Greed is Good: Guided Generation from a Greedy Perspective

Zander W. Blasingame   Chen Liu

Clarkson University

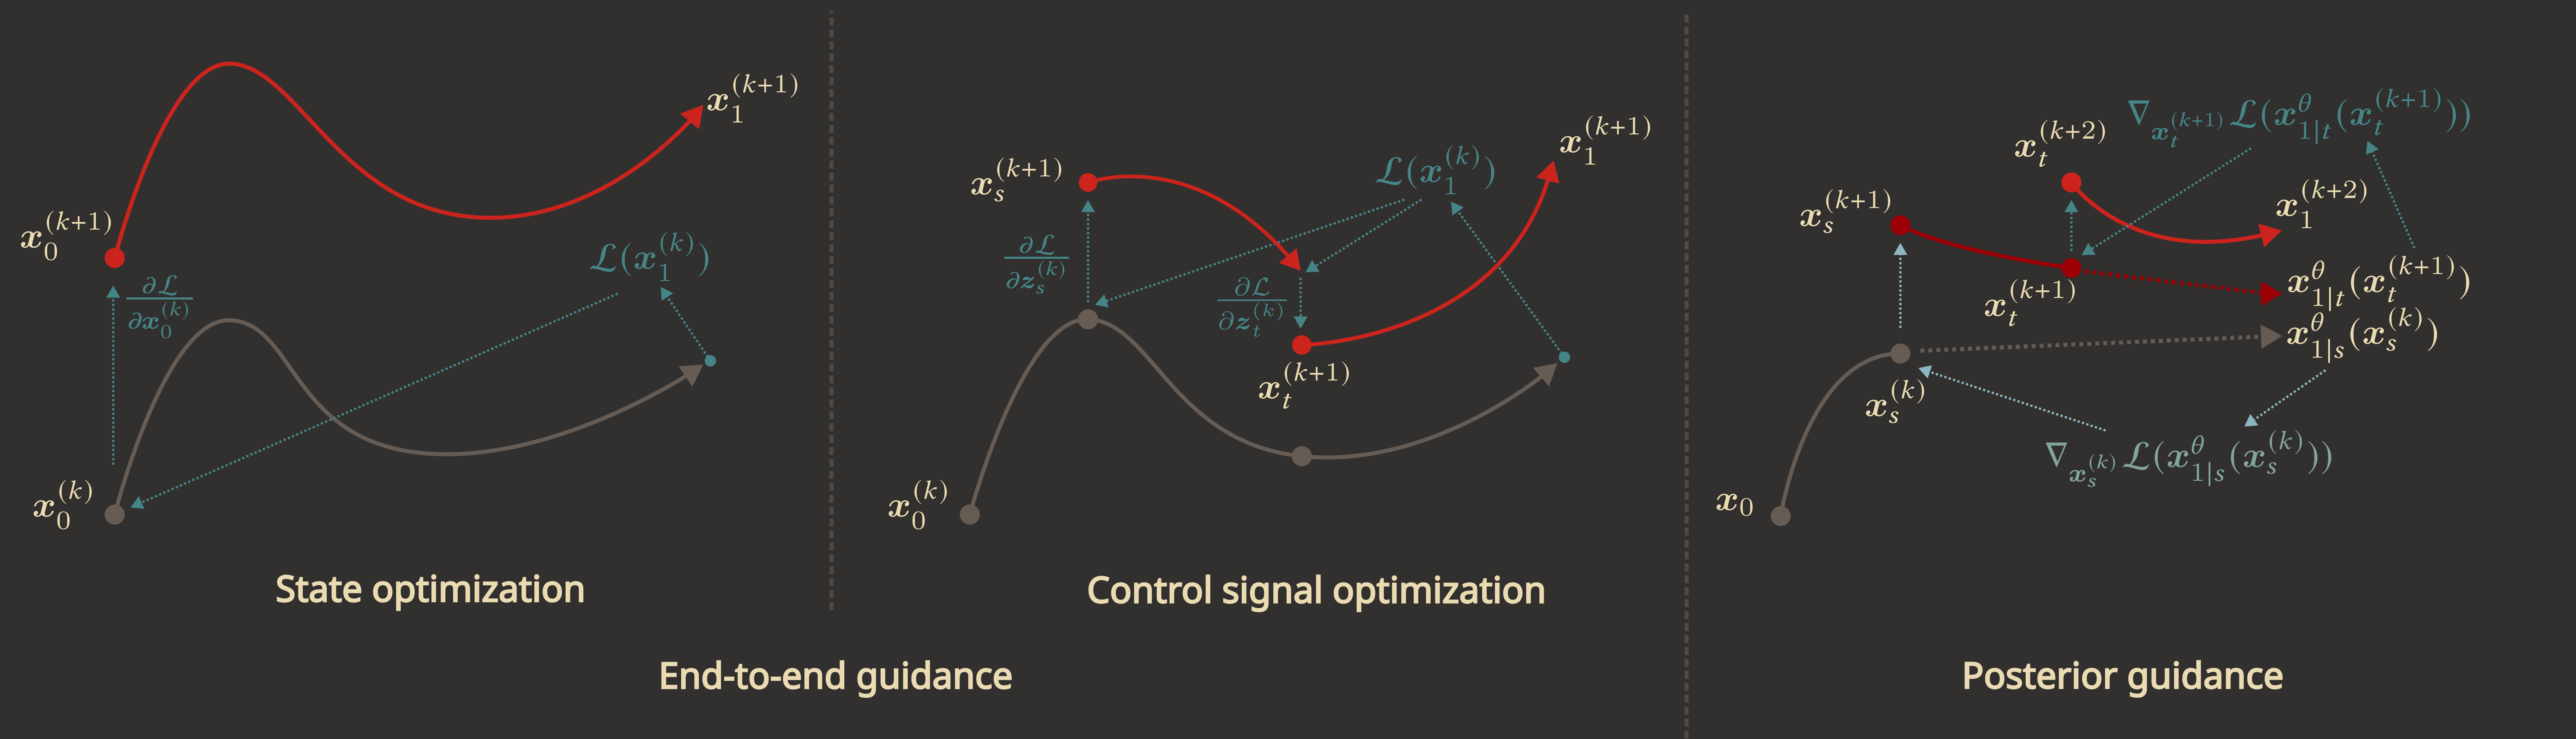

Motivation

Consider the usual *flow model*, let $(\boldsymbol{X}_0, \boldsymbol{X}_1) \sim p(\boldsymbol{x}_0)q(\boldsymbol{x}_1)$ where $q(\boldsymbol{x})$ is the target distribution and $p(\boldsymbol{x}_0)$ is the prior distribution. Define $\boldsymbol{X}_t$ as $\boldsymbol{X}_t = \alpha_t \boldsymbol{X}_1 + \sigma_t \boldsymbol{X}_0$, for schedule (α_t, σ_t). Then, the vector field of the *affine conditional flow* $\Phi_t(\boldsymbol{x}|\boldsymbol{x}_1) = \alpha_t \boldsymbol{x}_1 + \sigma_t \boldsymbol{x}$ is given by

$$\boldsymbol{u}_t(\boldsymbol{x}) = \mathbb{E}[\dot{\alpha}_t \boldsymbol{X}_1 + \dot{\sigma}_t \boldsymbol{X}_0 | \boldsymbol{X}_t = \boldsymbol{x}]. \tag{1}$$

Assume that $\boldsymbol{u}_t^\theta$ is trained to zero loss, so $\boldsymbol{u}_t^\theta = \boldsymbol{u}_t$.

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

$$\min_{\boldsymbol{x}_0} \mathcal{L}\left(\boldsymbol{x}_0 + \int_0^1 \boldsymbol{u}_\tau^\theta(\boldsymbol{x}_\tau)\,\mathrm{d}\tau\right). \tag{2}$$

Posterior guidance. We can use the *gradient* of the denoiser $\boldsymbol{x}_{1|t}^\theta(\boldsymbol{x}) = \mathbb{E}[\boldsymbol{X}_1 | \boldsymbol{X}_t = \boldsymbol{x}]$ for guidance [3], *i.e.*, for some iteration $\boldsymbol{x}_n$ in the numerical scheme

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

End-to-end guidance. Alternatively, optimize the initial point $\boldsymbol{x}_0$ [1, 2], *i.e.*,

$$\boldsymbol{x}_0^{(k+1)} = \boldsymbol{x}_0^{(k)} - \eta \nabla \mathcal{L}\left(\Phi_{0,1}^\theta(\boldsymbol{x}_0^{(k)})\right), \tag{4}$$

where $\Phi_{0,1}^\theta$ is the flow map from 0 to 1 induced by $\boldsymbol{u}_t^\theta$. 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 $\boldsymbol{x}_t$ at time t, the greedy gradient given by $\nabla_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{x}_{1|t}^\theta(\boldsymbol{x}))$ 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.*,

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

Proposition 2 (Dynamics of greedy gradient guidance). Consider the standard affine Gaussian probability paths model trained to zero loss. The Gateaux differential of $\boldsymbol{x}$ at some time $t \in [0, 1]$ in the direction of the gradient $\nabla_{\boldsymbol{x}} \mathcal{L}\left(\boldsymbol{x}_{1|t}^\theta(\boldsymbol{x})\right)$ is given by

$$\delta_{\boldsymbol{x}}^{\mathcal{G}} \Phi_{t,1}^\theta(\boldsymbol{x}) = -\nabla_{\boldsymbol{x}} \Phi_{t,1}^\theta(\boldsymbol{x}) \nabla_{\boldsymbol{x}} \boldsymbol{x}_{1|t}^\theta(\boldsymbol{x})^\top \nabla_{\boldsymbol{x}_1} \mathcal{L}(\boldsymbol{x}_1). \tag{6}$$

Theorem 3 (Greedy convergence). For affine probability paths, if there exists a sequence of states $\boldsymbol{x}_t^{(n)}$ at time t such that it converges to the locally optimal solution $\boldsymbol{x}_{1|t}^\theta(\boldsymbol{x}_t^{(n)}) \to \boldsymbol{x}_1^*$. Then the solution, $\Phi_{t,1}^\theta(\boldsymbol{x}_t^{(n)})$, converges to a neighborhood of size $O(h^2)$ centered at $\boldsymbol{x}_1^*$.

Beyond Euler

What if we take more than an Euler step when performing posterior guidance, perhaps the midpoint method?

Theorem 4 (Truncation error of single-step gradients). Let Φ be an explict Runge-Kutta solver of order $\alpha > 0$ of a flow model with flow $\Phi_{s,t}^\theta(\boldsymbol{x})$. Then for any $t \in [0, 1]$,

$$\left\| \nabla_{\boldsymbol{x}} \Phi_{t,1}^\theta(\boldsymbol{x}) - \nabla_{\boldsymbol{x}} \Phi_{t,1}(\boldsymbol{x}) \right\| = O(h^{\alpha+1}), \tag{7}$$

where $h = 1 - t$.

Corollary 4.1 (Convergence of a α-th order posterior gradient). For affine probability paths, if there exists a sequence of states $\boldsymbol{x}_t^{(n)}$ at time t such that it converges to the locally optimal solution $\Phi_{t,1}(\boldsymbol{x}_t^{(n)}) \to \boldsymbol{x}_1^*$. Then solution, $\Phi_{t,1}^\theta(\boldsymbol{x}_t^{(n)})$, converges to a neighborhood of size $O(h^{\alpha+1})$ centered at $\boldsymbol{x}_1^*$.

Numerical experiments

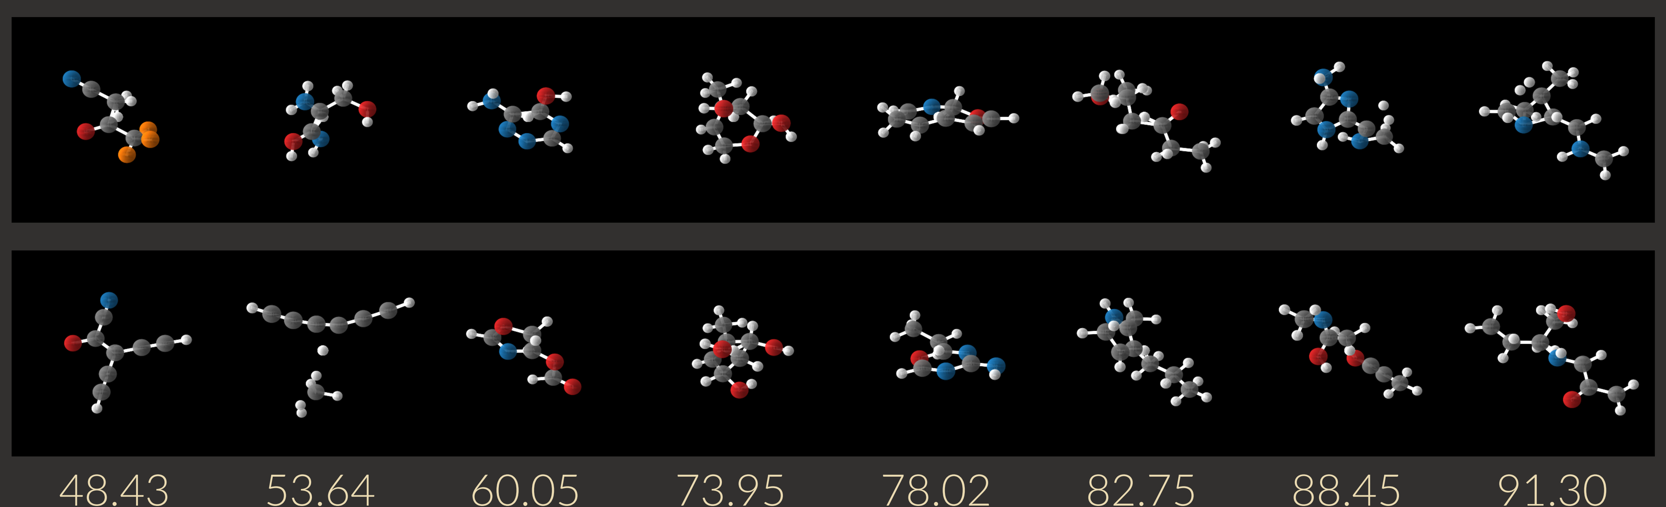

Figure 1. Visualization of controlled generated molecules for various polarizability (α) levels. Top row uses a DTO scheme; bottom row uses posterior guidance.

Table 1. Quantitative evaluation of conditional molecule generation. The MAE is reported for each molecule property (lower is better).

Property Unit	α Bohr^2	$\Delta\varepsilon$ meV	$\varepsilon_{\mathrm{HOMO}}$ meV	$\varepsilon_{\mathrm{LUMO}}$ meV	μ D	C_v $\frac{\mathrm{cal}}{\mathrm{K \cdot mol}}$
Greedy (Euler)	11.282	1265	725	1092	1.559	6.469
Greedy (2-step Euler)	5.377	1275	560	1204	1.563	2.975
Greedy (midpoint)	5.313	1196	599	1057	1.417	2.967
DTO	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.

Correspondence to blasinzw@clarkson.edu