdot \log q(\mathbf{y}|\tilde{\mathbf{x}}_{\varphi}(\mathbf{z}_t))\Big]$$

$$\mathcal{O}(KL)$$

(Straight-through estimator: Bengio et. al., 2013, Oord et. al., NeurIPS, 2017)

Anchored Remasking

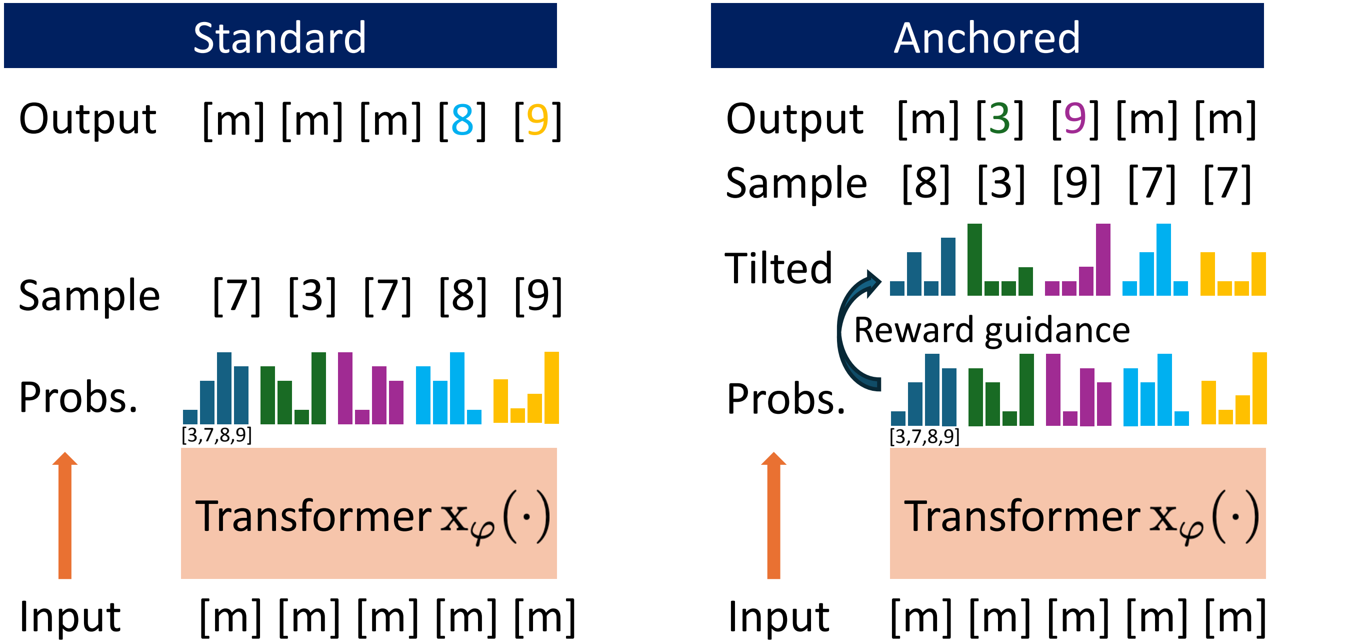

- Sample from product of marginals
- Selectively mask tokens with low marginal probability

- Sample from approximate joint supported on anchor tokens
- Selectively mask tokens with low tilted probability

Experiments: Linear Inverse Problems

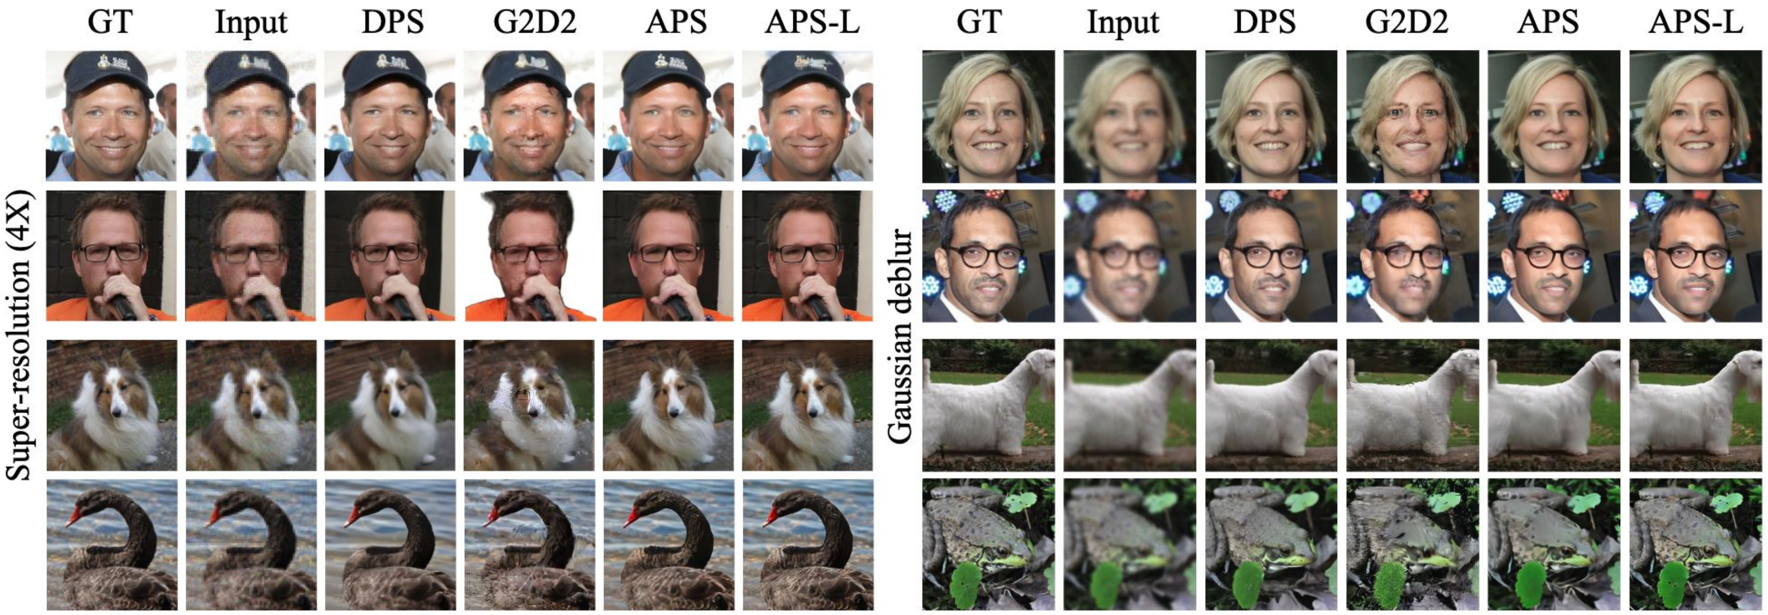

APS (ours) yields better results with sharper texture and refined facial features

Experiments: Linear Inverse Problems

Type	Method	(a) FFHQ				(b) ImageNet			
		SR (4×)		Deblur		SR (4×)		Deblur	
		LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑
Pixel	DPS	0.269	25.86	0.219	25.87	0.367	22.61	0.443	19.04
	DDRM	0.282	26.58	0.239	24.93	0.352	24.00	0.246	27.30
	DiffPIR	0.260	26.64	0.236	27.36	0.371	23.18	0.355	22.80
	DAPS	0.177	29.07	0.165	29.19	0.276	25.89	0.253	26.15
Latent	PSLD	0.276	27.62	0.304	27.37	0.332	24.43	0.365	24.04
	ReSample	0.507	22.98	0.329	25.69	0.382	22.63	0.438	22.32
	LatentDAPS	0.182	27.48	0.234	27.93	0.276	25.06	0.345	25.05
Uniform (Mask)	SVDD-PM	0.594	12.08	–	–	–	–	–	–
	G2D2	0.271	26.93	0.287	26.35	0.349	23.20	0.375	22.71
	SGDD	0.288	25.85	–	–	–	–	–	–
Mask	**APS**	**0.234**	**27.50**	**0.276**	**27.90**	**0.324**	**24.30**	**0.375**	**24.71**
	APS-L	**0.186**	**28.83**	**0.241**	**29.50**	**0.224**	**25.74**	**0.282**	**26.35**

APS (ours) outperforms discrete samplers and competes with continuous baselines

Applications

Style Guided Generation

Image Editing

…. wearing a vibrant, oversized, silk-screened graphic t-shirt … untucked, hanging loosely over the jeans …

Reward Guided Question Answering

Question: "Write a short poem about generative AI"

Dream-7B-Instruct

Answer:
In the realm of code and logic,
The future of AI is bright.
Solving problems with precision,
Transforming day and night.
With every step, it learns,
Solving complex problems.
The future of AI is here.

Anchored Posterior Sampling

Answer:
Generative AI, a marvel of the mind,
Creating new art, music, and design.
With deep learning, it's so bold,
To generate content, new and bold.

From poetry to code, it's clear,
Generative AI is here to steer.
It's creative, innovative, and neat,
A tool for humans, not a threat.

It's powerful, it's not mean,
It's here to help us dream.
With machine learning, it's clear,
Generative AI is here to steer.

So let's embrace it, let's not fear,
For Generative AI is here to steer.