Greed is Good: Guided Generation from a Greedy Perspective

Zander W. Blasingame Chen Liu

Clarkson University

Motivation

Consider the usual *flow model*, let $(X_0, X_1) \sim p(x_0)q(x_1)$ where $q(x)$ is the target distribution and $p(x_0)$ is the prior distribution. Define X_t as $X_t = \alpha_t X_1 + \sigma_t X_0$, for schedule (α_t, σ_t). Then, the vector field of the *affine conditional flow* $\Phi_t(x|x_1) = \alpha_t x_1 + \sigma_t x$ is given by

$$u_t(x) = \mathbb{E}[\dot\alpha_t X_1 + \dot\sigma_t X_0 | X_t = x]. \tag{1}$$

Assume that u_t^θ is trained to zero loss, so $u_t^\theta = u_t$.

Problem statement. Find the optimal trajectory, *i.e.*, given a continuously differentiable loss function, $\mathcal{L} \in C^1(\mathbb{R}^d; \mathbb{R})$, find the minimizer

$$\min_{x_0} \mathcal{L}\left(x_0 + \int_0^1 u_\tau^\theta(x_\tau)\, \mathrm{d}\tau\right). \tag{2}$$

Posterior guidance. We can use the *gradient* of the denoiser $x_{1|t}^\theta(x) = \mathbb{E}[X_1|X_t = x]$ for guidance [3], *i.e.*, for some iteration x_n in the numerical scheme

$$x_n^{(k+1)} = x_n^{(k)} - \eta\nabla\mathcal{L}\left(x_{1|t}^\theta(x_n^{(k)})\right). \tag{3}$$

End-to-end guidance. Alternatively, optimize the initial point x_0 [1, 2], i.e.,

$$x_0^{(k+1)} = x_0^{(k)} - \eta\nabla\mathcal{L}\left(\Phi_{0,1}^\theta(x_0^{(k)})\right), \tag{4}$$

where $\Phi_{0,1}^\theta$ is the flow map from 0 to 1 induced by u_t^θ. This gradient can be found by backproping through the numerical ODE solver (*discretize-then-optimize* (DTO)) or by solving the *continuous adjoint equations* (*optimize-then-discretize* (OTD)) [4].

The greedy strategy

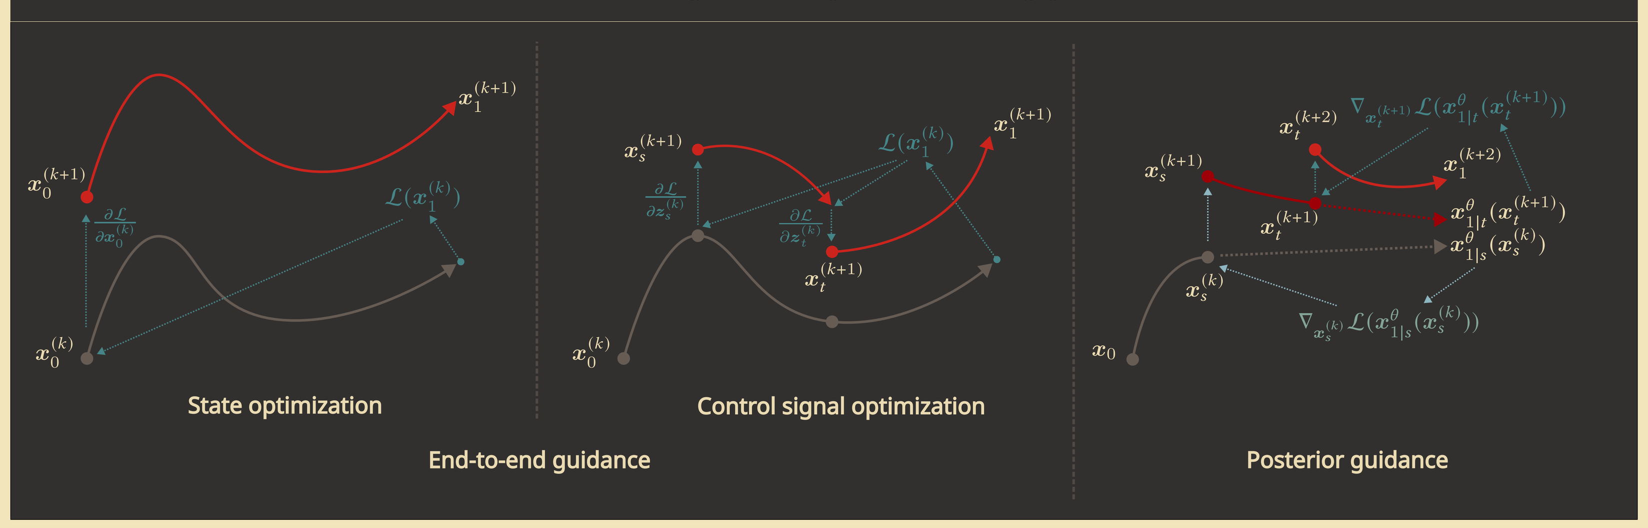

We can view the posterior guidance techinque as a greedy strategy of the end-to-end guidance techinque. In particular, we can view it as a single large Euler step with step size $h = \gamma_1 - \gamma_t$ with $\gamma_t = \alpha_t/\sigma_t$.

Theorem 1 (Greedy as an Euler scheme). For some trajectory state x_t at time t, the greedy gradient given by $\nabla_x\mathcal{L}(x_{1|t}^\theta(x))$ is:

1. a DTO scheme with an explicit Euler step of size $h = \gamma_1 - \gamma_t$, and
2. an OTD scheme with implicit Euler step of size $h = \gamma_1 - \gamma_t$.

Next, we consider how the output of the flow model will change under greedy guidance, *i.e.*,

$$x' = x - \eta\nabla_x\mathcal{L}\left(x_{1|t}^\theta(x)\right). \tag{5}$$

Theorem 2 (Jacobian matrices of affine Gaussian probability paths). For the standard affine Gaussian probability path with flow model $\Phi_{s,t}^\theta(x)$, the Jacobian matrix $\nabla_x\Phi_{s,t}(x)$ as function of x is given as the solution to

$$\nabla_x\Phi_{s,t}^\theta(x) = \frac{\sigma_t}{\sigma_s}I + \sigma_t\int_s^t \dot\gamma_u\frac{\gamma_u}{\sigma_u}\mathrm{Var}_{1|u}(\Phi_{s,u}^\theta(x))\nabla_x\Phi_{s,u}^\theta(x)\, \mathrm{d}u, \tag{6}$$

where

$$\mathrm{Var}_{1|t}(x) = \mathbb{E}_{p_{1|t}(x_1|x)}\left[(x_1 - x_{1|t}^\theta(x))(x_1 - x_{1|t}^\theta(x))^\top\right]. \tag{7}$$

Proposition 3 (Dynamics of greedy gradient guidance). Consider the standard affine Gaussian probability paths model trained to zero loss. The Gateaux differential of x at some time $t \in [0,1]$ in the direction of the gradient $\nabla_x\mathcal{L}\left(x_{1|t}^\theta(x)\right)$ is given by

$$\delta_x^{\mathcal{G}}\Phi_{t,1}^\theta(x) = -\nabla_x\Phi_{t,1}^\theta(x)\nabla_x x_{1|t}^\theta(x)^\top\nabla_{x_1}\mathcal{L}(x_1). \tag{8}$$

Theorem 4 (Greedy convergence). For affine probability paths, if there exists a sequence of states $x_t^{(n)}$ at time t such that it converges to the locally optimal solution $x_{1|t}^\theta(x_t^{(n)}) \to x_1^*$. Then the solution, $\Phi_{t,1}^\theta(x_t^{(n)})$, converges to a neighborhood of size $O(h^2)$ centered at x_1^*.

Inverse problems

Figure 1. Qualitative visualization of using greedy guidance to solve an inverse problem on the task of inpainting with a 70% random mask. Top row is the ground truth, middle row is the measurement, and the bottom row is the reconstruction.

Figure 2. Qualitative visualization of using greedy guidance to solve the HDR inverse problem. Top row is the ground truth, middle row is the measurement, and the bottom row is the reconstruction.

Table 1. A snapshot of the quantitative results for solving the non-linear HDR inverse problem on FFHQ. We report the mean performance (PSNR, SSIM, and LPIPS) across 100 validation images along with the FID. All tasks are using a noisy measurement with noise level $\beta_y = 0.05$.

Method	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
DAPS	27.12	0.752	0.162	42.97
DPS	22.73	0.591	0.264	112.82
RED-diff	22.16	0.512	0.258	108.32
Greedy (Euler)	25.07	0.776	0.173	43.25
Greedy (2-step Euler)	26.32	0.802	0.173	38.64
Greedy (3-step Euler)	27.17	0.820	0.154	36.07
Greedy (4-step Euler)	27.89	0.828	0.151	36.94
Greedy (5-step Euler)	**28.27**	**0.831**	**0.149**	**35.35**

Beyond Euler

What if we take more than an Euler step when performing posterior guidance, perhaps the midpoint method *à la* Moufad et al. [5].

Theorem 5 (Local truncation error of discretize-then-optimize gradients). Let Φ be an explicit Runge-Kutta solver of order $\alpha > 0$ to the ODE

$$x(0) = x_0, \qquad \frac{\mathrm{d}x}{\mathrm{d}t}(t) = u_\theta(t, x(t)), \tag{9}$$

on $[0,T]$ which satisfies the regularity conditions for the Picard-Lindelöf theorem. Let $\Phi_{s,t}^\theta(x)$ denote the flow from s to t, for any $s,t \in [0,T]$ admitted by the ODE. Then,

$$\left\|\nabla_x\Phi_{s,t}^\theta(x) - \nabla_x\Phi_{s,t}(x)\right\| = O(h^{\alpha+1}). \tag{10}$$

Corollary 5.1 (Convergence of a α-th order posterior gradient). For affine probability paths, if there exists a sequence of states $x_t^{(n)}$ at time t such that it converges to the locally optimal solution $\Phi_{t,1}(x_t^{(n)}) \to x_1^*$. Then solution, $\Phi_{t,1}^\theta(x_t^{(n)})$, converges to a neighborhood of size $O(h^{\alpha+1})$ centered at x_1^*.

Guided generation

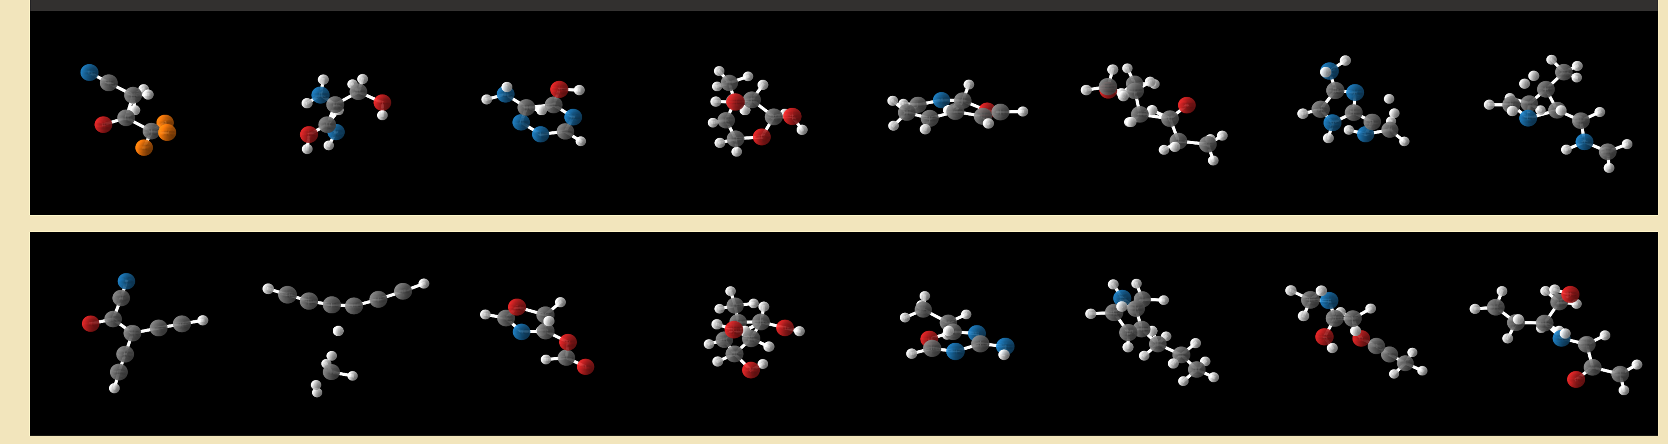

48.43 53.64 60.05 73.95 78.02 82.75 88.45 91.30

Figure 3. Visualization of controlled generated molecules for various polarizability (α) levels. Top row uses a DTO scheme; bottom row uses posterior guidance.

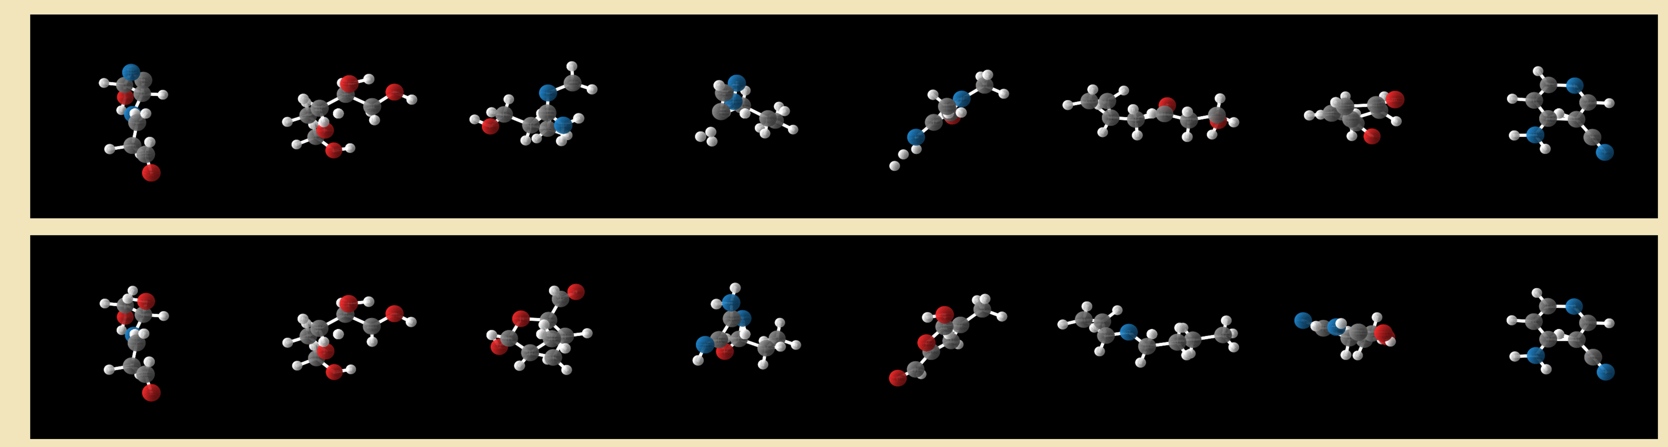

0.014 0.166 1.388 2.322 2.804 3.878 4.922 7.951

Figure 4. Qualitative visualization of controlled generated molecules for various dipole moments (μ). Top row is generated using a end-to-end guidance with a DTO scheme and the bottom row is generated using posterior guidance.

Table 2. Quantitative evaluation of conditional molecule generation on QM9. The MAE is reported for each molecule property (lower is better).

Property	α	$\Delta\varepsilon$	$\varepsilon_{\mathrm{HOMO}}$	$\varepsilon_{\mathrm{LUMO}}$	μ	C_v
Unit	Bohr^2	meV	meV	meV	D	$\frac{\mathrm{cal}}{\mathrm{K \cdot mol}}$
Greedy (Euler)	11.282	1265	725	1092	1.559	6.469
Greedy (midpoint)	5.313	1196	599	1057	1.417	2.967
Greedy (2-step Euler)	5.667	1205	695	1222	1.491	2.767
Greedy (3-step Euler)	5.098	1152	600	1152	1.384	3.229
Greedy (5-step Euler)	4.177	1083	571	939	1.328	2.332
DTO (1-step)	13.049	989×10^{12}	681	86.512	1.666	15.144
DTO (2-step)	6.113	1359	666	1199	1.533	3.757
DTO (4-step)	6.115	1294	668	1190	1.406	2.829
DTO (8-step)	4.549	1070	608	1078	1.247	2.594
DTO (16-step)	3.454	817	608	939	1.177	2.003
DTO (32-step)	2.912	750	410	666	0.721	1.566
DTO (50-step)	1.404	401	176	373	0.372	0.866
EquiFM	9.525	1494	622	1523	1.628	6.689
Lower bound	0.10	64	39	46	0.043	0.040

References

[1] H. Ben-Hamu, O. Puny, I. Gat, B. Karrer, U. Singer, and Y. Lipman. "D-Flow: Differentiating through Flows for Controlled Generation". In: *Forty-first International Conference on Machine Learning*. 2024. URL: https://openreview.net/forum?id=SE20BFqj6J.

[2] Blasingame and C. Liu. "AdjointDEIS: Efficient Gradients for Diffusion Models". In: *Advances in Neural Information Processing Systems*. Ed. by A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang. Vol. 37. Curran Associates, Inc., 2024, pp. 2449–2483. URL: https://proceedings.neurips.cc/paper_files/paper/2024/file/04badd3b048315c8c3a0ca17eff723d7-Paper-Conference.pdf.

[3] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. "Diffusion Posterior Sampling for General Noisy Inverse Problems". In: *The Eleventh International Conference on Learning Representations, ICLR 2023*. The International Conference on Learning Representations. 2023.

[4] P. Kidger. "On Neural Differential Equations". Available at https://arxiv.org/abs/2202.02435. Ph.D. thesis. Oxford University, 2022.

[5] B. Moufad, Y. Janati, L. Bedin, A. O. Durmus, R. Douc, E. Moulines, and J. Olsson. "Variational Diffusion Posterior Sampling with Midpoint Guidance". In: *The Thirteenth International Conference on Learning Representations*. 2025. URL: https://openreview.net/forum?id=6EUtjXAvmj.