
Learning Boltzmann Generators via Constrained Mass Transport

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Efficient sampling from high-dimensional and multimodal unnormalized proba-
2 bility distributions is a central challenge in many areas of science and machine
3 learning. We focus on Boltzmann generators (BGs) that aim to sample the Boltz-
4 mann distribution of physical systems, such as molecules, at a given temperature.
5 Classical variational approaches that minimize the reverse Kullback–Leibler di-
6 vergence are prone to mode collapse, while annealing-based methods, commonly
7 using geometric schedules, can suffer from mass teleportation and rely heavily on
8 schedule tuning. We introduce *Constrained Mass Transport* (CMT), a variational
9 framework that generates intermediate distributions under constraints on both the
10 KL divergence and the entropy decay between successive steps. These constraints
11 enhance distributional overlap, mitigate mass teleportation, and counteract prema-
12 ture convergence. Across standard BG benchmarks and the here introduced *ELIL*
13 *tetrapeptide*, the largest system studied without access to samples from molecular
14 dynamics, CMT consistently surpasses state-of-the-art variational methods, achiev-
15 ing more than 2.5× higher effective sample size while avoiding mode collapse.

16 1 Introduction

17 We consider the problem of sampling from a target probability measure $p \in \mathcal{P}(\mathbb{R}^d)$ given by
18 $p(x) = \tilde{p}(x)/\mathcal{Z}$ where $\tilde{p} : \mathbb{R}^d \rightarrow (0, \infty)$ can be evaluated pointwise but the normalization constant
19 $\mathcal{Z} = \int_{\mathbb{R}^d} \tilde{p}(x) dx$ is intractable. Sampling from unnormalized densities arises in many areas, including
20 Bayesian statistics [35], reinforcement learning [20], and the natural sciences [72]. A prominent
21 example is learning Boltzmann generators (BGs) [56], for which $\tilde{p}(x) = \exp(-E(x)/k_B T)$, with
22 E being an energy function, T the temperature, and k_B the Boltzmann constant. BGs enable
23 efficient sampling of thermodynamic ensembles, thereby bypassing costly molecular dynamics (MD)
24 simulations and accelerating the exploration of rare but physically important states. However, learning
25 BGs is challenging as the state space is typically high-dimensional, the target distribution is often
26 highly multimodal, and evaluating $E(x)$ can be very costly, especially when using accurate energies
27 such as those from density-functional theory [12].

28 A promising alternative to classical MD or Monte Carlo methods [39] is offered by variational
29 approaches [73], which aim to minimize a statistical divergence between a variational probability
30 measure $q \in \mathcal{P}(\mathbb{R}^d)$ and the target p , commonly the reverse Kullback–Leibler (KL) divergence

$$q^* = \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q \| p), \quad (1)$$

31 whose unique minimizer is $q^* = p$. Directly minimizing the reverse KL divergence tends to suffer
32 from mode collapse, ignoring low-probability modes of the target [15]. To counteract this, a number
33 of recent approaches have proposed to construct a sequence of intermediate distributions that transport
34 probability mass from a tractable base distribution q_0 to the target. This idea, which dates back
35 several decades to annealed importance sampling [54], is most often realized through a geometric
36 annealing path, which is defined as a sequence of $(q_i)_{i=1}^I$ which follows $q_i \propto q_0^{1-\beta_i} \tilde{p}^{\beta_i}$ where the

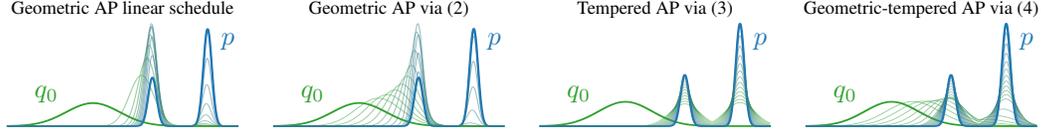


Figure 1: Illustration of the annealing paths (AP) obtained by solving the variational problems (2), (3), or (4).

37 corresponding annealing schedule $(\beta_i)_{i=1}^I$ ensures that $q_I = p$. Despite its simplicity, geometric
 38 annealing can suffer from mass teleportation, where large portions of the probability mass shift to
 39 disjoint regions between successive steps, complicating mass transport. Moreover, its performance
 40 critically depends on the choice of annealing schedule [75].

41 A recent study proposed selecting the annealing schedule using trust-region constraints that bound the
 42 KL divergence between successive distributions, thereby avoiding abrupt changes in the stochastic
 43 optimal control setting [14]. Adapting this idea to measures that admit densities, we further introduce
 44 a constraint that explicitly controls the rate at which the entropy of the variational distribution
 45 decreases along the transport path. This added degree of freedom enables deviations from the
 46 standard geometric annealing schedule, mitigating issues such as mass teleportation and premature
 47 convergence, while fostering greater overlap between consecutive distributions.

48 We demonstrate that our method, *Constrained Mass Transport* (CMT), consistently outperforms state-
 49 of-the-art approaches—often by a substantial margin—when learning Boltzmann generators solely
 50 from energy evaluations, without relying on additional MD samples. Furthermore, we introduce a
 51 new benchmark system, *ELIL tetrapeptide*, which, to the best of our knowledge, is the largest system
 52 studied to date in the setting of learning exclusively from energy evaluations. On this challenging
 53 system, CMT achieves a $2.5\times$ larger effective sample size compared to baselines, while avoiding
 54 mode collapse.

55 2 Constrained mass transport

56 Let $\mathcal{P}(\mathbb{R}^d)$ denote the space of probability measures on \mathbb{R}^d that are absolutely continuous with
 57 respect to Lebesgue measure and admit smooth densities. Here, we approach the sampling problem
 58 by dividing the problem (1) into a variety of constrained optimization problems that result in annealing
 59 paths of intermediate densities $(q_i)_{i=0}^I$ that bridge between a tractable prior q_0 and the target p .

60 To mitigate premature convergence, we employ two constraints: a trust-region constraint and an
 61 entropy constraint. We then combine both constraints into a unified approach.

62 **Trust-region constraint.** Trust-regions aim at dividing the problem (1) into subproblems by con-
 63 straining the updated density to be close to the old density in terms of KL divergence. Formally, this
 64 is given by the iterative optimization scheme ¹

$$q_{i+1} = \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q|p) \quad \text{s.t.} \quad D_{\text{KL}}(q|q_i) \leq \varepsilon_{\text{tr}}, \quad (2)$$

65 for $i \in \mathbb{N}$, trust region bound $\varepsilon_{\text{tr}} > 0$ and some $q_0 \in \mathcal{P}(\mathbb{R}^d)$. Due to the convexity of the KL
 66 divergence, we can show that in all but the last step we actually have an equality constraint in (2); see
 67 App. A.

68 **Entropy constraint.** In a similar fashion to (2), we can avoid premature convergence by regulating
 69 the entropy decay of the model by constructing a sequence of intermediate densities whose change in
 70 entropy is constrained. Formally, we aim to solve the following problem

$$q_{i+1} = \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q|p) \quad \text{s.t.} \quad H(q_i) - H(q) \leq \varepsilon_{\text{ent}}, \quad (3)$$

71 where $H(q) = -\int q(x) \log q(x) dx$ is the Shannon entropy and $\varepsilon_{\text{ent}} > 0$ the entropy bound.

72 Despite the potential of (3) for counteracting premature convergence, the KL divergence between q_0
 73 and q_1 can be arbitrarily large and therefore could cause instabilities due to a lack of overlap between
 74 the successive densities. This observation motivates a combined strategy, which incorporates both
 75 constraints into a single iterative optimization scheme.

¹To ensure that $q \in \mathcal{P}(\mathbb{R}^d)$ we need an additional constraint $\int q(x) dx = 1$ which we omitted in the main part of the paper for readability.

Algorithm 1 Constrained mass transport

Require: Initial measure q_0 , target measure p , divergence D , approximation family \mathcal{Q} , buffer size N
for $i \leftarrow 0, \dots, I - 1$ **do**
 Draw N samples $x_n \sim q_i$, evaluate $q_i(x_n), \tilde{p}(x_n)$ and initialize buffer $\mathcal{B}^{(i)} = (x_n, q_i(x_n), \tilde{p}(x_n))_{n=1}^N$
 Compute multiplier $\lambda_i, \eta_i = \arg \max_{\lambda, \eta \in \mathbb{R}^+} g_{\text{tr-ent}}^{(i+1)}(\lambda, \eta)$ using $\mathcal{B}^{(i)}$
 Compute $q_{i+1} \approx \hat{q}_{i+1} = \arg \min_{q \in \mathcal{Q}} D(q_{i+1}, q)$ using $\mathcal{B}^{(i)}$
return $\hat{q}_I \approx p$

76 **Combination of both constraints.** We can straightforwardly combine the constraints in (2) and (3)
77 into a single iterative optimization scheme defined as

$$q_{i+1} = \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q|p) \quad \text{s.t.} \quad \begin{cases} D_{\text{KL}}(q|q_i) \leq \varepsilon_{\text{tr}}, \\ H(q_i) - H(q) \leq \varepsilon_{\text{ent}}. \end{cases} \quad (4)$$

78 We introduce Lagrangian multiplier λ and η for the trust region and entropy constraint, respectively.
79 Indeed, one can obtain an analytical expression for the evolution of the optimal densities, see Prop. 2.1
80 and App. A for a proof.

81 **Proposition 2.1** (Optimal intermediate densities for entropy and trust region constraint). *Let $\tilde{p} \in$*
82 *$\mathcal{P}(\mathbb{R}^d)$ be the unnormalized target density. The intermediate optimal densities that solve (3) satisfy*

$$q_{i+1}(x, \lambda, \eta) = \frac{q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} \tilde{p}(x)^{\frac{1}{1+\lambda+\eta}}}{\mathcal{Z}_{i+1}(\lambda, \eta)} \quad \text{with} \quad \mathcal{Z}_{i+1}(\lambda, \eta) = \int q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} \tilde{p}(x)^{\frac{1}{1+\lambda+\eta}}(x) dx, \quad (5)$$

83 where q_{i+1} are the unique optima of the Lagrangian corresponding to (4).

84 Clearly, the trust-region constraint ensures that the KL divergence between q_0 and q_1 is at most ε_{tr}
85 and, therefore, for a suitable choice of ε_{tr} ensures that two consecutive densities have sufficient
86 overlap. Lastly, the specific Lagrangian multipliers λ_i, η_i that define $q_{i+1}(x)$, can be obtained by
87 solving the convex optimization problem

$$(\lambda_i, \eta_i) = \arg \max_{\lambda, \eta \in \mathbb{R}_{\geq 0}} -g_{\text{tr-ent}}^{(i+1)}(\lambda, \eta), \quad (6)$$

88 where $g_{\text{tr-ent}}^{(i+1)}(\lambda, \eta)$ is the Lagrangian dual function corresponding to (4), defined as,

$$g_{\text{tr-ent}}^{(i+1)}(\lambda, \eta) := -(1 + \lambda + \eta) \log \mathcal{Z}_{i+1}(\lambda, \eta) - \lambda \varepsilon_{\text{tr}} - \eta(H(q_i) - \varepsilon_{\text{ent}}). \quad (7)$$

89 Solving (6) can be done efficiently in practice; see Sec. 3 for additional details.

90 **Connection to annealing paths.** Iteratively solving (4) induces an *annealing path*, that is, a sequence
91 of densities $(q_i)_{i \in \mathbb{N}}$ that interpolates between q_0 and p . The problems (2) and (3) can be considered
92 special cases of (4), each yielding their own annealing path and respective one-dimensional dual
93 function. We characterize all annealing paths in Thm. 2.2; See App. A for a proof; The different
94 annealing paths are illustrated in App. A.

95 **Theorem 2.2** (Annealing paths). *Let $\tilde{p} \in \mathcal{P}(\mathbb{R}^d)$ be the unnormalized target density and $q_0 \in \mathcal{P}(\mathbb{R}^d)$*
96 *some initial density. The intermediate optimal densities $q_i \in \mathcal{P}(\mathbb{R}^d)$, $i \in \mathbb{N}_0$, that solve (2), (3) and*
97 *(4) satisfy*

$$q_i \propto q_0^{1-\beta_i} \tilde{p}^{\beta_i}, \quad q_i \propto \tilde{p}^{\alpha_i} \quad (i \geq 1), \quad \text{and} \quad q_i \propto q_0^{1-\beta_i} (\tilde{p}^{\alpha_i})^{\beta_i}, \quad (8)$$

98 respectively, with β and α being functions of the corresponding Lagrangian multiplier. Moreover,
99 the sequences $(\alpha_i)_{i \in \mathbb{N}_0}$ and $(\beta_i)_{i \in \mathbb{N}_0}$ take values in $[0, 1]$, satisfy $\alpha_0 = \beta_0 = 0$ and $\alpha_I = \beta_I = 1$ for
100 some $I \in \mathbb{N}_+$ and $(\beta_i)_{i \in \mathbb{N}_0}$ is monotonically increasing.

101 3 Learning the intermediate densities

102 So far, we discussed how one can construct a sequence of intermediate measures $(q_i)_{i \in \mathbb{N}}$ using our
103 constrained mass transport formulation. However, despite having access to the analytical form of
104 q_i , it is typically not possible to sample from it directly. As such, one can approximate each q_i by a
105 distribution from a tractable class $\mathcal{Q} \subset \mathcal{P}(\mathbb{R}^d)$ that permits efficient sampling and density evaluation.
106 Given an approximation family \mathcal{Q} , we select $\hat{q}_i \in \mathcal{Q}$ to approximate q_i by solving

$$\hat{q}_i = \arg \min_{q \in \mathcal{Q}} D(q_i, q), \quad (9)$$

107 where D is an arbitrary statistical divergence between probability measures. This formulation is
108 general: the choice of \mathcal{Q} and D determines the trade-off between expressivity, computational cost,
109 and statistical properties such as mode coverage or robustness.

SYSTEM	METHOD	TARGET EVALS ↓	NLL ↓	ESS [%] ↑
ALANINE DIPEPTIDE ($d = 60$)	FORWARD KL	5×10^9	-213.581 ± 0.000	$(82.16 \pm 0.09) \%$
	REVERSE KL	2.56×10^8	-213.609 ± 0.006	$(94.11 \pm 0.21) \%$
	FAB	2.13×10^8	-213.653 ± 0.000	$(94.81 \pm 0.04) \%$
	TA-BG	1×10^8	-213.666 ± 0.001	$(95.77 \pm 0.12) \%$
	CMT (OURS)	1×10^8	-213.677 ± 0.000	$(97.69 \pm 0.03) \%$
ALANINE TETRA- PEPTIDE ($d = 120$)	FORWARD KL	4.2×10^9	-330.069 ± 0.001	$(45.29 \pm 0.08) \%$
	REVERSE KL	2.56×10^8	-329.191 ± 0.122	$(74.88 \pm 3.65) \%$
	FAB	2.13×10^8	-330.100 ± 0.002	$(63.59 \pm 0.23) \%$
	TA-BG	1×10^8	-330.143 ± 0.002	$(64.87 \pm 0.21) \%$
	CMT (OURS)	1×10^8	-330.167 ± 0.002	$(68.49 \pm 0.14) \%$
ALANINE HEXA- PEPTIDE ($d = 180$)	FORWARD KL	4.2×10^9	-501.598 ± 0.005	$(10.97 \pm 0.11) \%$
	REVERSE KL	2.56×10^8	-497.378 ± 0.277	$(22.22 \pm 1.44) \%$
	FAB	4.2×10^8	-501.268 ± 0.008	$(14.64 \pm 0.08) \%$
	TA-BG	4×10^8	-501.582 ± 0.010	$(15.89 \pm 0.13) \%$
	CMT (OURS)	4×10^8	-501.761 ± 0.008	$(29.20 \pm 0.24) \%$
ELIL TETRA- PEPTIDE ($d = 219$)	FORWARD KL	4.2×10^9	-597.571 ± 0.004	$(5.80 \pm 0.04) \%$
	REVERSE KL	2.56×10^8	-583.381 ± 3.033	$(1.25 \pm 0.45) \%$
	FAB	8.43×10^8	-597.370 ± 0.006	$(7.30 \pm 0.08) \%$
	TA-BG	8×10^8	-597.830 ± 0.022	$(10.12 \pm 0.38) \%$
	CMT (OURS)	8×10^8	-598.440 ± 0.003	$(25.91 \pm 0.19) \%$

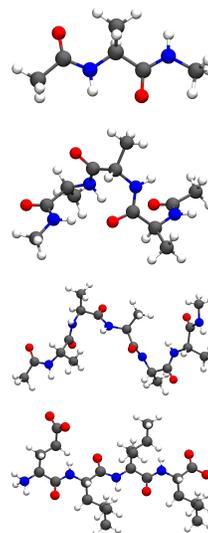


Table 1: Results for all systems of varying dimensionality d . Evaluation criteria include the number of target evaluations (*target evals*), the negative log-likelihood (*NLL*), and the reverse effective sample size (*ESS*). Each value is shown as the mean \pm standard error over four independent runs. The best results are highlighted in bold, except for the reverse KL, which is prone to mode collapse, making ESS values not directly comparable.

110 Because q_{i+1} is available in closed form from, the constrained transport updates (see Prop. 2.1)
 111 and importance weights $q_{i+1}(x)/q_i(x)$ can be computed solely from q_i and \tilde{p} . This admits efficient
 112 estimation of Z_{i+1} and the Lagrangian multipliers λ_i, η_i as well as sample reuse. Lastly, the trust-
 113 region constraint controls the variance of the importance weights, keeping it approximately constant,
 114 independent of the problem dimension d (see App. A), resulting in a highly scalable algorithm.

115 Details of the dual optimization procedure are provided in App. D.3, including a code example.
 116 Lastly, we refer to Algorithm 1 for an algorithmic overview of the constrained measure transport
 117 method.

118 4 Experiments

119 To approximate the intermediate distributions q_i , we select the forward Kullback–Leibler (KL)
 120 divergence as the divergence measure D , as it strongly penalizes underestimation of the support of
 121 q_i . This choice promotes mode coverage and mitigates the risk of mode collapse. The variational
 122 family \mathcal{Q} is instantiated as neural spline flows [31] with internal coordinates. CMT yields substantial
 123 improvements in performance across all considered molecular systems, as shown in Tab. 1. Additional
 124 experimental results and details on evaluation metrics are provided in Apps. C and D.4.

125 5 Conclusion

126 We have introduced *Constrained Mass Transport* (CMT), a variational framework for constructing
 127 intermediate distributions that transport probability mass from a tractable base measure to a complex,
 128 unnormalized target. By enforcing constraints on both the KL divergence and the entropy decay
 129 between successive steps, CMT balances exploration and convergence, thereby mitigating mass
 130 teleportation, reducing mode collapse, and promoting smooth distributional overlap. Our empir-
 131 ical evaluation across established Boltzmann generator benchmarks and the here proposed *ELIL*
 132 *tetrapeptide*, learned purely from energy evaluations without access to molecular dynamics samples,
 133 demonstrates that CMT consistently outperforms existing annealing-based and variational baselines,
 134 achieving over $2.5\times$ higher effective sample size while preserving mode diversity.

135 Promising directions for future work include exploring alternative approximation families \mathcal{Q} and
 136 divergences D for learning intermediate densities—for example, the log-variance divergence [34, 64],
 137 which may yield further performance gains. Another interesting avenue is to apply our method in
 138 Cartesian coordinate representations, which enables transferability across different molecular systems
 139 [45].

References

- 140
141 [1] A. Abdolmaleki, R. Lioutikov, J. R. Peters, N. Lau, L. Pualo Reis, and G. Neumann. Model-
142 based relative entropy stochastic search. *Advances in Neural Information Processing Systems*,
143 28, 2015.
- 144 [2] A. Abdolmaleki, J. T. Springenberg, J. Degraeve, S. Bohez, Y. Tassa, D. Belov, N. Heess, and
145 M. Riedmiller. Relative entropy regularized policy iteration. *arXiv preprint arXiv:1812.02256*,
146 2018.
- 147 [3] A. Abdolmaleki, J. T. Springenberg, Y. Tassa, R. Munos, N. Heess, and M. Riedmiller. Maxi-
148 mum a posteriori policy optimisation. *arXiv preprint arXiv:1806.06920*, 2018.
- 149 [4] J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In *International*
150 *conference on machine learning*, pages 22–31. PMLR, 2017.
- 151 [5] Z. Ahmed, N. Le Roux, M. Norouzi, and D. Schuurmans. Understanding the impact of entropy
152 on policy optimization. In *International conference on machine learning*, pages 151–160.
153 PMLR, 2019.
- 154 [6] T. Akhoun-Sadegh, J. Lee, J. Bose, V. D. Bortoli, A. Doucet, M. M. Bronstein, D. Beaini,
155 S. Ravanbakhsh, K. Neklyudov, and A. Tong. Progressive inference-time annealing of diffusion
156 models for sampling from boltzmann densities. In *ICML 2025 Generative AI and Biology*
157 *(GenBio) Workshop*, 2025.
- 158 [7] R. Akrou, A. Abdolmaleki, H. Abdulsamad, J. Peters, and G. Neumann. Model-free trajectory-
159 based policy optimization with monotonic improvement. *Journal of machine learning research*,
160 19(14):1–25, 2018.
- 161 [8] R. Akrou, G. Neumann, H. Abdulsamad, and A. Abdolmaleki. Model-free trajectory opti-
162 mization for reinforcement learning. In *International Conference on Machine Learning*, pages
163 2961–2970. PMLR, 2016.
- 164 [9] R. Akrou, J. Pajarinen, J. Peters, and G. Neumann. Projections for approximate policy iteration
165 algorithms. In *International Conference on Machine Learning*, pages 181–190. PMLR, 2019.
- 166 [10] O. Arenz, P. Dahlinger, Z. Ye, M. Volpp, and G. Neumann. A unified perspective on natural
167 gradient variational inference with gaussian mixture models. *arXiv preprint arXiv:2209.11533*,
168 2022.
- 169 [11] O. Arenz, M. Zhong, and G. Neumann. Trust-region variational inference with gaussian mixture
170 models. *Journal of Machine Learning Research*, 21(163):1–60, 2020.
- 171 [12] N. Argaman and G. Makov. Density functional theory: An introduction. *American Journal of*
172 *Physics*, 68(1):69–79, Jan. 2000.
- 173 [13] R. Benita, M. Elad, and J. Keshet. Spectral analysis of diffusion models with application to
174 schedule design. *arXiv preprint arXiv:2502.00180*, 2025.
- 175 [14] D. Blessing, J. Berner, L. Richter, C. Domingo-Enrich, Y. Du, A. Vahdat, and G. Neumann.
176 Trust region constrained measure transport in path space for stochastic optimal control and
177 inference, 2025.
- 178 [15] D. Blessing, X. Jia, J. Esslinger, F. Vargas, and G. Neumann. Beyond ELBOs: A large-scale
179 evaluation of variational methods for sampling. In R. Salakhutdinov, Z. Kolter, K. Heller,
180 A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *Proceedings of the 41st Inter-*
181 *national Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning*
182 *Research*, pages 4205–4229. PMLR, 21–27 Jul 2024.
- 183 [16] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge,
184 UK, 2004. Seventh printing with corrections, 2009.
- 185 [17] R. Brekelmans, V. Masrani, T. Bui, F. Wood, A. Galstyan, G. V. Steeg, and F. Nielsen. Annealed
186 importance sampling with q-paths. *arXiv preprint arXiv:2012.07823*, 2020.

- 187 [18] R. Brekelmans and F. Nielsen. Variational representations of annealing paths: Bregman
188 information under monotonic embedding. *Information Geometry*, 7(1):193–228, 2024.
- 189 [19] R. P. Brent. *Algorithms for minimization without derivatives*. Courier Corporation, 2013.
- 190 [20] O. Celik, Z. Li, D. Blessing, G. Li, D. Palenicek, J. Peters, G. Chaltatzaki, and G. Neu-
191 mann. Dime: Diffusion-based maximum entropy reinforcement learning. *arXiv preprint*
192 *arXiv:2502.02316*, 2025.
- 193 [21] J. Chen, D. Lu, Z. Xiu, K. Bai, L. Carin, and C. Tao. Variational inference with holder bounds.
194 *arXiv preprint arXiv:2111.02947*, 2021.
- 195 [22] J. Choi, Y. Chen, M. Tao, and G.-H. Liu. Non-equilibrium annealed adjoint sampler. *arXiv*
196 *preprint arXiv:2506.18165*, 2025.
- 197 [23] N. Chopin, F. R. Crucinio, and A. Korba. A connection between tempering and entropic mirror
198 descent. *arXiv preprint arXiv:2310.11914*, 2023.
- 199 [24] N. Chopin, O. Papaspiliopoulos, et al. *An introduction to sequential Monte Carlo*, volume 4.
200 Springer, 2020.
- 201 [25] A. R. Conn, N. I. Gould, and P. L. Toint. *Trust region methods*. SIAM, 2000.
- 202 [26] T. M. Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- 203 [27] P. Del Moral, A. Doucet, and A. Jasra. Sequential monte carlo samplers. *Journal of the Royal*
204 *Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.
- 205 [28] M. Dibak, L. Klein, A. Krämer, and F. Noé. Temperature steerable flows and boltzmann
206 generators. *Physical Review Research*, 4(4):L042005, 2022.
- 207 [29] M. Dibak, L. Klein, A. Krämer, and F. Noé. Temperature steerable flows and Boltzmann
208 generators. *Physical Review Research*, 4(4):L042005, Oct. 2022.
- 209 [30] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters*
210 *B*, 195(2):216–222, Sept. 1987.
- 211 [31] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios. Neural spline flows. *Advances in*
212 *neural information processing systems*, 32, 2019.
- 213 [32] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios. nflows: normalizing flows in PyTorch,
214 Nov. 2020.
- 215 [33] P. Eastman, R. Galvelis, R. P. Peláez, C. R. A. Abreu, S. E. Farr, E. Gallicchio, A. Gorenko,
216 M. M. Henry, F. Hu, J. Huang, A. Krämer, J. Michel, J. A. Mitchell, V. S. Pande, J. P. Rodrigues,
217 J. Rodriguez-Guerra, A. C. Simmonett, S. Singh, J. Swails, P. Turner, Y. Wang, I. Zhang, J. D.
218 Chodera, G. De Fabritiis, and T. E. Markland. OpenMM 8: Molecular Dynamics Simulation
219 with Machine Learning Potentials. *J. Phys. Chem. B*, 128(1):109–116, Jan. 2024.
- 220 [34] L. Felardos, J. Hénin, and G. Charpiat. Designing losses for data-free training of normalizing
221 flows on boltzmann distributions. *arXiv preprint arXiv:2301.05475*, 2023.
- 222 [35] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian data analysis*. Chapman and
223 Hall/CRC, 1995.
- 224 [36] V. Gómez, H. J. Kappen, J. Peters, and G. Neumann. Policy search for path integral control.
225 In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML*
226 *PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part I 14*, pages 482–497.
227 Springer, 2014.
- 228 [37] S. Goshtasbpour, V. Cohen, and F. Perez-Cruz. Adaptive annealed importance sampling with
229 constant rate progress. In *International Conference on Machine Learning*, pages 11642–11658.
230 PMLR, 2023.
- 231 [38] R. B. Grosse, C. J. Maddison, and R. R. Salakhutdinov. Annealing between distributions by
232 averaging moments. *Advances in Neural Information Processing Systems*, 26, 2013.

- 233 [39] J. Hammersley. *Monte carlo methods*. Springer Science & Business Media, 2013.
- 234 [40] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural*
235 *information processing systems*, 33:6840–6851, 2020.
- 236 [41] A. Jasra, D. A. Stephens, A. Doucet, and T. Tsagaris. Inference for lévy-driven stochastic
237 volatility models via adaptive sequential monte carlo. *Scandinavian Journal of Statistics*,
238 38(1):1–22, 2011.
- 239 [42] M. Kim, K. Seong, D. Woo, S. Ahn, and M. Kim. On scalable and efficient training of diffusion
240 samplers. *arXiv preprint arXiv:2505.19552*, 2025.
- 241 [43] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, Jan. 2017.
- 242 [44] L. Klein, A. Krämer, and F. Noé. Equivariant flow matching. *Advances in Neural Information*
243 *Processing Systems*, 36:59886–59910, 2023.
- 244 [45] L. Klein and F. Noé. Transferable boltzmann generators. *Advances in Neural Information*
245 *Processing Systems*, 37:45281–45314, 2024.
- 246 [46] J. Köhler, M. Invernizzi, P. De Haan, and F. Noé. Rigid body flows for sampling molecular
247 crystal structures. In *International Conference on Machine Learning*, pages 17301–17326.
248 PMLR, 2023.
- 249 [47] J. Köhler, A. Krämer, and F. Noé. Smooth normalizing flows. *Advances in Neural Information*
250 *Processing Systems*, 34:2796–2809, 2021.
- 251 [48] A. N. Kolmogorov and G. Castelnuovo. *Sur la notion de la moyenne*. G. Bardi, tip. della R.
252 Accad. dei Lincei, 1930.
- 253 [49] G.-H. Liu, J. Choi, Y. Chen, B. K. Miller, and R. T. Chen. Adjoint schrödinger bridge
254 sampler. *arXiv preprint arXiv:2506.22565*, 2025.
- 255 [50] W. Meng, Q. Zheng, Y. Shi, and G. Pan. An off-policy trust region policy optimization method
256 with monotonic improvement guarantee for deep reinforcement learning. *IEEE Transactions on*
257 *Neural Networks and Learning Systems*, 33(5):2223–2235, 2021.
- 258 [51] L. Midgley, V. Stimper, J. Antorán, E. Mathieu, B. Schölkopf, and J. M. Hernández-Lobato. Se
259 (3) equivariant augmented coupling flows. *Advances in Neural Information Processing Systems*,
260 36:79200–79225, 2023.
- 261 [52] L. I. Midgley, V. Stimper, G. N. Simm, B. Schölkopf, and J. M. Hernández-Lobato. Flow
262 annealed importance sampling bootstrap. *arXiv preprint arXiv:2208.01893*, 2022.
- 263 [53] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu.
264 Asynchronous methods for deep reinforcement learning. In *International conference on machine*
265 *learning*, pages 1928–1937. PmlR, 2016.
- 266 [54] R. M. Neal. Annealed importance sampling. *Statistics and computing*, 11(2):125–139, 2001.
- 267 [55] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *International*
268 *conference on machine learning*, pages 8162–8171. PMLR, 2021.
- 269 [56] F. Noé, S. Olsson, J. Köhler, and H. Wu. Boltzmann generators: Sampling equilibrium states of
270 many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- 271 [57] F. Noé and co workers. bgflow: Boltzmann generators and normalizing flows in pytorch, 2025.
272 Commit fbba56fac3eb88f6825d2bd4f745ee75ae9715e1.
- 273 [58] B. O’Donoghue, R. Munos, K. Kavukcuoglu, and V. Mnih. Pqg: Combining policy gradient
274 and q-learning. *CoRR*, 2016.
- 275 [59] F. Otto, P. Becker, N. A. Vien, H. C. Ziesche, and G. Neumann. Differentiable trust region
276 layers for deep reinforcement learning. *arXiv preprint arXiv:2101.09207*, 2021.
- 277 [60] J. Pajarinen, H. L. Thai, R. Akrouf, J. Peters, and G. Neumann. Compatible natural gradient
278 policy search. *Machine Learning*, 108:1443–1466, 2019.

- 279 [61] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin,
280 N. Gimselshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Te-
281 jani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative
282 style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer,
283 F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing*
284 *Systems*, volume 32. Curran Associates, Inc., 2019.
- 285 [62] J. Peters, K. Mulling, and Y. Altun. Relative entropy policy search. In *Proceedings of the AAAI*
286 *Conference on Artificial Intelligence*, volume 24, pages 1607–1612, 2010.
- 287 [63] D. J. Rezende, G. Papamakarios, S. Racaniere, M. Albergo, G. Kanwar, P. Shanahan, and
288 K. Cranmer. Normalizing Flows on Tori and Spheres. In *Proceedings of the 37th International*
289 *Conference on Machine Learning*, pages 8083–8092. PMLR, Nov. 2020.
- 290 [64] L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. Vargrad: a low-variance
291 gradient estimator for variational inference. *Advances in Neural Information Processing*
292 *Systems*, 33:13481–13492, 2020.
- 293 [65] H. Schopmans and P. Friederich. Conditional Normalizing Flows for Active Learning of Coarse-
294 Grained Molecular Representations. In *Forty-First International Conference on Machine*
295 *Learning*, June 2024.
- 296 [66] H. Schopmans and P. Friederich. Temperature-annealed boltzmann generators. In *Forty-second*
297 *International Conference on Machine Learning*, 2025.
- 298 [67] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization.
299 In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- 300 [68] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization
301 algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 302 [69] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based
303 generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*,
304 2020.
- 305 [70] V. Stimper, L. I. Midgley, G. N. C. Simm, B. Schölkopf, and J. M. Hernández-Lobato. Alanine
306 dipeptide in an implicit solvent at 300K. Aug. 2022.
- 307 [71] V. Stimper, B. Schölkopf, and J. M. Hernandez-Lobato. Resampling Base Distributions of Nor-
308 malizing Flows. In *Proceedings of The 25th International Conference on Artificial Intelligence*
309 *and Statistics*, pages 4915–4936. PMLR, May 2022.
- 310 [72] G. Stoltz, M. Rousset, et al. *Free energy computations: A mathematical perspective*. World
311 Scientific, 2010.
- 312 [73] M. Struwe and M. Struwe. *Variational methods*, volume 991. Springer, 2000.
- 313 [74] Y. Sun, D. Wierstra, T. Schaul, and J. Schmidhuber. Efficient natural evolution strategies. In
314 *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages
315 539–546, 2009.
- 316 [75] S. Syed, A. Bouchard-Côté, K. Chern, and A. Doucet. Optimised annealed sequential monte
317 carlo samplers. *arXiv preprint arXiv:2408.12057*, 2024.
- 318 [76] C. B. Tan, J. Bose, C. Lin, L. Klein, M. M. Bronstein, and A. Tong. Scalable Equilibrium
319 Sampling with Sequential Boltzmann Generators. In *Forty-Second International Conference on*
320 *Machine Learning*, June 2025.
- 321 [77] C. B. Tan, M. Hassan, L. Klein, S. Syed, D. Beaini, M. M. Bronstein, A. Tong, and K. Neklyudov.
322 Amortized sampling with transferable normalizing flows. In *ICML 2025 Generative AI and*
323 *Biology (GenBio) Workshop*, 2025.
- 324 [78] D. Thalmeier, H. J. Kappen, S. Totaro, and V. Gómez. Adaptive smoothing for path integral
325 control. *Journal of Machine Learning Research*, 21(191):1–37, 2020.

- 326 [79] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski,
327 P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman,
328 N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng,
329 E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero,
330 C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0
331 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature*
332 *Methods*, 17:261–272, 2020.
- 333 [80] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber. Natural evolution
334 strategies. *The Journal of Machine Learning Research*, 15(1):949–980, 2014.
- 335 [81] C. Williams, A. Campbell, A. Doucet, and S. Syed. Score-optimal diffusion schedules. *Advances*
336 *in Neural Information Processing Systems*, 37:107960–107983, 2024.
- 337 [82] Y. Wu, E. Mansimov, R. B. Grosse, S. Liao, and J. Ba. Scalable trust-region method for deep
338 reinforcement learning using kronecker-factored approximation. *Advances in neural information*
339 *processing systems*, 30, 2017.
- 340 [83] H. Xu, J. Xuan, G. Zhang, and J. Lu. Trust region policy optimization via entropy regularization
341 for kullback–leibler divergence constraint. *Neurocomputing*, 589:127716, 2024.
- 342 [84] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge. Projection-based constrained policy
343 optimization. *arXiv preprint arXiv:2010.03152*, 2020.
- 344 [85] L. Zhang. The cosine schedule is fisher-rao-optimal for masked discrete diffusion models. *arXiv*
345 *preprint arXiv:2508.04884*, 2025.
- 346 [86] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines
347 for large-scale bound-constrained optimization. *ACM Transactions on mathematical software*
348 *(TOMS)*, 23(4):550–560, 1997.

349	Appendix	
350	A Proofs	11
351	B Related work	14
352	C Numerical evaluation	14
353	D Experimental setup	18
354	D.1 Architecture	18
355	D.2 Target densities	18
356	D.3 Dual optimization in practice	19
357	D.4 Metrics	20
358	D.5 Hyperparameters	22

359 A Proofs

360 *Proof of Prop. 2.1.* Consider the constrained optimization problem

$$q_{i+1} = \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q|p) \quad \text{s.t.} \quad D_{\text{KL}}(q|q_i) \leq \varepsilon_{\text{tr}}, \quad H(q_i) - H(q) \leq \varepsilon_{\text{ent}}, \quad \int dq = 1. \quad (10)$$

361 and its corresponding Lagrangian

$$\mathcal{L}_{\text{tr}}^{(i+1)}(q, \lambda, \eta, \omega) = D_{\text{KL}}(q|p) + \lambda(D_{\text{KL}}(q|q_i) - \varepsilon_{\text{tr}}) + \eta(H(q_i) - H(q) - \varepsilon_{\text{ent}}) + \omega \left(\int dq - 1 \right). \quad (11)$$

362 Using the convexity of the Kullback-Leibler (KL) divergence in its arguments, the convexity of the
 363 negative Shannon entropy, and that the integral is a linear functional, the objective from Eq. (10)
 364 and its Lagrangian, given by Eq. (11) are convex in q . Using that $\mathcal{P}(\mathbb{R}^d)$ is continuous, there always
 365 exists a measure $\tilde{q} \neq q_i$ with $D_{\text{KL}}(\tilde{q}|q_i) < \varepsilon_{\text{tr}}$ and $H(q_i) - H(\tilde{q}) < \varepsilon_{\text{ent}}$ that satisfies the inequality
 366 constraints strictly. Following [16, 5.2.3], Slater's condition holds, implying strong duality. We can
 367 therefore instead solve the dual problem.

368 We start by setting up the Euler-Lagrange equation, given by

$$\frac{\partial}{\partial q} \mathcal{L}_{\text{tr}}^{(i+1)}(q, \lambda, \eta, \omega) = 0,$$

369 using

$$\begin{aligned} \mathcal{L}_{\text{tr}}^{(i+1)}(q, \lambda, \eta, \omega) = & \int q(x) \left((1 + \lambda + \eta) \log q(x) - \log p(x) - \lambda \log q_i(x) + \omega \right) dx \\ & - \lambda \varepsilon_{\text{tr}} + \eta(H(q_i) - \varepsilon_{\text{ent}}) - \omega \end{aligned}$$

370 and solve for q . Hence, we get

$$\begin{aligned} \frac{\partial}{\partial q} \mathcal{L}_{\text{tr}}^{(i+1)}(q, \lambda, \eta, \omega) = & (1 + \lambda + \eta)(\log q + 1) - \log p - \lambda \log q_i + \omega = 0 \\ \Leftrightarrow \log q = & \log \left(q_i^{\frac{\lambda}{1+\lambda+\eta}} p^{\frac{1}{1+\lambda+\eta}} \right) - \left(\frac{\omega}{1 + \lambda + \eta} + 1 \right). \end{aligned} \quad (12)$$

371 Using this, we can further determine ω using

$$\begin{aligned} \int dq = & \int q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} p(x)^{\frac{1}{1+\lambda+\eta}} dx / \exp \left(\frac{\omega}{1 + \lambda + \eta} + 1 \right) = 1 \\ \Leftrightarrow \omega = & (1 + \lambda + \eta)(\log \mathcal{Z}_{i+1}(\lambda, \eta) - 1), \quad \text{with} \quad \mathcal{Z}_{i+1}(\lambda, \eta) = \int q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} p(x)^{\frac{1}{1+\lambda+\eta}}(x) dx. \end{aligned}$$

372 Substituting ω back into (12) and simplifying the fraction using $\tilde{p} = \mathcal{Z}p$ yields

$$q_{i+1}(x, \lambda, \eta) = \frac{q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} \tilde{p}(x)^{\frac{1}{1+\lambda+\eta}}}{\mathcal{Z}_{i+1}(\lambda, \eta)} \quad \text{with} \quad \mathcal{Z}_{i+1}(\lambda, \eta) = \int q_i(x)^{\frac{\lambda}{1+\lambda+\eta}} \tilde{p}(x)^{\frac{1}{1+\lambda+\eta}}(x) dx,$$

373 which uses the unnormalized target \tilde{p} , proving Prop. 2.1.

374 □

375 *Proof of Thm. 2.2.* We divide the proof into three parts and start with the most general formulation
 376 using both constraints. The first two parts will show form and monotonicity while part three will
 377 derive the special cases with just the trust-region and just the entropy constraint.

378 **Part 1:** Given are the sequences of Lagrangian multipliers $(\lambda_i)_{i \in \mathbb{N}_0} \geq 0$ and $(\eta_i)_{i \in \mathbb{N}_0} \geq 0$. We now
 379 aim to proof that the sequence $(q_i)_{i \in \mathbb{N}_0}$, given by

$$\tilde{q}_i = \begin{cases} q_0 & , \quad i = 0 \\ \frac{q_{i-1}^{\lambda_{i-1}}}{\tilde{q}_{i-1}^{1+\lambda_{i-1}+\eta_{i-1}}} \tilde{p}^{\frac{1}{1+\lambda_{i-1}+\eta_{i-1}}} & , \quad i \geq 1 \end{cases}, \quad (13)$$

380 takes the form

$$\begin{aligned} \tilde{q}_i = q_0^{1-\beta_i} (\tilde{p}^{\alpha_i})^{\beta_i} \quad \text{with} \quad \beta_i = & 1 - \prod_{j=0}^{i-1} \frac{\lambda_j}{1 + \lambda_j + \eta_j} \\ \text{and} \quad \alpha_i = & \begin{cases} 0 & , \quad i = 0 \\ 1 - \frac{1}{\beta_i} \sum_{k=0}^{i-1} \frac{\eta_k}{1 + \lambda_k + \eta_k} \prod_{j=k+1}^{i-1} \frac{\lambda_j}{1 + \lambda_j + \eta_j} & , \quad i \geq 1. \end{cases} \end{aligned}$$

381 We use the common convention that empty products evaluate to one.

382 **Base case** ($i = 0$): The simplest case

$$\tilde{q}_0 = q_0^{1-\beta_0} (\tilde{p}^{\alpha_0})^{\beta_0}$$

383 holds due to $\beta_0 = 0$ (using empty product convention).

384 **Inductive step:** We start from Eq. (13) and transform it using the assumption that $\tilde{q}_i = q_0^{1-\beta_i} (\tilde{p}^{\alpha_i})^{\beta_i}$
 385 holds for some arbitrary but fixed $i \in \mathbb{N}_0$, yielding

$$\begin{aligned} \tilde{q}_{i+1} &= \tilde{q}_i^{\frac{\lambda_i}{1+\lambda_i+\eta_i}} \tilde{p}^{\frac{1}{1+\lambda_i+\eta_i}} \\ &= \left(q_0^{1-\beta_i} (\tilde{p}^{\alpha_i})^{\beta_i} \right)^{\frac{\lambda_i}{1+\lambda_i+\eta_i}} \tilde{p}^{\frac{1}{1+\lambda_i+\eta_i}} \\ &= q_0^{1-\beta_{i+1}} \tilde{p}^{\alpha_i \beta_i \frac{\lambda_i}{1+\lambda_i+\eta_i} + \frac{1}{1+\lambda_i+\eta_i}}. \end{aligned}$$

386 Using

$$\begin{aligned} \beta_i \frac{\lambda_i}{1+\lambda_i+\eta_i} &= \frac{\lambda_i}{1+\lambda_i+\eta_i} - \prod_{j=0}^i \frac{\lambda_j}{1+\lambda_j+\eta_j} \\ &= 1 - \frac{1+\eta_i}{1+\lambda_i+\eta_i} - \prod_{j=0}^i \frac{\lambda_j}{1+\lambda_j+\eta_j} \\ &= \beta_{i+1} - \frac{1+\eta_i}{1+\lambda_i+\eta_i}, \end{aligned}$$

387 we now can rewrite the exponent of p yielding

$$\begin{aligned} &\alpha_i \beta_i \frac{\lambda_i}{1+\lambda_i+\eta_i} + \frac{1}{1+\lambda_i+\eta_i} \\ &= \left(\beta_i - \sum_{k=0}^{i-1} \frac{\eta_k}{1+\lambda_k+\eta_k} \prod_{j=k+1}^{i-1} \frac{\lambda_j}{1+\lambda_j+\eta_j} \right) \frac{\lambda_i}{1+\lambda_i+\eta_i} + \frac{1}{1+\lambda_i+\eta_i} \\ &= \beta_{i+1} - \frac{1+\eta_i}{1+\lambda_i+\eta_i} - \sum_{k=0}^{i-1} \frac{\eta_k}{1+\lambda_k+\eta_k} \prod_{j=k+1}^i \frac{\lambda_j}{1+\lambda_j+\eta_j} + \frac{1}{1+\lambda_i+\eta_i} \\ &= \beta_{i+1} - \sum_{k=0}^i \frac{\eta_k}{1+\lambda_k+\eta_k} \prod_{j=k+1}^i \frac{\lambda_j}{1+\lambda_j+\eta_j}, \\ &= \alpha_{i+1} \beta_{i+1}. \end{aligned}$$

388 again using the convention that an empty product evaluates to one. It directly follows

$$q_{i+1} \propto \tilde{q}_{i+1} = q_0^{1-\beta_{i+1}} (\tilde{p}^{\alpha_{i+1}})^{\beta_{i+1}},$$

389 which completes the induction.

390 **Part 2:** It remains to show that $(\alpha_i)_{i \in \mathbb{N}_0}$ and $(\beta_i)_{i \in \mathbb{N}_0}$ take values in $[0, 1]$ and are monotonically
 391 increasing with $\alpha_0 = \beta_0 = 0$ and $\alpha_I = \beta_I = 1$ for some $I \in \mathbb{N}_+$.

392 The first case ($\alpha_0 = \beta_0 = 0$) holds by definition. Assuming that there exists an $I \in \mathbb{N}_+$, such that
 393 $\lambda_{I-1} = \eta_{I-1} = 0$,

$$\beta_i = 1 - \prod_{j=0}^{i-1} \frac{\lambda_j}{1+\lambda_j+\eta_j} \stackrel{i \geq I}{=} 1$$

394 and

$$\alpha_i = 1 - \frac{1}{\beta_i} \sum_{k=0}^{i-1} \frac{\eta_k}{1+\lambda_k+\eta_k} \prod_{j=k+1}^{i-1} \frac{\lambda_j}{1+\lambda_j+\eta_j} \stackrel{i \geq I}{=} 1$$

395 follow directly for all $i \geq I$. Both the trust-region and entropy constraints become inactive at
 396 the optimum and do not prevent $(q_i)_{i \in \mathbb{N}_0}$ from reaching the unique optimum p ($\varepsilon_{\text{tr}}, \varepsilon_{\text{ent}} > 0$).
 397 Consequently, both Lagrange multipliers will eventually vanish, motivating the existence of some
 398 $I \in \mathbb{N}_+$, such that $\lambda_{I-1} = \eta_{I-1} = 0$.

399 We now show monotonicity of $(\beta_i)_{i \in \mathbb{N}_0}$ using $(\lambda_i)_{i \in \mathbb{N}_0} \geq 0$ and $(\eta_i)_{i \in \mathbb{N}_0} \geq 0$. Let $i \in \mathbb{N}_0$ be
 400 arbitrary. As a direct consequence of

$$\begin{aligned} \beta_{i+1} - \beta_i &= \prod_{j=0}^{i-1} \frac{\lambda_j}{1 + \lambda_j + \eta_j} - \prod_{j=0}^i \frac{\lambda_j}{1 + \lambda_j + \eta_j} \\ &= \left(\prod_{j=0}^{i-1} \frac{\lambda_j}{1 + \lambda_j + \eta_j} \right) \left(1 - \frac{\lambda_i}{1 + \lambda_i + \eta_i} \right) \\ &= \left(\prod_{j=0}^{i-1} \frac{\lambda_j}{1 + \lambda_j + \eta_j} \right) \left(\frac{1 + \eta_i}{1 + \lambda_i + \eta_i} \right) \stackrel{\lambda_j, \eta_j \geq 0}{\geq} 0, \end{aligned}$$

401 the sequence $(\beta_i)_{i \in \mathbb{N}_0}$ must be monotonically increasing.

402 **Part 3:** We now consider the special cases in which only the trust-region constraint or the entropy
 403 constraint is active by setting the Lagrange multiplier sequence of the other constraint to zero.

404 We first consider only the trust-region constraint (2), which corresponds to setting the Lagrangian
 405 multiplier of the entropy constraint to zero, i.e., $\eta_i = 0$ for all $i \in \mathbb{N}_0$. In this scenario, α_i simplifies
 406 to $\alpha_0 = 0$ and $\alpha_i = 1$ for all $i \geq 1$. Consequently, and using $\beta_0 = 0$, the iterates take the form

$$q_i \propto \tilde{q}_i = q_0^{1-\beta_i} \tilde{p}^{\beta_i}, \quad i \in \mathbb{N}_0,$$

407 as claimed.

408 Analogously, the trust-region constraint can be rendered inactive by setting $\lambda_i = 0$ for all $i \in \mathbb{N}_0$,
 409 leaving only the entropy constraint active, corresponding to Eq. (3). In this case, $\beta_0 = 0$ and $\beta_i = 1$
 410 for all $i \geq 1$, yielding

$$q_i \propto \tilde{q}_i = \begin{cases} q_0 & , \quad i = 0, \\ \tilde{p}^{\alpha_i} & , \quad i \geq 1, \end{cases}$$

411 which concludes the proof. \square

412 *Proof of bounded importance-weight variance its effect on the effective sample size.* In this section,
 413 we show that using the trust-region constraint yields an approximate lower bound for the effective
 414 sample between any two consecutive distributions q_i and q_{i+1} . This approximate lower bound only
 415 depends on ε_{tr} .

416 The variance of the importance weights

$$\frac{q_{i+1}(x)}{q_i(x)} = \begin{cases} \frac{1}{\mathcal{Z}_{i+1}(\lambda_i)} \left(\frac{p(x)}{q_i(x)} \right)^{\frac{1}{1+\lambda_i}} & \text{with trust-region constraint (2)} \\ \frac{1}{\mathcal{Z}_{i+1}(\lambda_i, \eta_i)} \left(\frac{p(x)}{q_i(x)^{1+\eta_i}} \right)^{\frac{1}{1+\lambda_i+\eta_i}} & \text{with trust-region + entropy constraint (4)} \end{cases}$$

417 between two normalized consecutive distributions is closely connected to the effective sample size
 418 via

$$\text{ESS}(q_i, q_{i+1}) = \frac{1}{1 + \text{Var}_{q_i} \left(\frac{q_{i+1}(x)}{q_i(x)} \right)},$$

419 also explained in App. D.4. The relation $\text{Var}_{q_i}(q_{i+1}(x)/q_i(x)) = \chi^2(q_{i+1}|q_i)$ [24] and the well-
 420 known Taylor approximation $\chi^2(q_{i+1}|q_i) \approx 2D_{\text{KL}}(q_{i+1}|q_i)$ [26] lets us rewrite the effective sample
 421 size in terms of the KL divergence between q_{i+1} and q_i yielding

$$\text{ESS}(q_i, q_{i+1}) \approx \frac{1}{1 + 2D_{\text{KL}}(q_{i+1}|q_i)}$$

422 as approximation for the effective sample size. This approximation is justified under the assumption
 423 that q_{i+1} is close to q_i , a condition that is satisfied by the design of the problem for a small trust-
 424 region bound $\varepsilon_{\text{tr}} > 0$. Due to q_{i+1} being the optimal solution to an objective with the constraint
 425 $D_{\text{KL}}(q|q_i) \leq \varepsilon_{\text{tr}}$, the constraint must also hold for $q = q_{i+1}$ resulting in the approximate lower
 426 bound

$$\text{ESS}(q_i, q_{i+1}) \gtrsim \frac{1}{1 + 2\varepsilon_{\text{tr}}} \quad (14)$$

427 for the effective sample size of the importance weights with equality in all but the last step.

428 This approximate lower bound justifies the use of Monte Carlo approximations in Sec. 3, helping to
 429 stabilize training independent of the problem's dimensionality. \square

430 *Proof of uniqueness and tightness of the trust-region solution.* Closely following [14], we now es-
 431 tablish the uniqueness of the trust-region solution and show that the trust-region constraint is tight for
 432 all but the the final step. Specifically, we show

$$D_{\text{KL}}(q_i|p) < \varepsilon_{\text{tr}} \implies q_i = p$$

$$q_i = \arg \min D_{\text{KL}}(q|p) \quad \text{s.t.} \quad D_{\text{KL}}(q|q_{i-1}) \leq \varepsilon_{\text{tr}} \quad \text{is unique}$$

433 If $D_{\text{KL}}(q_i||p) < \varepsilon_{\text{tr}}$, the KKT conditions imply that the Lagrangian multiplier satisfies $\lambda_i = 0$, so
 434 the constraint is inactive. Consequently, q_i must solve the strictly convex unconstrained problem

$$\min_{q \in \mathcal{P}(\mathbb{R}^d)} D_{\text{KL}}(q|p),$$

435 which has the unique minimizer p . Since p is feasible ($D_{\text{KL}}(p|p) = 0 \leq \varepsilon_{\text{tr}}$), it follows that $q_i = p$.

436 Uniqueness of q_i further follows from the convexity of the feasible set $\{q \in \mathcal{P} | D_{\text{KL}}(q|q_i) \leq \varepsilon_{\text{tr}}\}$
 437 together with the strict convexity of the objective in q when p is fixed. \square

438 B Related work

439 **Boltzmann generators.** Learning molecular Boltzmann generators [56] from energy evaluations has
 440 been explored using internal coordinates with flow-based [28, 46, 47, 52, 66] and diffusion-based
 441 methods [22, 42, 49]. While flow-based approaches have demonstrated strong performance, their
 442 diffusion-based counterparts remain less competitive, often struggling with mode collapse, even on
 443 relatively small systems. Cartesian-coordinate approaches [6, 44, 51, 76] enable transferability across
 444 systems [45, 77].

445 **Trust regions and entropy constraints.** Trust region methods have a long history as robust opti-
 446 mization algorithms that iteratively minimize an objective within an adaptively sized “trust region”;
 447 see [25] for an overview. Beyond classical optimization, these methods have been extended to
 448 operate over spaces of probability distributions, with applications in reinforcement learning [2–
 449 4, 9, 50, 59, 60, 62, 67, 68, 82–84], black-box optimization [1, 74, 80], variational inference [10, 11],
 450 and path integral control [36, 78]. The first explicit link between trust-region optimization and
 451 geometric annealing paths was established in [14] for path space measures in the setting of stochastic
 452 optimal control. Entropy constraints, often introduced as entropy regularization, have also been stud-
 453 ied in policy optimization and reinforcement learning, either in the form of soft constraints [5, 53, 58]
 454 or hard constraints [1, 7–9, 60]. However, prior work typically constrains the absolute entropy value,
 455 which is problematic for inference tasks, since it requires prior knowledge of the target density’s
 456 entropy. To the best of our knowledge, such methods have not yet been extended to sampling
 457 problems. Furthermore, the connection between entropy-constrained optimization and annealing
 458 paths has not previously been established.

459 **Improved annealing paths.** Research on improving annealing paths (APs) has largely focused
 460 on geometric APs in the context of annealed importance sampling (AIS) [54] and their extensions
 461 to sequential Monte Carlo (SMC) [27]; see [23, 37, 41, 75]. Beyond the standard geometric AP,
 462 alternative constructions have been proposed, such as the moment-averaging path for exponential
 463 family distributions [38] and the arithmetic mean path [21]. The geometric path itself can be
 464 interpreted as a quasi-arithmetic mean [48] under the natural logarithm, which motivated [17] to
 465 propose APs based on the deformed logarithm transformation. A variational characterization of these
 466 paths was later analysed in [18]. Related work also explores improved schedules for the diffusion
 467 coefficient in ergodic Ornstein–Uhlenbeck processes used to train denoising diffusion models [40, 69];
 468 see, e.g., [13, 55, 81, 85].

469 C Numerical evaluation

470 In this section, we compare our approach against state-of-the-art methods on four challenging
 471 molecular systems. We provide a brief overview of the experimental setup here, with full details in
 472 App. D. We further present several ablation study on the effect of both constraints, and an analysis of
 473 the trust-region constraint across systems of different dimensionality.

474 **Benchmark problems.** We evaluate on alanine dipeptide ($d = 60$) [29, 52, 71, 76], alanine
 475 tetrapeptide ($d = 120$), alanine hexapeptide ($d = 180$) [66], and the newly introduced ELIL
 476 tetrapeptide ($d = 219$), which, to the best of our knowledge, is the largest molecular system studied
 477 with variational approaches to date. See App. D.2 for details.

SYSTEM	METHOD	TARGET EVALS ↓	NLL ↓	ESS [%] ↑	RAM KL ↓	RAM KL w. RW ↓
ALANINE DIPEPTIDE ($d = 60$)	FORWARD KL	5×10^9	-213.581 ± 0.000	$(82.16 \pm 0.09) \%$	$(2.21 \pm 0.05) \times 10^{-3}$	$(1.99 \pm 0.07) \times 10^{-3}$
	REVERSE KL	2.56×10^8	-213.609 ± 0.006	$(94.11 \pm 0.21) \%$	$(1.75 \pm 0.28) \times 10^{-2}$	$(1.65 \pm 0.29) \times 10^{-2}$
	FAB	2.13×10^8	-213.653 ± 0.000	$(94.81 \pm 0.04) \%$	$(1.50 \pm 0.03) \times 10^{-3}$	$(1.25 \pm 0.01) \times 10^{-3}$
	TA-BG	1×10^8	-213.666 ± 0.001	$(95.77 \pm 0.12) \%$	$(1.94 \pm 0.07) \times 10^{-3}$	$(1.36 \pm 0.02) \times 10^{-3}$
	CMT (OURS)	1×10^8	-213.677 ± 0.000	$(97.69 \pm 0.03) \%$	$(1.49 \pm 0.02) \times 10^{-3}$	$(1.41 \pm 0.02) \times 10^{-3}$
ALANINE TETRA- PEPTIDE ($d = 120$)	FORWARD KL	4.2×10^9	-330.069 ± 0.001	$(45.29 \pm 0.08) \%$	$(2.26 \pm 0.06) \times 10^{-3}$	$(2.50 \pm 0.03) \times 10^{-3}$
	REVERSE KL	2.56×10^8	-329.191 ± 0.122	$(74.88 \pm 3.65) \%$	$(3.00 \pm 0.35) \times 10^{-1}$	$(2.87 \pm 0.40) \times 10^{-1}$
	FAB	2.13×10^8	-330.100 ± 0.002	$(63.59 \pm 0.23) \%$	$(6.89 \pm 0.25) \times 10^{-3}$	$(1.25 \pm 0.01) \times 10^{-3}$
	TA-BG	1×10^8	-330.143 ± 0.002	$(64.87 \pm 0.21) \%$	$(2.47 \pm 0.23) \times 10^{-3}$	$(1.71 \pm 0.07) \times 10^{-3}$
	CMT (OURS)	1×10^8	-330.167 ± 0.002	$(68.49 \pm 0.14) \%$	$(1.99 \pm 0.05) \times 10^{-3}$	$(1.65 \pm 0.07) \times 10^{-3}$
ALANINE HEXA- PEPTIDE ($d = 180$)	FORWARD KL	4.2×10^9	-501.598 ± 0.005	$(10.97 \pm 0.11) \%$	$(4.16 \pm 0.26) \times 10^{-3}$	$(7.69 \pm 0.03) \times 10^{-3}$
	REVERSE KL	2.56×10^8	-497.378 ± 0.277	$(22.22 \pm 1.44) \%$	$(5.41 \pm 0.38) \times 10^{-1}$	$(5.32 \pm 0.38) \times 10^{-1}$
	FAB	4.2×10^8	-501.268 ± 0.008	$(14.64 \pm 0.08) \%$	$(2.09 \pm 0.02) \times 10^{-2}$	$(1.12 \pm 0.02) \times 10^{-2}$
	TA-BG	4×10^8	-501.582 ± 0.010	$(15.89 \pm 0.13) \%$	$(8.33 \pm 0.59) \times 10^{-3}$	$(8.15 \pm 0.44) \times 10^{-3}$
	CMT (OURS)	4×10^8	-501.761 ± 0.008	$(29.20 \pm 0.24) \%$	$(1.25 \pm 0.04) \times 10^{-2}$	$(1.21 \pm 0.01) \times 10^{-2}$
ELIL TETRA- PEPTIDE ($d = 219$)	FORWARD KL	4.2×10^9	-597.571 ± 0.004	$(5.80 \pm 0.04) \%$	$(4.12 \pm 0.03) \times 10^{-3}$	$(9.38 \pm 0.06) \times 10^{-3}$
	REVERSE KL	2.56×10^8	-583.381 ± 3.033	$(1.25 \pm 0.45) \%$	$(1.22 \pm 0.28) \times 10^0$	$(1.14 \pm 0.31) \times 10^0$
	FAB	8.43×10^8	-597.370 ± 0.006	$(7.30 \pm 0.08) \%$	$(2.56 \pm 0.10) \times 10^{-2}$	$(9.01 \pm 0.57) \times 10^{-3}$
	TA-BG	8×10^8	-597.830 ± 0.022	$(10.12 \pm 0.38) \%$	$(7.35 \pm 0.89) \times 10^{-3}$	$(7.65 \pm 0.59) \times 10^{-3}$
	CMT	8×10^8	-598.440 ± 0.003	$(25.91 \pm 0.19) \%$	$(5.74 \pm 0.05) \times 10^{-3}$	$(5.22 \pm 0.10) \times 10^{-3}$

Table 2: Comparison of metrics obtained for all four peptide systems. The table reports the number of target evaluations (Target Evals), the negative log-likelihood (NLL), the effective sample size (ESS), the average forward KL divergence to the ground-truth Ramachandran plots (Ram KL) and its importance-weighted version (Ram KL w. RW.), along with the corresponding system and method. All values are presented as the mean and standard error across four independent experiments. The best-performing variational method for each metric is highlighted in bold. Reverse KL was excluded, as it tends to suffer from mode collapse, making ESS values not directly comparable.

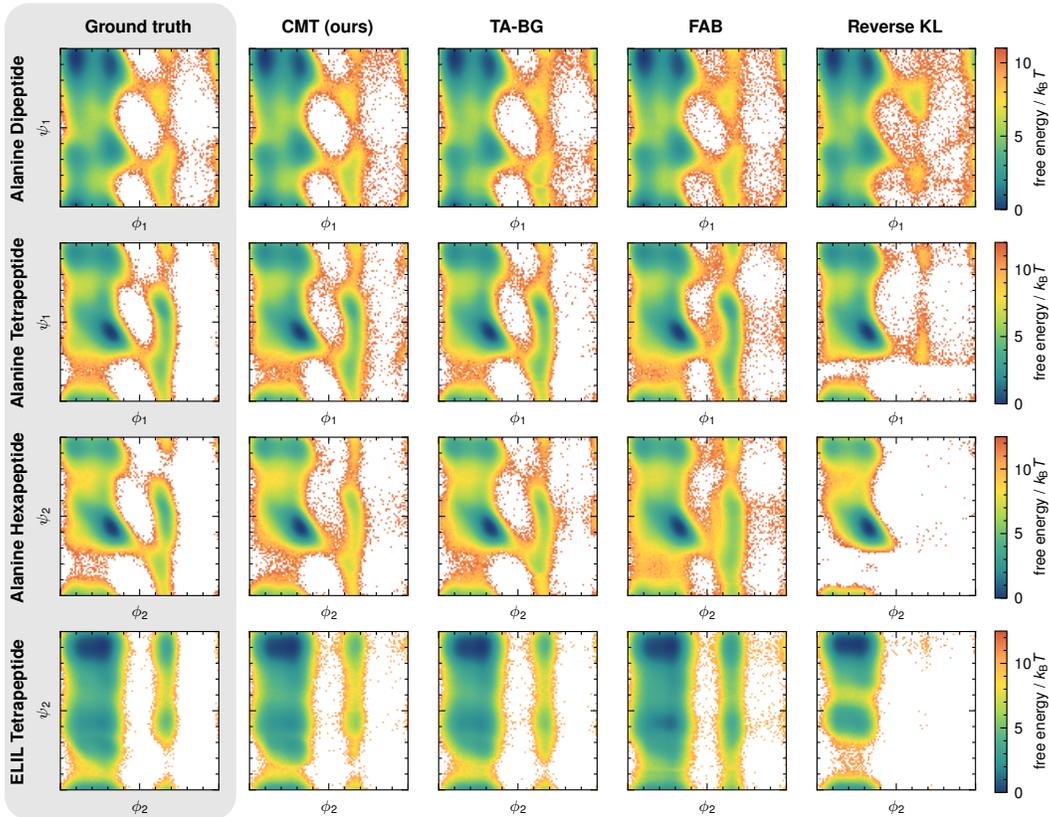


Figure 2: Comparison of Ramachandran plots of backbone dihedral angle pairs obtained with different methods. See App. D.4 for details.

478 **Baseline methods.** We compare against Flow Annealed Importance Sampling Bootstrap (FAB) [52],
 479 Temperature-Annealed Boltzmann Generators (TA-BG) [66], reverse KL and forward KL, trained
 480 on MD samples. All methods use neural spline flows [31] with identical architectures and internal
 481 coordinates.

482 **Performance criteria.** Methods are evaluated using two metrics. First, the negative log-likelihood
 483 (NLL) computed with ground truth MD samples, which corresponds (up to an additive constant) to
 484 the forward KL divergence and is effective for detecting mode collapse [15]. Second, we consider
 485 the reverse effective sample size (ESS), a common measure of sample quality but less reliable for
 486 assessing mode coverage [15]. For full details, see App. D.4. We also report the number of target
 487 density evaluations, which are typically expensive for molecular systems. Ram KL and Ram KL
 488 w. RW report the forward KL divergence between the ground truth Ramachandran distribution
 489 obtained from MD simulations and the model-predicted Ramachandran distribution, with and without
 490 reweighting, respectively. Substantially elevated values are a strong indicator for the presence of
 491 mode collapse. For a detailed discussion of all evaluation metrics, see App. D.4.

492 **Main results.** The findings are summarized in Tab. 2. While the performance gap between our
 493 method and the baselines is less pronounced for smaller systems, it widens substantially for the larger
 494 ones. In particular, on alanine hexapeptide and ELIL tetrapeptide, our method attains approximately
 495 twice the ESS of competing approaches, while also avoiding mode collapse, as reflected in improved
 496 NLL values.

SYSTEM	CONSTRAINT		TARGET EVALS ↓	NLL ↓	ESS [%] ↑ ★	RAM KL ↓	RAM KL w. RW ↓
	TRUST-REGION	ENTROPY					
ALANINE DIPEPTIDE ($d = 60$)	✗	✗	1×10^8	-213.667 ± 0.000	(97.71 ± 0.01) %	$(1.50 \pm 0.02) \times 10^{-3}$	$(1.50 \pm 0.01) \times 10^{-3}$
	✓	✗	1×10^8	-213.673 ± 0.000	(97.46 ± 0.02) %	$(1.52 \pm 0.03) \times 10^{-3}$	$(1.47 \pm 0.01) \times 10^{-3}$
	✗	✓	1×10^8	-213.677 ± 0.000	(97.76 ± 0.04) %	$(1.50 \pm 0.02) \times 10^{-3}$	$(1.38 \pm 0.02) \times 10^{-3}$
	✓	✓	1×10^8	-213.677 ± 0.000	(97.69 ± 0.03) %	$(1.49 \pm 0.02) \times 10^{-3}$	$(1.41 \pm 0.02) \times 10^{-3}$
ALANINE TETRA- PEPTIDE ($d = 120$)	✗	✗	1×10^8	-329.787 ± 0.193	(69.60 ± 2.20) %	$(6.92 \pm 3.44) \times 10^{-2}$	$(6.64 \pm 3.25) \times 10^{-2}$
	✓	✗	1×10^8	-330.154 ± 0.001	(66.11 ± 0.02) %	$(2.08 \pm 0.04) \times 10^{-3}$	$(2.05 \pm 0.04) \times 10^{-3}$
	✗	✓	1×10^8	-330.130 ± 0.003	(63.34 ± 0.24) %	$(2.27 \pm 0.04) \times 10^{-3}$	$(1.77 \pm 0.02) \times 10^{-3}$
	✓	✓	1×10^8	-330.167 ± 0.002	(68.49 ± 0.14) %	$(1.99 \pm 0.05) \times 10^{-3}$	$(1.65 \pm 0.07) \times 10^{-3}$
ALANINE HEXA- PEPTIDE ($d = 180$)	✗	✗	4×10^8	-499.746 ± 0.183	(25.63 ± 1.05) %	$(2.52 \pm 0.33) \times 10^{-1}$	$(2.55 \pm 0.36) \times 10^{-1}$
	✓	✗	4×10^8	-501.322 ± 0.232	(33.45 ± 2.45) %	$(4.21 \pm 1.34) \times 10^{-2}$	$(4.14 \pm 1.40) \times 10^{-2}$
	✗	✓	4×10^8	-501.333 ± 0.017	(15.03 ± 1.00) %	$(1.27 \pm 0.08) \times 10^{-2}$	$(1.39 \pm 0.11) \times 10^{-2}$
	✓	✓	4×10^8	-501.761 ± 0.008	(29.20 ± 0.24) %	$(1.25 \pm 0.04) \times 10^{-2}$	$(1.21 \pm 0.01) \times 10^{-2}$

Table 3: Performance of CMT with the trust-region and entropy constraints selectively enabled or disabled. For all metrics except ESS, the best-performing variant is highlighted in bold. ★ The reverse ESS is only partially comparable, as training without the entropy constraint led to partial mode collapse on alanine tetrapeptide and alanine hexapeptide, rendering direct comparisons unreliable. To avoid confusion, no best-performing method is highlighted in bold for ESS.

497 **Ablation study on the constraints.** Tab. 3 presents the performance of our method under different
 498 configurations, with the trust-region and entropy constraints selectively enabled or disabled. The
 499 absence of both constraints effectively corresponds to importance-weighted forward KL training.
 500 Considering the NLL, which serves as a forward metric, it becomes clear that both constraints are
 501 necessary to achieve optimal performance. Variants of the method without the entropy constraint
 502 exhibited at least partial mode collapse, rendering the ESS largely incomparable. Partial mode
 503 collapse is further reflected in the RAM KL and RAM KL w. RW. metrics, which take substantially
 504 higher values when the entropy constraint is omitted.

505 Fig. 3 depicts the evolution of model entropy and the gradient norm (prior to clipping) during
 506 training across different systems. Training with only the entropy constraint yields an approximately
 507 linear decay of entropy for both alanine dipeptide and alanine hexapeptide. In the case of alanine
 508 hexapeptide, however, the entropy constraint is noticeably violated, likely due to the system’s higher
 509 dimensionality and the pronounced discrepancy between the initial model distribution q_0 and the
 510 first intermediate distribution q_1 . Larger system sizes also tend to increase the gradient norm, most
 511 prominently in alanine hexapeptide. The combination of the trust-region and entropy constraints
 512 produces the most stable gradient norms, while the approximately linear entropy decay indicates
 513 that the entropy constraint is effectively enforced, thereby enabling its practical application even in
 514 the case of alanine hexapeptide. By contrast, the trust-region constraint alone leads to a more rapid
 515 entropy collapse, which reduces exploration and ultimately limits the algorithm’s final performance
 516 in practice.

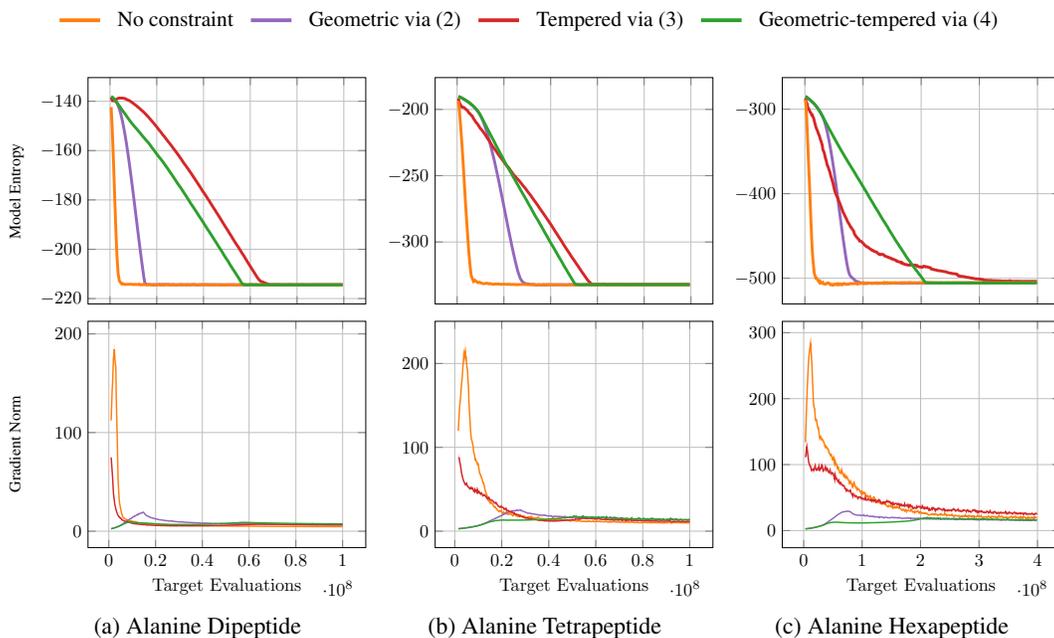


Figure 3: Effect of trust-region and entropy constraint on the model entropy (top row) and the gradient norm (bottom row) across different molecular systems.

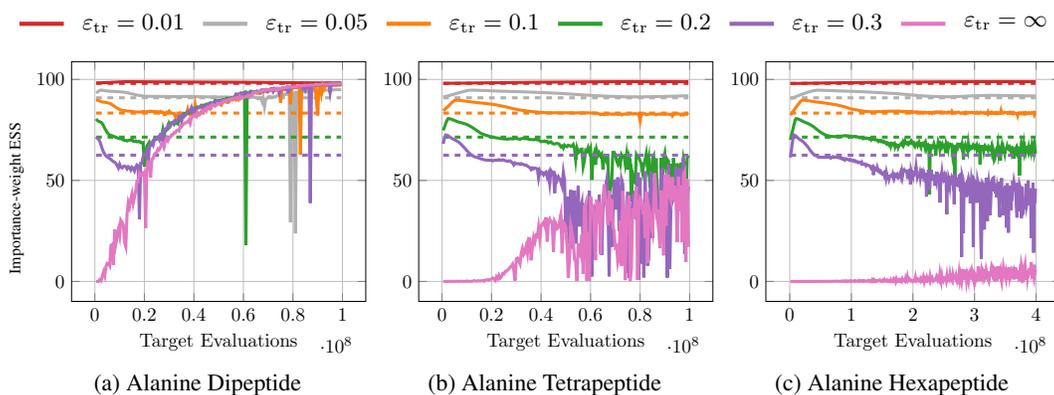


Figure 4: Importance-weight variance between successive intermediate distributions, shown in terms of effective sample size (ESS), for different trust-region bounds and system sizes. Each trust-region bound ϵ_{tr} defines an approximate lower bound on the ESS, indicated by dashed lines.

517 **Ablation study on the trust-region bound.** Fig. 4 illustrates the importance-weight variance of
 518 CMT across different trust-region bounds and system sizes, highlighting the approximate direct
 519 relationship between the trust-region bound and the variance of importance weights between con-
 520 secutive intermediate distributions. Importance-weight variance is expressed in terms of effective
 521 sample size (ESS). In the absence of a trust-region constraint ($\varepsilon_{\text{tr}} = \infty$), the ESS decreases with
 522 increasing system size. By contrast, finite trust-region bounds yield a substantially larger and more
 523 stable ESS, with the approximate lower bound on the ESS becoming increasingly well realized as
 524 the trust-region bound ε_{tr} decreases. Notably, this approximate lower bound is independent of the
 525 problem’s dimensionality, a property that is empirically supported.

526 D Experimental setup

527 D.1 Architecture

528 Our normalizing flow architecture closely follows the one used in previous works [52, 65, 66]. We
 529 represent the conformations of the studied molecular systems using internal coordinates based on
 530 bond lengths, angles, and dihedral angles.

531 We use 8 pairs of neural spline coupling layers based on monotonic rational-quadratic splines [31].
 532 The splines map from $[0, 1]$ to $[0, 1]$ using 8 bins. We use a random mask to select transformed and
 533 conditioned dimensions in the first coupling of each pair, and the corresponding inverted mask for the
 534 second coupling. The dihedral angle dimensions are modeled with circular splines [63] to respect
 535 their topology, with a random (fixed) periodic shift applied after each coupling layer. The parameter
 536 networks that calculate the spline parameters in each coupling are fully connected neural networks
 537 with hidden dimensions [256, 256, 256, 256, 256] and ReLU activation functions. To capture their
 538 periodicity, dihedral angles ψ_i are encoded as $(\cos \psi_i, \sin \psi_i)$ when passing them to the parameter
 539 network.

540 As the base distribution of the normalizing flow, we use a uniform distribution in $[0, 1]$ for the dihedral
 541 angles and a Gaussian truncated to $[0, 1]$ with mean $\mu = 0.5$ and standard deviation $\sigma = 0.1$ for the
 542 bond lengths and angles.

543 We follow [66] to map the internal coordinates to the range $[0, 1]$ of the spline transformations:
 544 Dihedral angles are divided by 2π . Bond lengths and angles are shifted and scaled as $\eta'_i = (\eta_i -$
 545 $\eta_{i;\text{min}})/\sigma + 0.5$, where $\eta_{i;\text{min}}$ is obtained from a minimum energy structure after energy minimization.
 546 σ was set to 0.07 nm for bond lengths and 0.5730 for angle dimensions.

547 The studied molecular systems have two chiral forms (mirror images), L- and R-chirality, while in
 548 nature, one almost only finds the L-chirality. To constrain the generated molecular configurations
 549 to the L-chirality, we constrain the spline output ranges of the relevant dihedral angles (see [66]
 550 for details). Similarly, some atoms and groups (such as the hydrogen atoms in CH_3 groups) are
 551 permutation invariant in the force field energy parametrization, but have a preference in the ground
 552 truth molecular dynamics data due to very large barriers. Similarly to the chirality constraints, we
 553 constrain the splines such that only the permutation found in the ground truth data can be generated
 554 [66].

555 D.2 Target densities

556 The goal of all our experiments is to sample molecular systems at 300 K. An overview of the studied
 557 molecular systems, including their force field parametrization, is given in Table 4. We explicitly note
 558 that the largest studied system, ELIL, does not contain capping groups, in contrast to the other three
 559 systems.

560 The energy evaluations during training were performed with the OpenMM 8.0.0 [33] CPU platform,
 561 using 18 workers in parallel.

562 Following previous work [52, 66], we use a regularized energy function to avoid large van der Waals
 563 energies due to atom clashes:

$$E_{\text{reg.}}(E) = \begin{cases} E, & \text{if } E \leq E_{\text{high}}, \\ \log(E - E_{\text{high}} + 1) + E_{\text{high}}, & \text{if } E_{\text{high}} < E \leq E_{\text{max}}, \\ \log(E_{\text{max}} - E_{\text{high}} + 1) + E_{\text{high}}, & \text{if } E > E_{\text{max}}. \end{cases} \quad (15)$$

564 We set $E_{\text{high}} = 1 \times 10^8$ and $E_{\text{max}} = 1 \times 10^{20}$ [52].

NAME	SEQUENCE	NO. ATOMS	FORCE FIELD	CONSTRAINTS
ALANINE DIPEPTIDE	ACE-ALA-NME	22	AMBER FF96 WITH OBC1 IMPLICIT SOLVATION	NONE
ALANINE TETRAPEPTIDE	ACE-3·ALA-NME	42	AMBER99SB-ILDN WITH AMBER99 OBC IMPLICIT SOLVATION	HYDROGEN BOND LENGTHS
ALANINE HEXAPEPTIDE	ACE-5·ALA-NME	62	AMBER99SB-ILDN WITH AMBER99 OBC IMPLICIT SOLVATION	HYDROGEN BOND LENGTHS
ELIL	GLU-LEU-ILE-LEU	75	AMBER99SB-ILDN WITH AMBER99 OBC IMPLICIT SOLVATION	HYDROGEN BOND LENGTHS

Table 4: Overview of the molecular systems and corresponding force field parametrization.

565 Ground truth datasets

566 We use ground truth test datasets obtained from extensive molecular dynamics simulations to calculate
567 the metrics reported in Table 1.

568 1. For alanine dipeptide, we use the ground truth test dataset previously published by [52] [70].
569 Furthermore, we use the datasets published by [66] as additional validation and training datasets
570 (for the forward KL experiments).

571 2. For alanine tetrapeptide and alanine hexapeptide, we use the test, validation, and training datasets
572 published by [66].

573 3. To generate ground truth data for the tetrapeptide ELIL, we followed the simulation protocol by
574 [66]. We performed two replica-exchange molecular dynamics simulations with replicas at 300 K,
575 332 K, 368 K, 408 K, 451 K, and 500 K, each using a time step of 2 fs. Each simulation used
576 200 ns equilibration without exchanges, 200 ns equilibration with exchanges, and 1 μ s production
577 simulation time. One of the simulations was used for the ground truth test dataset, the other
578 simulation for the training and validation datasets.

579 For each system, the test dataset contains 1×10^7 samples, and the training and validation datasets
580 contain 1×10^6 samples.

581 D.3 Dual optimization in practice

582 **Monte Carlo estimation.** Maximizing the concave dual function (7) requires evaluating intermediate
583 normalization constants \mathcal{Z}_{i+1} . This can be done efficiently by expressing \mathcal{Z}_{i+1} as an expectation
584 under q_i and using Monte Carlo estimation. For instance, the expression for \mathcal{Z}_{i+1} in (5) can be
585 estimated as

$$\mathcal{Z}_{i+1}(\lambda, \eta) = \mathbb{E}_{x \sim q_i} \left[\left(\frac{p(x)}{q_i(x)^{1+\eta}} \right)^{\frac{1}{1+\lambda+\eta}} \right] \approx \frac{1}{N} \sum_{x_n \sim q_i} \left(\frac{p(x_n)}{q_i(x_n)^{1+\eta}} \right)^{\frac{1}{1+\lambda+\eta}}. \quad (16)$$

586 We note that samples $x_n \sim q_i$ and the corresponding evaluations $q_i(x_n)$ and $p(x_n)$ are typically
587 already computed when solving (9), so the additional cost of determining the Lagrange multipliers is
588 minimal.

589 **Optimization.** The concavity of the dual functions permits the use of any suitable nonlinear opti-
590 mization algorithm. For one-dimensional dual optimization, we employ the bounded Brent method
591 [19], implemented via `scipy.optimize.minimize_scalar` [79], which is the library’s default 1D
592 algorithm due to its robustness and efficiency. A minimal working example on how a Lagrangian mul-
593 tiplier is estimated is given in Code Example 1. For 2D duals, we use `scipy.optimize.minimize`
594 with the L-BFGS-B algorithm [86], one of SciPy’s default quasi-Newton algorithms. There, we
595 additionally passed the dual gradient function, which we obtained through automatic differentiation.
596 Due to the constraints $\lambda, \eta \geq 0$, and to avoid numerical overflow, we bound both optimizers to stay
597 within the interval $[0, 10^{10}]$. The method `scipy.optimize.minimize` requires an initial guess,
598 which we set to 1×10^{-20} , a value chosen to be close to the lower bound.

Code Example 1: Minimal working example of the dual optimization for objective (2).

```

1  import numpy as np
2  import torch
3  from scipy.optimize import minimize_scalar
4
5  def estimate_log_Z(
6      model_log_prob: torch.Tensor,
7      target_log_prob: torch.Tensor,
8      tr_mul: float,
9  ) -> torch.Tensor:
10     """Estimate log-partition function of next intermediate density"""
11     log_N = torch.tensor(target_log_prob.shape[0]).log()
12     log_iw = (target_log_prob - model_log_prob) / (1 + tr_mul)
13     log_Z = torch.logsumexp(log_iw, dim=0) - log_N
14     return log_Z
15
16
17  def find_best_kl_multiplier(
18      model_log_prob: torch.Tensor,
19      target_log_prob: torch.Tensor,
20      eps_tr: float,
21      max_multiplier: float = 1e10,
22  ) -> float:
23     """Finds the best Lagrangian multiplier by maximizing the dual"""
24     # define dual function (dependent on Lagrangian multiplier)
25     def dual(tr_mul: float):
26         log_Z = estimate_log_Z(
27             model_log_prob=model_log_prob,
28             target_log_prob=target_log_prob,
29             tr_mul=tr_mul,
30         )
31         dual_value = -(1 + tr_mul) * log_Z - tr_mul * eps_tr
32         return dual_value.item()
33
34     neg_dual = lambda mul: -dual(mul) # concave -> convex
35
36     res = minimize_scalar(
37         neg_dual,
38         bounds=(0.0, max_multiplier),
39         method="Bounded"
40     )
41     best_tr_mul: float = res.x
42     return float(best_tr_mul)

```

599 D.4 Metrics

600 In this section, we present several commonly used metrics for both theoretical analysis and experi-
601 mental evaluation. For more details, we refer to [15].

602 NEGATIVE LOG-LIKELIHOOD (NLL)

603 The negative-log-likelihood (NLL) is a forward metric computed between ground truth samples and
604 the learned model distribution \hat{q} . It is directly related to the forward KL divergence and the evidence
605 upper bound (EUBO) via

$$\begin{aligned}
 D_{\text{KL}}(p|q) &= \underbrace{\mathbb{E}_{p(x)} \left[\log \frac{\tilde{p}(x)}{q(x)} \right]}_{\text{EUBO}} - \underbrace{\log \mathcal{Z}}_{\text{const. w.r.t. } q} \\
 &= \underbrace{-\mathbb{E}_{p(x)} [\log q(x)]}_{\text{NLL}} - \underbrace{H(p)}_{\text{const. w.r.t. } q} .
 \end{aligned}$$

606 The metrics NLL, EUBO and forward KL therefore only deviate from each other by an additive
607 constant.

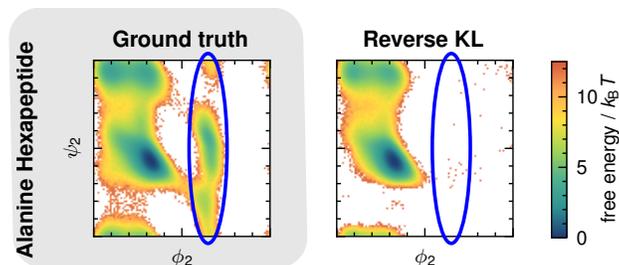


Figure 5: Illustration of mode collapse in Ramachandran plots. The figure highlights high-energy regions where mode collapse often occurs (e.g., under reverse KL training) but which remain important for accurately capturing the peptide backbone distribution. Mode collapse is visible as the absence of density in these regions.

608 EFFECTIVE SAMPLE SIZE (ESS)

609 The effective sample size (ESS) is defined as

$$\text{ESS}(a, b) = \frac{1}{1 + \text{Var}_{a(x)} \left[\frac{b(x)}{a(x)} \right]}, \quad a, b \in \mathcal{P}(\mathbb{R}^d).$$

610 Closely following the notation of [15], the reverse ESS

$$\text{ESS}(q, p) = \frac{\mathcal{Z}_r}{\mathbb{E}_{q(x)} \left[\left(\frac{\tilde{p}(x)}{q(x)} \right)^2 \right]}, \quad \text{with } \mathcal{Z}_r = \mathbb{E}_{q(x)} \left[\frac{\tilde{p}(x)}{q(x)} \right]$$

611 can be directly estimated via Monte Carlo using samples from the model q and the unnormalized
612 target \tilde{p} .

613 Following [52, 66], we clipped the top 0.01% importance-weights, setting them to the smallest value
614 among them for numerical reasons. Furthermore, reverse ESS is computed using the regularized
615 energy function, defined in Eq. (15).

616 Although forward ESS could be computed using samples from the target distribution, [66] found it to
617 be extremely sensitive to the chosen clipping threshold and prone to instability. Consequently, only
618 the reverse ESS was used, even though it may not fully capture phenomena such as mode collapse.

619 RAMACHANDRAN PLOTS

620 A Ramachandran plot visualizes the 2D log-density of the joint distribution of a pair of dihedral
621 angles in a peptide’s backbone. For more details, we refer to [66]. These plots are used to visualize a
622 peptide’s main degrees of freedom and are likely to show mode collapse if it occurs. A Ramachandran
623 plot is effectively a histogram of the occurrence of dihedral angles and is computed solely from model
624 or ground-truth samples.

625 How mode collapse can be detected and where the high-energy density region is located in a
626 Ramachandran plot is illustrated in App. D.4.

627 For alanine tetrapeptide, alanine hexapeptide, and ELIL tetrapeptide, which contain multiple backbone
628 dihedral angle pairs, we always show the pair exhibiting the most pronounced deviation from the
629 ground truth, which is the same across methods. Among the four runs made per method in Fig. 2, we
630 selected the one with the lowest Ram KL value.

631 RAMACHANDRAN KL DIVERGENCE (RAM KL AND RAM KL W.Rw.)

632 To obtain quantitative estimates of the quality of such Ramachandran plots, we used two metrics,
633 following the approach of [52, 66]. We computed the forward KL divergence between the Ramachan-
634 dran plots from ground truth and model samples (RAM KL). For this, we used 100×100 bins and
635 1×10^7 samples. Additionally, we also calculated a reweighted version of the metric (Ram KL
636 w.Rw.) where the model samples were first reweighted to the target distribution before generation of
637 Ramachandran plots.

638 For the larger systems, where more than one Ramachandran plot exists, we reported the average
639 Ramachandran KL.

640 **D.5 Hyperparameters**

641 Hyperparameters play a crucial role in the performance of all models. Common hyperparameters
 642 include the choice of optimizer, learning rate, batch size, gradient steps, and weight decay. Below,
 643 we provide a description of the hyperparameters for each method, emphasizing any method-specific
 644 choices.

645 All experiments employed the Adam optimizer [43]. Our implementation builds on the Python
 646 packages *bgflow* [57], *nflows* [32], and *PyTorch* [61]. The number of parameters in the normalizing
 647 flow architecture for each system is summarized in Tab. 5.

Table 5: Number of flow parameters for each system. The number of parameters is completely determined by a molecular system’s size, as the architecture is the same across all systems.

	ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
NUMBER OF PARAMETERS	7 421 512	9 452 376	12 124 616	13 727 952

648 **CMT**

649 We refer to Tab. 6 for the general and method-specific hyperparameters of CMT.

	HYPERPARAMETERS	ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	BATCH SIZE	1000	1000	2000	2000
	LEARNING RATE	4×10^{-5}	5×10^{-5}	5×10^{-5}	5×10^{-5}
	LR SCHEDULER	COSINE	COSINE	COSINE	COSINE
	GRADIENT DESCENT STEPS	400 000	400 000	800 000	1 600 000
	WEIGHT-DECAY	1×10^{-5}	1×10^{-5}	1×10^{-5}	1×10^{-5}
	LR LINEAR WARMUP STEPS	1000	1000	1000	1000
METHOD-SPECIFIC	MAX GRAD NORM	100.0	100.0	100.0	100.0
	TRUST-REGION BOUND	0.3	0.3	0.3	0.3
	ENTROPY BOUND	0.8	1.8	1.4	0.7
	BUFFER SIZE	500 000	500 000	1 000 000	1 000 000
	GRADIENT DESCENT STEPS PER ANNEALING STEP	2000	2000	2000	2000

Table 6: Hyperparameter settings for CMT (general and method-specific) for all systems.

650 **TA-BG**

651 Tab. 7 summarizes the hyperparameters for the pre-training of TA-BG [66] using the reverse KL
 652 divergence.

	HYPERPARAMETERS	ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	TARGET TEMPERATURE	1200 K	1200 K	1200 K	3000 K
	BATCH SIZE	256	256	512	512
	LEARNING RATE	1×10^{-4}	1×10^{-4}	1×10^{-4}	1×10^{-4}
	LR SCHEDULER	COSINE	COSINE	COSINE	COSINE
	GRADIENT DESCENT STEPS	100 000	100 000	250 000	250 000
	WEIGHT-DECAY	1×10^{-5}	1×10^{-5}	1×10^{-5}	1×10^{-5}
	LR LINEAR WARMUP STEPS	1000	1000	1000	1000
	MAX GRAD NORM	100.0	100.0	100.0	100.0
	NO. HIGHEST ENERGY VALUES REMOVED	10	10	20	20

Table 7: Hyperparameter settings for TA-BG pre-training for all systems.

653 After pre-training, the temperature is annealed with a geometrically decaying temperature sequence
 654 and the hyperparameters summarized in Tab. 8. The TA-BG experiments on alanine dipeptide and
 655 alanine tetrapeptide used the geometric temperature annealing sequence

$$1200 \text{ K} \rightarrow 1028.69 \text{ K} \rightarrow 881.84 \text{ K} \rightarrow 755.95 \text{ K} \rightarrow 648.04 \text{ K} \rightarrow 555.52 \text{ K} \\ \rightarrow 476.22 \text{ K} \rightarrow 408.24 \text{ K} \rightarrow 349.96 \text{ K} \rightarrow 300.00 \text{ K} \rightarrow 300.00 \text{ K}.$$

656 Including an additional finetuning step per temperature, TA-BG employs the temperature sequence

$$1200 \text{ K} \rightarrow 1028.69 \text{ K} \rightarrow 1028.69 \text{ K} \rightarrow 881.84 \text{ K} \rightarrow 881.84 \text{ K} \rightarrow 755.95 \text{ K} \\ \rightarrow 755.95 \text{ K} \rightarrow 648.04 \text{ K} \rightarrow 648.04 \text{ K} \rightarrow 555.52 \text{ K} \rightarrow 555.52 \text{ K} \rightarrow 476.22 \text{ K} \\ \rightarrow 476.22 \text{ K} \rightarrow 408.24 \text{ K} \rightarrow 408.24 \text{ K} \rightarrow 349.96 \text{ K} \rightarrow 349.96 \text{ K} \rightarrow 300.00 \text{ K} \rightarrow 300.00 \text{ K}$$

657 on alanine hexapeptide. On ELIL, reverse KL pre-training suffers from mode-collapse at 1200 K.
 658 Therefore, the temperature annealing starts at 3000 K, resulting in the temperature sequence

3000.00 K → 2573.09 K → 2573.09 K → 2573.09 K → 2206.93 K → 2206.93 K
 → 2206.93 K → 1892.88 K → 1892.88 K → 1892.88 K → 1623.52 K → 1623.52 K
 → 1623.52 K → 1392.49 K → 1392.49 K → 1392.49 K → 1194.33 K → 1194.33 K
 → 1194.33 K → 1024.37 K → 1024.37 K → 1024.37 K → 878.60 K → 878.60 K
 → 878.60 K → 753.57 K → 753.57 K → 753.57 K → 646.34 K → 646.34 K
 → 646.34 K → 554.36 K → 554.36 K → 554.36 K → 475.48 K → 475.48 K
 → 475.48 K → 407.81 K → 407.81 K → 407.81 K → 349.78 K → 349.78 K
 → 349.78 K → 300.00 K → 300.00 K.

HYPERPARAMETERS		ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	BATCH SIZE	2048	4096	2048	2048
	LEARNING RATE	5×10^{-6}	1×10^{-5}	5×10^{-6}	5×10^{-6}
	LR SCHEDULER	COSINE	-	-	-
	GRADIENT DESCENT STEPS	(PER TEMPERATURE STEP) 300 000	200 000	360 000	600 000
METHOD-SPECIFIC	BUFFER SIZE	7 440 000	7 440 000	15 111 111	22 400 000
	BUFFER RESAMPLED TO	2 000 000	2 000 000	2 000 000	10 000 000
	GRADIENT DESCENT STEPS PER ANNEALING STEP	30 000	20 000	20 000	20 000

Table 8: Hyperparameter settings for TA-BG (general and method-specific) for all systems.

659 **FAB**

660 The used hyperparameters for FAB [52] can be found in Tab. 9. Furthermore, we used a step
 661 size of 0.05 for the Hamiltonian Monte Carlo [30] transitions. For details on the method and its
 662 hyperparameters, we refer to [52].

HYPERPARAMETERS		ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	BATCH SIZE	1024	1024	1024	2048
	LEARNING RATE	1×10^{-4}	1×10^{-4}	1×10^{-4}	2×10^{-4}
	LR SCHEDULER	COSINE	COSINE	COSINE	COSINE
	GRADIENT DESCENT STEPS	50 000	50 000	50 000	25 000
	WEIGHT-DECAY	1×10^{-5}	1×10^{-5}	1×10^{-5}	1×10^{-5}
	LR LINEAR WARMUP STEPS	1000	1000	1000	1000
METHOD-SPECIFIC	MAX GRAD NORM	1000.0	1000.0	1000.0	1000.0
	NO. INTERMED. DIST.	8	8	8	16
	NO. INNER HMC STEPS	4	4	8	8

Table 9: Hyperparameter settings of FAB (general and method-specific) for all systems.

663 **Forward and Reverse KL**

664 This section reports the used hyperparameters for training with the forward KL divergence on MD
 665 data (Tab. 10) and the hyperparameters for training with the reverse KL divergence (Tab. 11). A
 666 description on how the MD data was obtained can be found in App. D.2.

HYPERPARAMETERS		ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	BATCH SIZE	1024	1024	1024	1024
	LEARNING RATE	5×10^{-5}	5×10^{-5}	5×10^{-5}	5×10^{-5}
	LR SCHEDULER	COSINE	COSINE	COSINE	COSINE
	GRADIENT DESCENT STEPS	100 000	100 000	120 000	140 000

Table 10: Hyperparameter settings of forward KL training using MD data for all systems.

HYPERPARAMETERS		ALANINE DIPEPTIDE	ALANINE TETRAPEPTIDE	ALANINE HEXAPEPTIDE	ELIL
GENERAL	BATCH SIZE	1024	1024	1024	1024
	LEARNING RATE	1×10^{-4}	1×10^{-4}	1×10^{-4}	1×10^{-4}
	LR SCHEDULER	COSINE	COSINE	COSINE	COSINE
	GRADIENT DESCENT STEPS	250 000	250 000	250 000	250 000
	WEIGHT-DECAY	1×10^{-5}	1×10^{-5}	1×10^{-5}	1×10^{-5}
	LR LINEAR WARMUP STEPS	1000	1000	1000	1000
	MAX GRAD NORM	100.0	100.0	100.0	100.0
	NO. HIGHEST ENERGY VALUES REMOVED	40	40	40	40

Table 11: Hyperparameter settings of reverse KL training for all systems.