
Transferable Boltzmann Generators

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The generation of equilibrium samples of molecular systems has been a long-
2 standing problem in statistical physics. Boltzmann Generators are a generative
3 machine learning method that addresses this issue by learning a transformation via
4 a normalizing flow from a simple prior distribution to the target Boltzmann distri-
5 bution of interest. Recently, flow matching has been employed to train Boltzmann
6 Generators for small molecular systems in Cartesian coordinates. We extend this
7 work and propose a first framework for Boltzmann Generators that are transferable
8 across chemical space, such that they predict zero-shot Boltzmann distributions
9 for test molecules without being retraining for these systems. These transferable
10 Boltzmann Generators allow approximate sampling from the target distribution
11 of unseen systems, as well as efficient reweighting to the target Boltzmann distri-
12 bution. The transferability of the proposed framework is evaluated on dipeptides,
13 where we show that it generalizes efficiently to unseen systems. Furthermore, we
14 demonstrate that our proposed architecture enhances the efficiency of Boltzmann
15 Generators trained on single molecular systems.

16 1 Introduction

17 Generative models have demonstrated remarkable success in the physical sciences, including protein
18 structure prediction [1, 2, 3], generation of de novo molecules [4, 5, 6], and efficiently generating
19 samples from the Boltzmann distribution [7, 8, 9]. In this work, we will focus on the latter for
20 molecular systems, which represents a promising avenue for addressing the sampling problem. The
21 sampling problem refers to the long-standing challenge in statistical physics to generate samples
22 from equilibrium Boltzmann distributions $\mu(x) \propto \exp(-U(x)/k_B T)$, where $U(x)$ is the potential
23 energy of the system, k_B the Boltzmann constant, and T the temperature. Traditionally, samples are
24 generated with sequential sampling algorithms such as Markov Chain Monte Carlo and Molecular
25 Dynamics (MD) simulations. However, these algorithms require a significant amount of time to
26 generate uncorrelated samples from the target distribution. This is due to the necessity of performing
27 small update steps, in the order of femtoseconds, for stability. This is especially challenging in
28 the presence of well-separated metastable states, where transitions are unlikely due to high energy
29 barriers. In recent years, numerous machine learning methods have emerged to address this challenge
30 [7, 10]. One such method is the Boltzmann Generators (BG) [7]. In this work, we refer to BGs
31 as a model that allows for the approximate sampling of the Boltzmann distribution of interest and
32 the subsequent reweighting to the unbiased target distribution. If the model is only capable of
33 generating approximate samples, which may stem from a subset of the Boltzmann distribution, we
34 refer to them as Boltzmann Emulators¹. Boltzmann Generators transform an often simple, prior
35 distribution via a normalizing flow [11, 12, 13, 14] to an approximation of the target Boltzmann
36 distribution. Subsequently, generated samples can be reweighted to the unbiased target distribution.
37 The effectiveness of the reweighting depends on how close the generated distribution matches the

¹To the best of our knowledge Bowen Jing introduced the name first.

38 target distribution. Hence, it is possible to generate uncorrelated and unbiased samples from the target
39 Boltzmann distribution, potentially generating significant speed-up over classical MD simulations.
40 There are many ways to build a BG, because of the various available realizations of normalizing flows.
41 In this work, we will focus on continuous normalizing flows (CNFs) [15, 16], rather than coupling
42 flows [17]. Recently, flow matching [18, 19, 20, 21] emerged as an alternative training method for
43 CNFs, which is simulation free, allowing for more efficient training of CNFs.

44 Thus far, Boltzmann Generators have been found to be limited by the necessity of training them on
45 the system of interest. This training process, which requires a significant amount of time, makes
46 it challenging to achieve any significant speed-up over classical MD simulations. Furthermore,
47 the training time must be taken into account, which may even necessitate the execution of MD
48 simulations of the system of interest on its own. It is therefore desirable to have a transferable
49 Boltzmann Generator that can be trained on one set of molecules and generalize to another set, where
50 Boltzmann samples can be efficiently generated at inference time without retraining. Only recently,
51 reliable Boltzmann Generators in Cartesian coordinates for molecules were introduced [22, 23],
52 which paved the way for transferable Boltzmann Generators, as they do not depend on the molecule
53 specific internal coordinate representation, which make it difficult to construct transferable models.

54 In this work, we introduce a framework for *transferable* Boltzmann Generators based on CNFs,
55 allowing effective sample generation from unseen Boltzmann distributions. Transferable Boltzmann
56 Generators are desirable, as they do not require retraining for similar systems and can be trained on
57 short training trajectories that miss metastable states. The Boltzmann Generator can still learn these
58 from other similar trajectories of other systems.

59 We make the following main contributions:

- 60 1. We introduce, to the best of our knowledge, the first *transferable* Boltzmann Generator.
61 We demonstrate the transferability on dipeptides, where we are able to generate unbiased
62 samples from Boltzmann distributions of unseen dipeptides.
- 63 2. We describe a general framework for training and sampling with transferable Boltzmann
64 Generators based on continuous normalizing flows. This includes also the post-processing
65 of generated samples.
- 66 3. We perform several ablation studies to investigate the effect of different architectures,
67 training set sizes, as well as biasing the training data. The results demonstrate that small
68 training sets can be sufficient to train transferable Boltzmann Generators.

69 2 Related work

70 The initial work on Boltzmann Generators [7] has led to a great deal of subsequent research. The
71 most common application of BGs is to generate samples from Boltzmann distributions of molecules
72 [24, 25, 26, 27, 28, 29], as well as lattice systems [25, 30, 31]. Most BGs for molecular systems
73 require system-specific featurizations such as internal coordinates [7, 32, 25, 26, 33, 34, 29]. Only
74 recently, BGs for small molecular systems in Cartesian coordinates were introduced [22, 23], using
75 CNFs and coupling flows, respectively. Equivariant normalizing flows [35, 36, 37, 27, 22, 38] played
76 a pivotal role in the success of Boltzmann Generators in Cartesian coordinates, not only for molecular
77 systems. The majority of BGs employ a Gaussian prior distribution, but it is also possible to start
78 from prior distributions close to the target distribution [39, 34, 40], which makes the learning task
79 simpler. However, all previous Boltzmann Generators are not transferable. Arguably, the work of [6]
80 represents an exception, as they are able to generate samples from unseen conditional (Boltzmann)
81 distributions in torsion space. However, the distribution is conditioned on a single local structure
82 for each molecule, namely fixed bonds and angles. Consequently, in contrast to our work, they
83 are unable to generate samples from the full Boltzmann distribution in Euclidean space. The first
84 transferable deep generative model that was able to generate asymptotically unbiased samples from
85 the Boltzmann distribution is [10]. Instead of generating independent samples, they learn large
86 time steps and combine these with Metropolis-Hastings acceptance steps, to ensure asymptotically
87 unbiased samples. However, in contrast to our work, they do not generate uncorrelated samples.

88 Boltzmann Emulators are analogous to Boltzmann Generators, yet they are not designed to generate
89 unbiased equilibrium samples from the target Boltzmann distribution. Instead, they are intended
90 to generate approximate samples that do not undergo reweighting. Furthermore, the generation of
91 all metastable states may not be a necessary requirement, depending on the system. Boltzmann

92 Emulators do not need to be based on flow models, as they do not aim to do reweighting to the
 93 target distribution. They are often similar to Boltzmann Generators and use normalizing flows or
 94 diffusion models for the architecture, but due to removing the constraint of sampling the unbiased
 95 Boltzmann distribution, they can target significantly larger systems or are transferable. One example
 96 is [41], who propose a three stage transferable CNF model to learn peptide ensembles. [42] use
 97 flow matching to learn distributions of proteins, while [43] utilize diffusion models. [44] build
 98 a transferable Boltzmann Emulator for small molecules. Others aim to additionally also capture
 99 the correct dynamics of the molecular systems, such as [45], who use a diffusion model to predict
 100 transition probabilities. Scaling to larger systems often requires coarse graining [46, 47, 42], e.g.
 101 describing amino acids by a single bead rather than the individual atoms. However, this approach
 102 precludes the possibility of reweighting to the Boltzmann distribution.

103 A distinct, though related, learning objective is to generate novel molecular conformations. However,
 104 approximations from the Boltzmann distribution are not necessary; it is sufficient to generate a few
 105 (or even a single) conformation per molecule. The utilized architectures are once again analogous, as
 106 flow and diffusion models are employed [4, 5, 48, 49, 50, 6].

107 3 Boltzmann Generators and Normalizing Flows

108 Here, we describe Boltzmann Generators and normalizing flows, which are a central part of our
 109 proposed transferable Boltzmann Generator framework. We follow the notation of [22].

110 3.1 Boltzmann Generators

111 Boltzmann Generators (BGs) [7] combine an exact likelihood deep generative model and a reweight-
 112 ing algorithm to reweight the generated distribution to the target Boltzmann distribution. The exact
 113 likelihood deep generative model is trained to generate samples from a distribution $\tilde{p}(x)$ that is close
 114 to the target Boltzmann distribution $\mu(x)$. A common choice for the exact likelihood model are
 115 normalizing flows.

116 The Boltzmann Generator can be used to generate unbiased samples by first sampling $x \sim \tilde{p}(x)$ with
 117 the exact likelihood model and then computing corresponding importance weights $w(x) = \mu(x)/\tilde{p}(x)$
 118 for each sample. These allow to reweight generated samples to the target Boltzmann distribution
 119 $\mu(x)$. It is possible to estimate observables of interest (asymptotically unbiased) using the weights
 120 $w(x)$ with importance sampling via

$$\langle O \rangle_\mu = \frac{\mathbb{E}_{x \sim \tilde{p}(x)}[w(x)O(x)]}{\mathbb{E}_{x \sim \tilde{p}(x)}[w(x)]}. \quad (1)$$

121 Furthermore, these reweighting weights can be employed to assess the efficiency of trained BGs by
 122 computing the effective sample size (ESS) with Kish’s equation [51]. In this work, we will compute
 123 the relative ESS, rather than the absolute one, and refer to it as ESS.

124 3.2 Continuous Normalizing Flows (CNFs)

125 Normalizing flows [14, 52] are a type of deep generative model used to learn complex probability
 126 densities $\mu(x)$ by transforming a prior distribution $q(x)$ through an invertible transformation $f_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^n$, resulting in the push-forward distribution $\tilde{p}(x)$.

128 Continuous Normalizing Flows (CNFs) [15, 16] are a specific kind of normalizing flow. In CNFs, the
 129 invertible transformation $f_\theta^t(x)$ is defined by the ordinary differential equation

$$\frac{df_\theta^t(x)}{dt} = v_\theta(t, f_\theta^t(x)), \quad f_\theta^0(x) = x_0, \quad (2)$$

130 where $v_\theta(t, x) : \mathbb{R}^n \times [0, 1] \rightarrow \mathbb{R}^n$ is a time-dependent vector field. The solution to this initial value
 131 problem provides the transformation equation

$$f_\theta^t(x) = x_0 + \int_0^t dt' v_\theta(t', f_\theta^{t'}(x)), \quad (3)$$

132 with $f_\theta^1(x) = \tilde{p}^t(x)$. The corresponding change in log density from the prior to the push-forward
 133 distribution is described by the continuous change of variable equation

$$\log \tilde{p}(x) = \log q(x) - \int_0^1 dt \nabla \cdot v_\theta(t, f_\theta^t(x)). \quad (4)$$

134 **Equivariant flows** The energy of molecular systems is typically invariant under permutations of
 135 interchangeable particles and global rotations and translations. Consequently, it is advantageous for
 136 the push-forward distribution of a Boltzmann generator to possess the same symmetries as the target
 137 system. In [35, 36] the authors demonstrate that the push-forward distribution $\tilde{p}(x)$ of a permutation
 138 and rotation equivariant normalizing flow with a permutation and rotation invariant prior distribution,
 139 is again rotation and permutation invariant. Furthermore, [35] present a method to construct such
 140 equivariant CNFs by using an equivariant vector field v_θ .

141 3.3 Flow matching

142 Flow matching [18, 19, 20, 21] enables efficient, simulation-free training of CNFs. The conditional
 143 flow matching training objective allows for the direct training of the vector field $v_\theta(t, x)$ through

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t \sim [0,1], x \sim p_t(x|z)} \|v_\theta(t, x) - u_t(x|z)\|_2^2. \quad (5)$$

144 There are many possible parametrizations for the conditional vector field $u_t(x|z)$ and the conditional
 145 probability path $p_t(x|z)$. One of the most simple, but powerful possible parametrization is

$$z = (x_0, x_1) \quad \text{and} \quad p(z) = q(x_0)\mu(x_1) \quad (6)$$

$$u_t(x|z) = x_1 - x_0 \quad \text{and} \quad p_t(x|z) = \mathcal{N}(x|t \cdot x_1 + (1-t) \cdot x_0, \sigma^2), \quad (7)$$

146 which we use in this work to train our models. For a more detailed description refer to [18, 21, 22, 48].

147 4 Transferable Boltzmann Generators

148 This section presents our proposed framework for transferable Boltzmann Generators (TBGs).

149 4.1 Architecture

150 Our proposed transferable Boltzmann Generator is based on a CNF. The corresponding vector
 151 field $v_\theta(t, x)$ is parametrized by an $O(D)$ - and $S(N)$ -equivariant graph neural network (EGNN)
 152 [37, 53, 41], as commonly used in prior work, e.g. [22, 37]. Although, less expressive than other
 153 equivariant networks such as [54, 55, 56, 57], it is faster to evaluate, which is important for CNFs as
 154 there can be hundreds of vector field calls during inference.

155 The vector field $v_\theta(t, x)$ consists of L consecutive layers. The position of the i -th particle x_i is
 156 updated according to the following equations:

$$h_i^0 = (t, a_i, b_i, c_i), \quad m_{ij}^l = \phi_e(h_i^l, h_j^l, d_{ij}^2), \quad (8)$$

$$x_i^{l+1} = x_i^l + \sum_{j \neq i} \frac{(x_i^l - x_j^l)}{d_{ij} + 1} \phi_d(m_{ij}^l), \quad (9)$$

$$h_i^{l+1} = \phi_h(h_i^l, m_i^l), \quad m_i^l = \sum_{j \neq i} \phi_m(m_{ij}^l) m_{ij}^l, \quad (10)$$

$$v_\theta(t, x^0)_i = x_i^L - x_i^0 - \frac{1}{N} \sum_j (x_i^L - x_j^0), \quad (11)$$

157 where ϕ_α represents different neural networks, d_{ij} is the Euclidean distance between particle i and
 158 j , t is the time, a_i is an embedding for the particle type, b_i for the amino acid, and c_i or the amino
 159 acid position in the peptide. In the final step, the geometric center is subtracted to ensure that the
 160 center of positions is conserved. When combined with a symmetric mean-free prior distribution, the
 161 push-forward distribution of the CNF will be $O(D)$ - and $S(N)$ -invariant, as demonstrated in [58].

162 The embedding of each atom is constructed from three parts. The first part is the atom type a_i , which
163 is a one-hot vector of 54 classes. The classes are defined by the atom type in the topology for a
164 classical force field. Therefore, only a few atoms are indistinguishable, such as hydrogen atoms that
165 are bound to the same carbon or nitrogen atom. The second part is the amino acid to which the atom
166 belongs, which is divided into 20 classes. The third part is the position of the amino acid in the peptide
167 sequence. This embedding is similar to the embedding used in [41] for the rotamer embeddings.
168 The amino acid and positional embeddings are only used for the transferable experiments. For more
169 details see Appendix B.5. In this study, we refer to this transferable Boltzmann Generator architecture
170 as *TBG + full*, and we use this name even when we apply it to a non-transferable setting.

171 The proposed architecture in [22] employs distinct encodings for all backbone atoms and the atom
172 types for the remainder. This represents a special case of our architecture, wherein b_i and c_i are
173 omitted and a_i is simply the atom types or the backbone encoding. We refer to this architecture
174 as *TBG + backbone*. Furthermore, we refer to the specific architecture employed in [22] as *BG +*
175 *backbone* for the alanine dipeptide experiments.

176 Moreover, we employ a model that utilizes the atom type as the sole encoding (there are only five
177 distinct atom types). This model is referred to as simply *TBG*.

178 4.2 Training transferable Boltzmann Generators

179 All transferable Boltzmann Generators are trained with flow matching. As there are different peptides
180 in each batch, the individual flow matching loss is divided by the number of atoms in each peptide.
181 For the 2AA dataset, all training peptides in each batch are used. All different architectures are
182 trained in the same way; for more details, see Appendix B.

183 4.3 Inference with transferable Boltzmann Generators

184 Sampling with a transferable Boltzmann Generator, especially on unseen peptides, poses multiple
185 challenges: (i) Some generated samples may not correspond to the molecule of interest, but rather to
186 a molecule that contains the same atom number and types but has a different bonding graph. This is
187 because the model has never encountered such a configuration during training. These configurations
188 might even have much lower quantum Chemical potential energies than the molecule of interest. For
189 some examples see Appendix A.4. However, as we are in this work interested in sampling from
190 the equilibrium Boltzmann distribution for a given molecular bonding graph, rather than sampling
191 distinct molecules, we would like to avoid these cases. Nevertheless, this effect can be largely
192 mitigated by our proposed *TBG + full* architecture. (ii) When working with classical force fields,
193 the correct ordering with respect to topology is crucial for evaluating energies. This is not a concern
194 for semi-empirical force fields, as they respect the permutation symmetry of particles of the same
195 type. As we typically use Gaussian prior distribution, it is common that the generated samples are
196 not arranged according to the topology. Consequently, in order to evaluate the energy, it is necessary
197 to reorder the generated samples according to the topology. (iii) It is possible that the chirality of
198 generated samples may differ from that of the peptide of interest. This can be rectified by simple
199 mirroring if all chirality centers of a peptide are flipped. Otherwise, these samples are assigned high
200 energies, as they are not from the target Boltzmann distribution of interest.

201 We resolve (i) and (ii) by generating a bond graph for the generated samples, based on empirical
202 bond distances and atom types. This graph is then compared with a reference bond graph. If the
203 two graphs are isomorphic, we can conclude that the configuration is correct. For more details, see
204 Appendix B.1. For (iii), we employ the code of [10] to check all chiral centers. It should be noted that
205 only if the generated samples possess the correct configuration and chirality will they be considered
206 valid samples from the Boltzmann distribution of interest.

207 5 Experiments

208 In this section, we compare our model with similar previous work [22] on equivariant Boltzmann
209 Generators for alanine dipeptide. Moreover, we show the transferability of our model on dipeptides.
210 More experimental details, such as dataset details, the specifics of the employed models, and the
211 utilized computing infrastructure can be found in Appendix B.

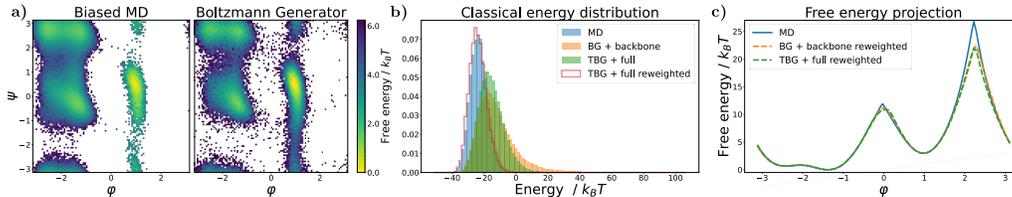


Figure 1: Results for the alanine dipeptide system simulated with a classical force field (a) Ramachandran plots for the biased MD distribution (left) and for samples generate with the TBG + full model (right). (b) Energies of samples generated with different methods. (c) Free energy projection along the slowest transition (φ angle), computed with different methods.

Table 1: Comparison of Boltzmann Generators with different architectures for the single molecular system alanine dipeptide. Errors are computed over five runs. The results for Boltzmann Generator and backbone encoding (BG + backbone) for the semi-empirical force field are taken from [22].

Model	NLL (\downarrow)	ESS (\uparrow)
Alanine dipeptide - semi-empirical force field		
BG + backbone [22]	-107.56 ± 0.09	$0.50 \pm 0.13\%$
TBG + full (ours)	-124.71 ± 0.08	$1.03 \pm 0.17\%$
Alanine dipeptide - classical force field		
BG + backbone [22]	-109.02 ± 0.01	$1.56 \pm 0.30\%$
TBG + full (ours)	-127.06 ± 0.12	$6.03 \pm 1.34\%$

212 5.1 Alanine dipeptide

213 In our first experiment, we investigate the single molecule alanine dipeptide in implicit solvent,
 214 described in Cartesian coordinates. The dataset was introduced in [22], for more details see Ap-
 215 pendix B.2. The training trajectory was generated by sampling with respect to a classical force field,
 216 and subsequently, 10^5 random samples were relaxed with respect to the semi-empirical *GFN2-xTB*
 217 force-field [59] for 100fs each. The objective is to train a Boltzmann Generator capable of sampling
 218 from the equilibrium Boltzmann distribution defined by the semi-empirical *GFN2-xTB* force-field
 219 efficiently and to recover the free energy surface along the slowest transition, i.e. the φ angle.
 220 Following the methodology outlined in [22], the training data is biased towards the less probable
 221 (positive) φ state. It is evident that any trained model on this set will be biased in comparison to
 222 the true Boltzmann distribution defined by the semi-empirical energy. However, the reweighting
 223 technique allows for the debiasing of the samples. The model is trained in the same way as described
 224 in [22]. Overall, the likelihoods and ESS values observed for the TBG + full model are superior to
 225 those reported in [22] (Table 1). This is achieved with nearly the same amount of parameters and
 226 maintaining comparable training and inference times (see Appendix B.3). Furthermore, the correct
 227 free energy difference is recovered, as demonstrated in Appendix A.1.

228 In [22] the authors used a semi-empirical potential to avoid the required ordering of the atoms to the
 229 topology for classical force fields. As the prior distribution of the Boltzmann Generator is usually
 230 a multivariate standard Gaussian distribution, generated samples will almost certainly not have the
 231 correct ordering. As we have introduced an efficient way to reorder samples in Section 4.3, we can
 232 now also evaluate alanine dipeptide for a classical force field. Therefore, we retrain the model in [22]
 233 on the classical MD trajectory and compare with our TBG + full architecture. We bias the training
 234 data as before towards the unlikely φ state. As expected, the likelihood and ESS for the classical
 235 force field are much better than for the semi-empirical one, as the training data stems from the target
 236 distribution. Our proposed architecture again performs significantly better, as shown in Section 5.1
 237 and Figure 1. The majority of generated samples with the TBG + full model and the BG + backbone
 238 sample nearly exclusively correct configurations, i.e. configurations with the correct bond graph,
 239 namely nearly 100% and about 98%, respectively. As presented in Figure 1, both models recover the
 240 free energy landscape correctly.

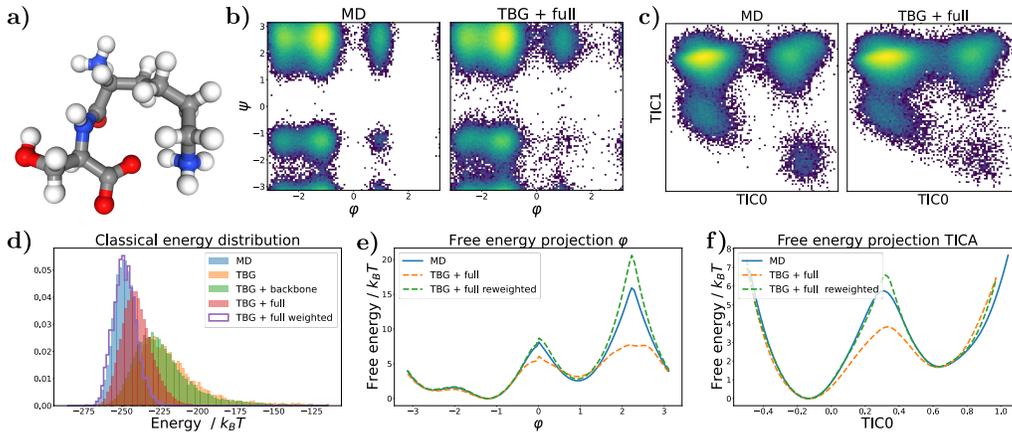


Figure 2: Results for the KS dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of samples generated with different methods and architectures. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

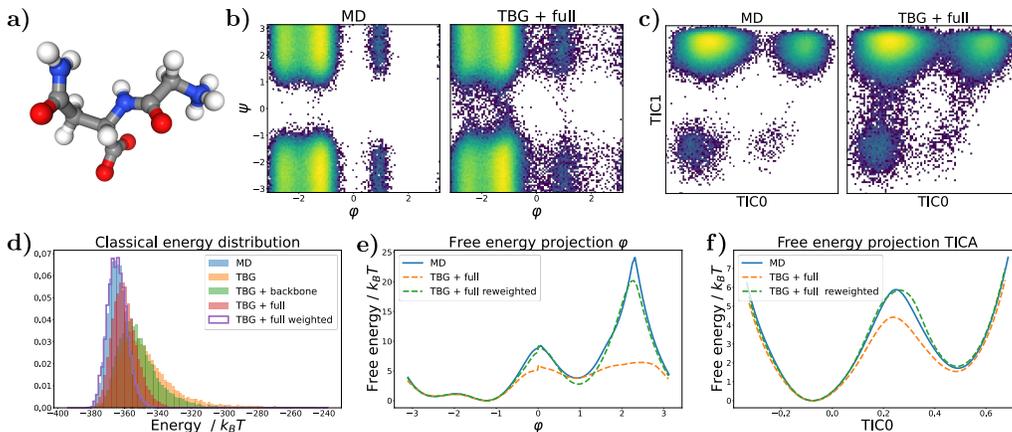


Figure 3: Results for the GN dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of samples generated with different methods and architectures. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

241 5.2 Dipeptides (2AA)

242 In our second experiment, we evaluate our model on dipeptides and show transferability. The dataset
 243 was introduced in [10]. The training set consists of 200 dipeptides, which were simulated each with a
 244 classical force field for 50 ns and, therefore, may not have reached convergence. Nevertheless, as
 245 previously demonstrated, it is not necessary to train on unbiased data in order to obtain unbiased
 246 samples with a Boltzmann Generator.

247 We compare the three different transferable architectures described in Section 4.1 and use the same
 248 training procedure for all of them. Similar to the alanine dipeptide experiments, we obtain significantly
 249 better results for the TBG + full model in terms of ESS (Table 2 and Figure 4a), energies (Figure 2d),
 250 the ratio of correct configurations (Table 2), and likelihoods of test set samples (Appendix A.5).
 251 In particular, the extremely low number of correct configurations for numerous test peptides for the

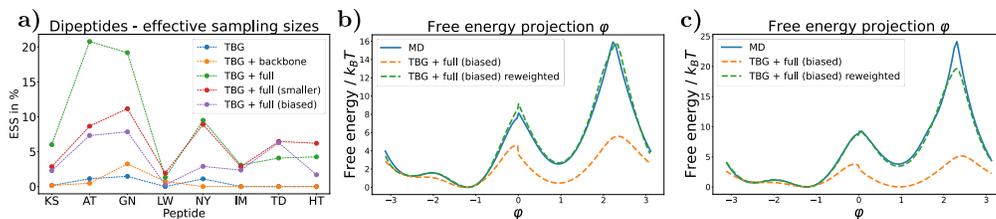


Figure 4: (a) Effective samples sizes (ESS) for the first 8 test peptides for different transferable architectures and training sets. (b) Free energy projection along the φ angle for the TBG + full model trained on the *biased* dataset for the KS dipeptide. The weighted free energy projection demonstrates a superior fit compared to the TBG + full model (see Figure 2e). (c) Free energy projection along the φ angle for the TBG + full model trained on the *biased* dataset for the GN dipeptide. The weighted free energy projection demonstrates a superior fit compared to the TBG + full model (see Figure 3e).

252 TBG and TBG + backbone models renders them unsuitable as Boltzmann Generators for this setting
 253 (Table 2 and Appendix A.5). Furthermore, the TBG + full model always find all metastable states for
 254 unseen test peptides (see also Appendix A.5).

255 The results for the well-performing TBG + full model are presented for two exemplary peptides from
 256 the test set in Figure 2 and Figure 3. They were chosen as all architectures sample relevant amounts of
 257 correct configurations. Detailed results for other evaluated test peptides are shown in Appendix A.5.

258 The TBG + full model is an exemplary Boltzmann Emulator, as it is capable of capturing all
 259 metastable states of the target Boltzmann distribution (Figure 2b,c and Figure 3b,c). However, it is
 260 furthermore also a capable Boltzmann Generator, as it allows for efficient reweighting (Figure 2d,e,f
 261 and Figure 3d,e,f). To identify different metastable states, we employ time-lagged independent
 262 component analysis (TICA) [60], a dimensionality reduction technique that separates metastable
 263 states. We show this analysis in addition to the Ramachandran plots for the dihedral angles.

264 Moreover, we investigate the influence of the training set in two ablation studies.

265 **Training on a biased training set** Our alanine dipeptide results as well as [22] indicate that it can
 266 be advantageous to bias the training data towards states that are less probable, such as positive φ
 267 states, to recover free energy landscapes. Therefore, we bias the training data by weighting positive
 268 φ states for each training peptide, such that they have nearly equal weight to the negative states (see
 269 also Appendix B.4). We show that a TBG + full model trained on this dataset (TBG + full (biased))
 270 produces even more accurate free energy landscapes for both the Ramachandran and TICA projections
 271 (Figure 4bc). Notably, the unweighted projection shows a clear bias, as expected. However, as the
 272 training data is now biased, the effective sample size (ESS) is generally lower (Table 2 and Figure 4a).

273 **Training on a smaller training set** Additionally, we examine the impact of smaller training sets
 274 on the generalization results. To this end, we train the TBG + full model on two smaller datasets with
 275 shorter simulation times: (i) 5ns for each training simulation and (ii) only 500ps of each training
 276 simulation. Consequently, the training trajectories are 10 times and 100 times smaller than before.
 277 As we utilize only the initial portion of each trajectory, a greater number of metastable states are
 278 missed during the brief simulations, as illustrated in Appendix A.2. While the training on the tenfold
 279 smaller trainings set, we refer to the model as TBG + full (smaller), shows similar results to training
 280 on the whole trainings set (Table 2 and Appendix A.5), the even smaller trainings set leads to inferior
 281 results, with several metastable states being missed as presented in Appendix A.3. Nevertheless, we
 282 demonstrated that TBGs can be trained with very small datasets, with trajectories that individually
 283 miss many metastable states.

284 6 Discussion

285 For the first time, we demonstrated the feasibility of training *transferable* Boltzmann Generators.
 286 We introduced a general framework for training and evaluating transferable Boltzmann Generators
 287 based on continuous normalizing flows. Furthermore, we developed a transferable architecture
 288 based on equivariant graph neural networks and demonstrated the importance of including topology

Table 2: Effective samples size and correct configuration rate for unseen dipeptides across different transferable Boltzmann Generator (TBG) architectures.

Model	ESS (\uparrow)		Correct configurations (\uparrow)	
	Mean	Range	Mean	Range
TBG	$0.48 \pm 0.59\%$	(0.0%, 1.47%)	$13 \pm 18\%$	(1%, 48%)
TBG + backbone	$0.58 \pm 1.04\%$	(0.0%, 3.24%)	$17 \pm 21\%$	(1%, 52%)
TBG + full	$8.53 \pm 6.99\%$	(1.31%, 20.79%)	$95 \pm 4\%$	(86%, 100%)
TBG + full (smaller)	$6.13 \pm 3.13\%$	(1.93%, 11.16%)	$96 \pm 3\%$	(88%, 100%)
TBG + full (biased)	$3.86 \pm 2.67\%$	(0.24%, 7.84%)	$96 \pm 4\%$	(87%, 100%)

289 information in the architecture to enable efficient generalization to unseen, but similar systems. The
 290 transferable Boltzmann Generator was evaluated on dipeptides, where significant effective sample
 291 sizes were demonstrated on unseen test peptides and accurate sampling of physical properties, such
 292 as the free energy difference between metastable states, was achieved. Moreover, we have shown
 293 in ablation studies that transferable Boltzmann Generators can be extremely data efficient, with
 294 even small training trajectories being sufficient. Future research will determine whether and how
 295 transferable Boltzmann Generators can be scaled to larger systems.

296 7 Limitations / Future work

297 We leave the scaling to larger system for future work. Notably, this usually requires large amounts of
 298 computational resources, as e.g. shown in [10], where they are able to train their transferable model
 299 on tetrapeptides, but use more than 100 times more parameters than us.

300 Instead of flow matching, one could use optimal transport flow matching [21] or equivariant optimal
 301 transport flow matching [22] for training, but as indicated in [22] the effect for molecular systems,
 302 especially in the presence of many distinguishable particles, are small.

303 Throughout our work, we utilize a standard Gaussian prior distribution. However, as recently
 304 introduced, an alternative is to use a Harmonic prior distribution [61, 62], where atoms that are
 305 close in the bond graph are sampled in the vicinity of each other. Notably, we experimented with
 306 this different prior distribution, but did not find relevant improvements for our transferable model.
 307 This finding is in alignment with the results of [63] that chemical informed prior distributions do
 308 not enhance performance significantly compared to simpler uninformed prior distributions for flow
 309 matching for molecules. Instead, the network architecture and inductive bias are more important.

310 Despite conducting a series of ablation studies, we did not pursue the impact of a training set
 311 comprising a smaller number of peptides. Instead, we opted for a shorter trajectory approach.
 312 Furthermore, we could consider relaxing the 2AA dataset with the semi-empirical force field and
 313 training on this modified version, analogous to the alanine dipeptide experiment.

314 The EGNN architecture was employed for the vector field, as it permits fast evaluation. However, a
 315 promising avenue for future research is to explore alternative architectures for the vector field, such as
 316 [54, 55, 56, 57, 64, 65, 62], to ascertain whether this enhances performance, which may be necessary
 317 to enable scaling to larger systems than those considered. We hope that our provided framework will
 318 enable the scaling of transferable Boltzmann Generators to larger systems in future research.

319 8 Broader Impact

320 This work represents foundational research with no immediate societal impact. However, if our
 321 method is scalable to larger, more relevant systems, it could facilitate the acceleration of drug and
 322 material discovery by replacing or enhancing MD simulations, which often play a crucial part in the
 323 process. A potential risk is that it could then be used to identify new diseases or biological weapons.
 324 Another potential risk associated with this method is that, at present, no convergence criterion is
 325 known. This implies that it is not possible to be certain that all potential configurations have been
 326 identified, even if an infinite number of samples are taken. This could result in false claims regarding
 327 the results, which could have an impact on subsequent applications.

328 **References**

- 329 [1] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf
330 Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure
331 prediction of biomolecular interactions with alphafold 3. *Nature*, pages 1–3, 2024.
- 332 [2] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ron-
333 neberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al.
334 Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- 335 [3] Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Allan dos
336 Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Sal Candido, et al. Language models of
337 protein sequences at the scale of evolution enable accurate structure prediction. *BioRxiv*,
338 2022:500902, 2022.
- 339 [4] Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant
340 diffusion for molecule generation in 3D. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song,
341 Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International
342 Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*,
343 pages 8867–8887. PMLR, 17–23 Jul 2022.
- 344 [5] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A
345 geometric diffusion model for molecular conformation generation. In *International Conference
346 on Learning Representations*, 2022.
- 347 [6] Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Tor-
348 sional diffusion for molecular conformer generation. In S. Koyejo, S. Mohamed, A. Agarwal,
349 D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*,
350 volume 35, pages 24240–24253. Curran Associates, Inc., 2022.
- 351 [7] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators — sampling
352 equilibrium states of many-body systems with deep learning. *Science*, 365:eaaw1147, 2019.
- 353 [8] Alessandro Coretti, Sebastian Falkner, Jan Weinreich, Christoph Dellago, and O Anatole von
354 Lilienfeld. Boltzmann generators and the new frontier of computational sampling in many-body
355 systems. *arXiv preprint arXiv:2404.16566*, 2024.
- 356 [9] Grant M Rotskoff. Sampling thermodynamic ensembles of molecular systems with generative
357 neural networks: Will integrating physics-based models close the generalization gap? *Current
358 Opinion in Solid State and Materials Science*, 30:101158, 2024.
- 359 [10] Leon Klein, Andrew Y. K. Foong, Tor Erlend Fjelde, Bruno Kacper Mlodozieniec, Marc
360 Brockschmidt, Sebastian Nowozin, Frank Noe, and Ryota Tomioka. Timewarp: Transferable
361 acceleration of molecular dynamics by learning time-coarsened dynamics. In *Thirty-seventh
362 Conference on Neural Information Processing Systems*, 2023.
- 363 [11] Esteban G Tabak, Eric Vanden-Eijnden, et al. Density estimation by dual ascent of the log-
364 likelihood. *Communications in Mathematical Sciences*, 8(1):217–233, 2010.
- 365 [12] Esteban G Tabak and Cristina V Turner. A family of nonparametric density estimation algo-
366 rithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- 367 [13] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji
368 Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *CoRR*, 2019.
- 369 [14] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In
370 *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- 371 [15] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
372 differential equations. In *Advances in neural information processing systems*, pages 6571–6583,
373 2018.
- 374 [16] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, and David Duvenaud. Scalable re-
375 versible generative models with free-form continuous dynamics. In *International Conference
376 on Learning Representations*, 2019.

- 377 [17] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components
378 estimation. *CoRR*, 2014.
- 379 [18] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
380 matching for generative modeling. In *The Eleventh International Conference on Learning*
381 *Representations*, 2023.
- 382 [19] Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic
383 interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.
- 384 [20] Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate
385 and transfer data with rectified flow. In *The Eleventh International Conference on Learning*
386 *Representations*, 2023.
- 387 [21] Alexander Tong, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks,
388 Kilian Fatras, Guy Wolf, and Yoshua Bengio. Conditional flow matching: Simulation-free
389 dynamic optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- 390 [22] Leon Klein, Andreas Krämer, and Frank Noe. Equivariant flow matching. In *Thirty-seventh*
391 *Conference on Neural Information Processing Systems*, 2023.
- 392 [23] Laurence Illing Midgley, Vincent Stimper, Javier Antoran, Emile Mathieu, Bernhard Schölkopf,
393 and José Miguel Hernández-Lobato. SE(3) equivariant augmented coupling flows. In *Thirty-*
394 *seventh Conference on Neural Information Processing Systems*, 2023.
- 395 [24] Peter Wirnsberger, Andrew J Ballard, George Papamakarios, Stuart Abercrombie, Sébastien
396 Racanière, Alexander Pritzel, Danilo Jimenez Rezende, and Charles Blundell. Targeted free
397 energy estimation via learned mappings. *The Journal of Chemical Physics*, 153(14):144112,
398 2020.
- 399 [25] Manuel Dibak, Leon Klein, Andreas Krämer, and Frank Noé. Temperature steerable flows and
400 Boltzmann generators. *Phys. Rev. Res.*, 4:L042005, Oct 2022.
- 401 [26] Jonas Köhler, Andreas Krämer, and Frank Noé. Smooth normalizing flows. In M. Ranzato,
402 A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural*
403 *Information Processing Systems*, volume 34, pages 2796–2809. Curran Associates, Inc., 2021.
- 404 [27] Laurence Illing Midgley, Vincent Stimper, Gregor N. C. Simm, Bernhard Schölkopf, and
405 José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. In *The Eleventh*
406 *International Conference on Learning Representations*, 2023.
- 407 [28] Xinqiang Ding and Bin Zhang. Deepbar: A fast and exact method for binding free energy
408 computation. *Journal of Physical Chemistry Letters*, 12:2509–2515, 3 2021.
- 409 [29] Joseph C. Kim, David Bloore, Karan Kapoor, Jun Feng, Ming-Hong Hao, and Mengdi Wang.
410 Scalable normalizing flows enable boltzmann generators for macromolecules. In *International*
411 *Conference on Learning Representations (ICLR)*, 2024.
- 412 [30] Rasool Ahmad and Wei Cai. Free energy calculation of crystalline solids using normaliz-
413 ing flows. *Modelling and Simulation in Materials Science and Engineering*, 30(6):065007,
414 September 2022.
- 415 [31] Kim A Nicoli, Christopher J Anders, Lena Funcke, Tobias Hartung, Karl Jansen, Pan Kessel,
416 Shinichi Nakajima, and Paolo Stornati. Estimation of thermodynamic observables in lattice
417 field theories with deep generative models. *Physical review letters*, 126(3):032001, 2021.
- 418 [32] Hao Wu, Jonas Köhler, and Frank Noe. Stochastic normalizing flows. In H. Larochelle,
419 M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information*
420 *Processing Systems*, volume 33, pages 5933–5944. Curran Associates, Inc., 2020.
- 421 [33] Xinqiang Ding and Bin Zhang. Computing absolute free energy with deep generative models.
422 *Biophysical Journal*, 120(3):195a, 2021.
- 423 [34] Andrea Rizzi, Paolo Carloni, and Michele Parrinello. Multimap targeted free energy estimation,
424 2023.

- 425 [35] Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative
426 learning for symmetric densities. In *International conference on machine learning*, pages
427 5361–5370. PMLR, 2020.
- 428 [36] Danilo Jimenez Rezende, Sébastien Racanière, Irina Higgins, and Peter Toth. Equivariant
429 Hamiltonian flows. *arXiv preprint arXiv:1909.13739*, 2019.
- 430 [37] Victor Garcia Satorras, Emiel Hoogetboom, Fabian Fuchs, Ingmar Posner, and Max Welling.
431 E(n) equivariant normalizing flows. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and
432 J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34,
433 pages 4181–4192. Curran Associates, Inc., 2021.
- 434 [38] Jonas Köhler, Michele Invernizzi, Pim de Haan, and Frank Noé. Rigid body flows for sampling
435 molecular crystal structures. In *International Conference on Machine Learning, ICML 2023*,
436 volume 202 of *Proceedings of Machine Learning Research*, pages 17301–17326. PMLR, 2023.
- 437 [39] Andrea Rizzi, Paolo Carloni, and Michele Parrinello. Targeted free energy perturbation revisited:
438 Accurate free energies from mapped reference potentials. *Journal of Physical Chemistry Letters*,
439 12:9449–9454, 2021.
- 440 [40] Michele Invernizzi, Andreas Krämer, Cecilia Clementi, and Frank Noé. Skipping the replica
441 exchange ladder with normalizing flows. *The Journal of Physical Chemistry Letters*, 13:11643–
442 11649, 2022.
- 443 [41] Osama Abidin and Philip M Kim. Pepflow: direct conformational sampling from peptide energy
444 landscapes through hypernetwork-conditioned diffusion. *bioRxiv*, pages 2023–06, 2023.
- 445 [42] Bowen Jing, Bonnie Berger, and Tommi Jaakkola. Alphafold meets flow matching for generating
446 protein ensembles. *arXiv preprint arXiv:2402.04845*, 2024.
- 447 [43] Shuxin Zheng, Jiyan He, Chang Liu, Yu Shi, Ziheng Lu, Weitao Feng, Fusong Ju, Jiayi Wang,
448 Jianwei Zhu, Yaosen Min, et al. Predicting equilibrium distributions for molecular systems with
449 deep learning. *Nature Machine Intelligence*, pages 1–10, 2024.
- 450 [44] Juan Viguera Diez, Sara Romeo Atance, Ola Engkvist, and Simon Olsson. Generation of
451 conformational ensembles of small molecules via surrogate model-assisted molecular dynamics.
452 *Machine Learning: Science and Technology*, 5(2):025010, 2024.
- 453 [45] Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: Multiple
454 time-resolution models for molecular dynamics. In *Thirty-seventh Conference on Neural
455 Information Processing Systems*, 2023.
- 456 [46] Nicholas E Charron, Felix Musil, Andrea Guljas, Yaoyi Chen, Klara Bonneau, Aldo S Pados-
457 Trejo, Jacopo Venturin, Daria Gusew, Iryna Zaporozhets, Andreas Krämer, et al. Navigating
458 protein landscapes with a machine-learned transferable coarse-grained model. *arXiv preprint
459 arXiv:2310.18278*, 2023.
- 460 [47] Jonas Köhler, Yaoyi Chen, Andreas Krämer, Cecilia Clementi, and Frank Noé. Flow-matching:
461 Efficient coarse-graining of molecular dynamics without forces. *Journal of Chemical Theory
462 and Computation*, 19(3):942–952, 2023.
- 463 [48] Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou,
464 and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule
465 generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- 466 [49] Yan Wang, Lihao Wang, Yuning Shen, Yiqun Wang, Huizhuo Yuan, Yue Wu, and Quanquan
467 Gu. Protein conformation generation via force-guided se (3) diffusion models. *arXiv preprint
468 arXiv:2403.14088*, 2024.
- 469 [50] Felix Draxler, Peter Sorrenson, Lea Zimmermann, Armand Rousselot, and Ullrich Köthe.
470 Free-form flows: Make any architecture a normalizing flow. In Sanjoy Dasgupta, Stephan
471 Mandt, and Yingzhen Li, editors, *Proceedings of The 27th International Conference on Artificial
472 Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages
473 2197–2205. PMLR, 02–04 May 2024.

- 474 [51] H Wiegand. Kish, I.: Survey sampling. John Wiley & Sons, Inc., New York, London 1965, ix+
475 643 s., 31 abb., 56 tab., Preis 83 s., 1968.
- 476 [52] George Papamakarios, Eric T Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji
477 Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *J. Mach.*
478 *Learn. Res.*, 22(57):1–64, 2021.
- 479 [53] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural
480 networks. In *International conference on machine learning*, pages 9323–9332. PMLR, 2021.
- 481 [54] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron
482 Dror. Learning from protein structure with geometric vector perceptrons. In *International*
483 *Conference on Learning Representations*, 2021.
- 484 [55] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the
485 prediction of tensorial properties and molecular spectra. In *International Conference on*
486 *Machine Learning*, pages 9377–9388. PMLR, 2021.
- 487 [56] Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d
488 atomistic graphs. In *The Eleventh International Conference on Learning Representations*, 2023.
- 489 [57] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional
490 graph neural networks for molecules. *Advances in Neural Information Processing Systems*,
491 34:6790–6802, 2021.
- 492 [58] Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative
493 learning for symmetric densities. In *International Conference on Machine Learning*, pages
494 5361–5370. PMLR, 2020.
- 495 [59] Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. Gfn2-xtb—an accurate and
496 broadly parametrized self-consistent tight-binding quantum chemical method with multipole
497 electrostatics and density-dependent dispersion contributions. *Journal of Chemical Theory and*
498 *Computation*, 15(3):1652–1671, 2019. PMID: 30741547.
- 499 [60] Guillermo Pérez-Hernández, Fabian Paul, Toni Giorgino, Gianni De Fabritiis, and Frank Noé.
500 Identification of slow molecular order parameters for Markov model construction. *The Journal*
501 *of chemical physics*, 139(1):07B604_1, 2013.
- 502 [61] Bowen Jing, Ezra Erives, Peter Pao-Huang, Gabriele Corso, Bonnie Berger, and Tommi Jaakkola.
503 Eigenfold: Generative protein structure prediction with diffusion models. *arXiv preprint*
504 *arXiv:2304.02198*, 2023.
- 505 [62] Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Harmonic self-
506 conditioned flow matching for multi-ligand docking and binding site design. *arXiv preprint*
507 *arXiv:2310.05764*, 2023.
- 508 [63] Dina A Sharon, Yining Huang, Motolani Oyewole, and Sammy Mustafa. How to go with
509 the flow: an analysis of flow matching molecular docking performance with priors of varying
510 information content. In *ICLR 2024 Workshop on Generative and Experimental Perspectives for*
511 *Biomolecular Design*, 2024.
- 512 [64] Weitao Du, Yuanqi Du, Limei Wang, Dieqiao Feng, Guifeng Wang, Shuiwang Ji, Carla Gomes,
513 and Zhi-Ming Ma. A new perspective on building efficient and expressive 3d equivariant graph
514 neural networks. *arXiv preprint arXiv:2304.04757*, 2023.
- 515 [65] Han Yang, Chenxi Hu, Yichi Zhou, Xixian Liu, Yu Shi, Jielan Li, Guanzhi Li, Zekun Chen,
516 Shuizhou Chen, Claudio Zeni, et al. Mattersim: A deep learning atomistic model across
517 elements, temperatures and pressures. *arXiv preprint arXiv:2405.04967*, 2024.
- 518 [66] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
519 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas
520 Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,
521 Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style,
522 high-performance deep learning library. In *Advances in Neural Information Processing Systems*
523 32, pages 8024–8035. Curran Associates, Inc., 2019.

- 524 [67] Michael Poli, Stefano Massaroli, Atsushi Yamashita, Hajime Asama, Jinkyoo Park, and Stefano
525 Ermon. Torchdyn: Implicit models and neural numerical methods in pytorch. In *Neural*
526 *Information Processing Systems, Workshop on Physical Reasoning and Inductive Biases for the*
527 *Real World*, volume 2, 2021.
- 528 [68] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and
529 function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos,
530 NM (United States), 2008.
- 531 [69] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
532 *arXiv:1412.6980*, 2014.

Table 3: Dimensionless free energy differences for the slowest transition of alanine dipeptide estimated with various methods. Umbrella sampling yields a converged reference solution. Errors are calculated over five runs. Values for BG + backbone and Umbrella sampling are taken from [22].

	Umbrella sampling	BG + backbone [22]	TBG + full (ours)
Free energy difference / $k_B T$	4.10 ± 0.26	4.10 ± 0.08	4.09 ± 0.05

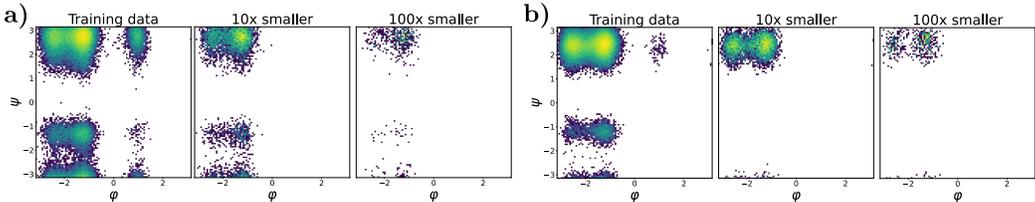


Figure 5: Example Ramachandran plots for different trajectory lengths for the training data. It can be observed that as the trajectory length decreases, the number of metastable states that are missed increases, thereby making the learning task more challenging. (a) AY dipeptide (b) IH dipeptide.

533 Appendix

534 A Additional results and experiments

535 A.1 Semi-empirical force field for alanine dipeptide

536 We report the free energy differences for the slowest transitions of alanine dipeptide for a semi-
537 empirical force field in Table 3. See Section 5 for more details.

538 A.2 Dipeptide training data

539 When training on smaller training sets, i.e. with shorter trajectories, additional metastable states will
540 not be visited during the short simulation times. We show this for two example training peptides in
541 Figure 5. Nevertheless, the TBG + full (smaller) model trained on 10 times shorter trajectories, is
542 nearly as good as the model trained on the full trajectories, see Section 5. However, for the 100 times
543 smaller trajectories, the TBG + full model perform significantly worse, see Appendix A.3.

544 A.3 Smaller dataset

545 We investigate the effect of 100 times smaller trainings trajectories, i.e. simulation time of only
546 500ps. As shown in Appendix A.2, these trajectories miss many metastable states. This can be also
547 observed for the so trained models, which we refer as TBG + full (smaller500), as they do not capture
548 especially unlikely metastable states well as presented in Figure 6 and Figure 7. In contrast, models
549 trained on larger trajectories find all metastable states and allow for efficient reweighting, as discussed
550 in Section 5 and Appendix A.5.

551 A.4 Sampled dipeptide configurations

552 For some amino acid combinations, both the TBG and TBG + backbone models sample only a
553 small number of correct configurations. Although the generated configurations are potentially valid
554 molecular configurations, they are not the one of the target dipeptide as shown in Figure 8. Only the
555 various TBG + full architectures samples nearly exclusively correct configurations.

556 A.5 Additional results for dipeptides

557 Inference is a costly process, and extensive sampling is necessary to obtain reliable estimates for the
558 expected sample size (ESS). Therefore, we only evaluate the transferable models on a subset of the
559 test set. The dipeptides are randomly selected, but it is ensured that all amino acids are represented at
560 least once. However, we evaluate the best-performing model, namely TBG + full, for twice as many

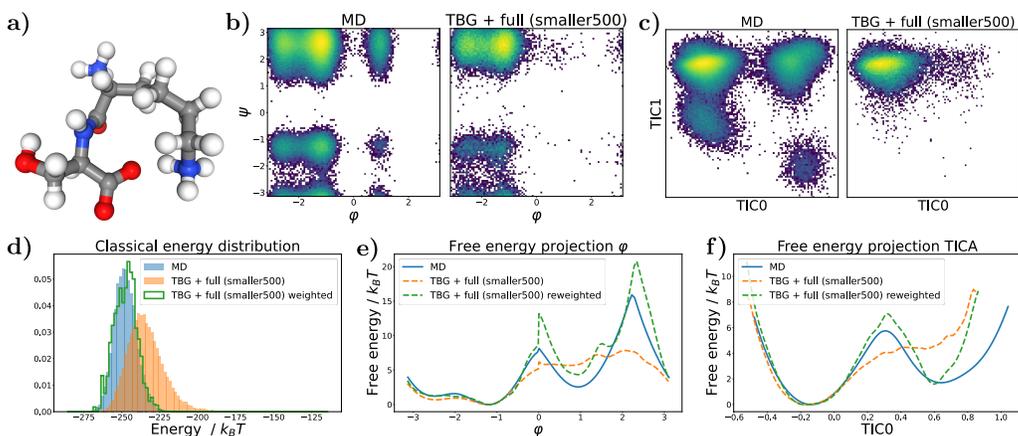


Figure 6: Results for the KS dipeptide for TBG + full model trained on 100 times smaller training trajectories. As can be seen in Figure 2, the results for the TBG + full model trained on the whole trajectories are much better. (a) KS dipeptide (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the model (right). (d) Energies of samples generated with the model. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

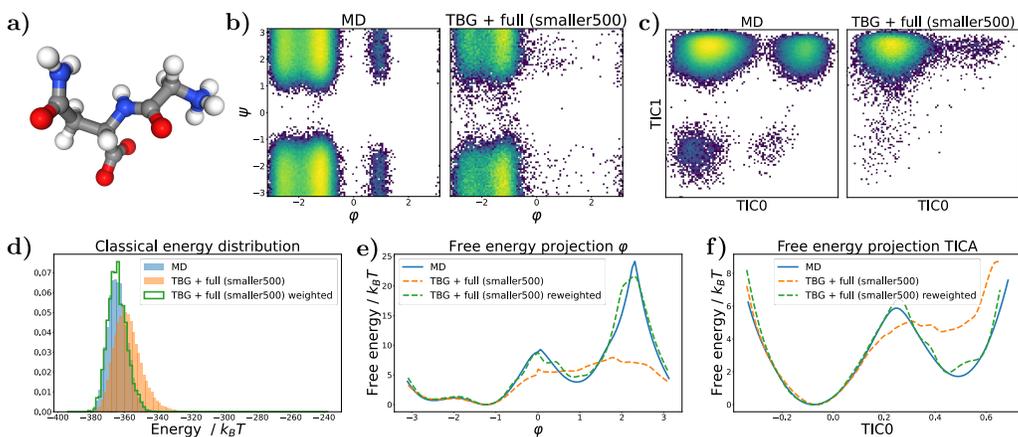


Figure 7: Results for the GN dipeptide for TBG + full model trained on 100 times smaller training trajectories. As can be seen in Figure 3, the results for the TBG + full model trained on the whole trajectories are much better. (a) GN dipeptide (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the model (right). (d) Energies of samples generated with the model. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

561 test peptides as the rest. The results for the additional test peptides are in good agreement with the
 562 first half as presented in Table 4 and Figure 9.

563 We report individual results for the different architectures in Figure 9.

564 To illustrate the performance of the TBG + full model, we present additional examples of dipeptides
 565 from the test set in Figure 10a-f and Figure 11a-f. Furthermore, we also again show results for
 566 training on the biased dataset in the same figures (Figure 10g,h,i and Figure 11g,h,i). As observed
 567 previously, the TBG + full (biased) model recovers the free energy landscape better than the TBG +
 568 full model, especially for the φ projections. We present additional examples of dipeptides from the
 569 test set for the TBG + full model in Figure 12, Figure 13, Figure 14, Figure 15, and Figure 16.

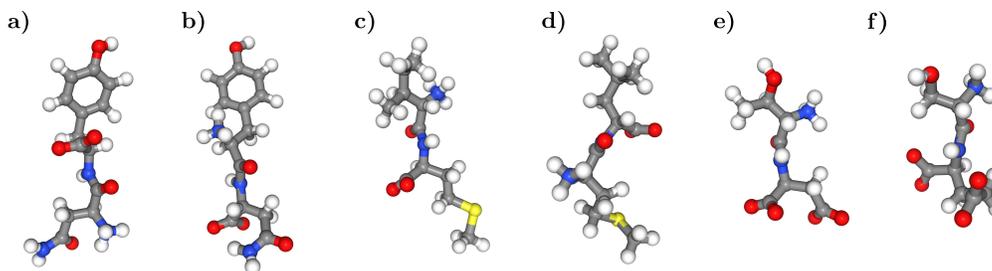


Figure 8: Sampled molecules with the TBG and TBG + backbone models, which do not have the correct topology. (a) NY dipeptide reference (b) Generated molecule with NY atoms by the TBG model. (c) IM dipeptide reference (d) Generated molecule with IM atoms by the TBG model. (e) TD dipeptide reference (f) Generated molecule with TD atoms by the TBG + backbone model.

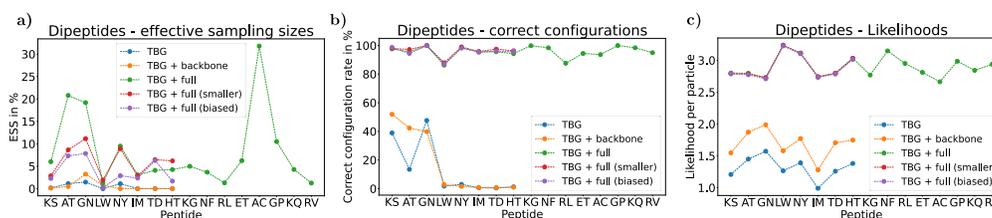


Figure 9: Performance comparison for different transferable architectures and training sets on the test set (a) Effective samples sizes (ESS) (b) Correct configuration rate (c) Likelihood per particle.

570 Furthermore, we present results for two example peptides from the test set for the TBG + full (smaller)
 571 model, which is trained on trajectories that are tenfold smaller than those used for the TBG + full
 572 model. These results are shown in Figure 17 and Figure 18.

573 A.6 Transferable Boltzmann Generators as Boltzmann Emulators

574 Given the high cost of sampling with CNFs, which necessitates integrating the Jacobian trace along
 575 the positions, we did not evaluate all available test peptides (see Appendix B.2). However, since
 576 sampling without the Jacobian trace is less expensive and we do not require as many samples as for
 577 estimating the ESS, we also employ the TBG + full (smaller) model as a Boltzmann Emulator to
 578 ascertain whether we have identified all metastable states, despite the fact that it was only trained
 579 on the 10 times smaller training set. The Boltzmann Emulator is evaluated on a diverse set of
 580 test peptides, and nearly always finds all metastable states within less than one hour of wall clock
 581 time. This is a notable improvement over MD simulations, which often take longer to explore due
 582 to the iterative nature of MD. Some examples are shown in Figure 19. This experiment shares
 583 similarities with the exploration mode of [10], where they employ their model without the acceptance
 584 step and therefore also explore a potentially biased distribution rather than the unbiased Boltzmann
 585 distribution.

Table 4: Effective samples size and correct configuration rate for unseen dipeptides for the TBG + full architecture for different number of test peptides.

Model	ESS (\uparrow)		Correct configurations (\uparrow)	
	Mean	Range	Mean	Range
TBG + full (8 test peptides)	$8.53 \pm 6.99\%$	(1.31%, 20.79%)	$95 \pm 4\%$	(86%, 100%)
TBG + full (16 test peptides)	$8.27 \pm 8.29\%$	(1.26%, 31.80%)	$96 \pm 4\%$	(86%, 100%)

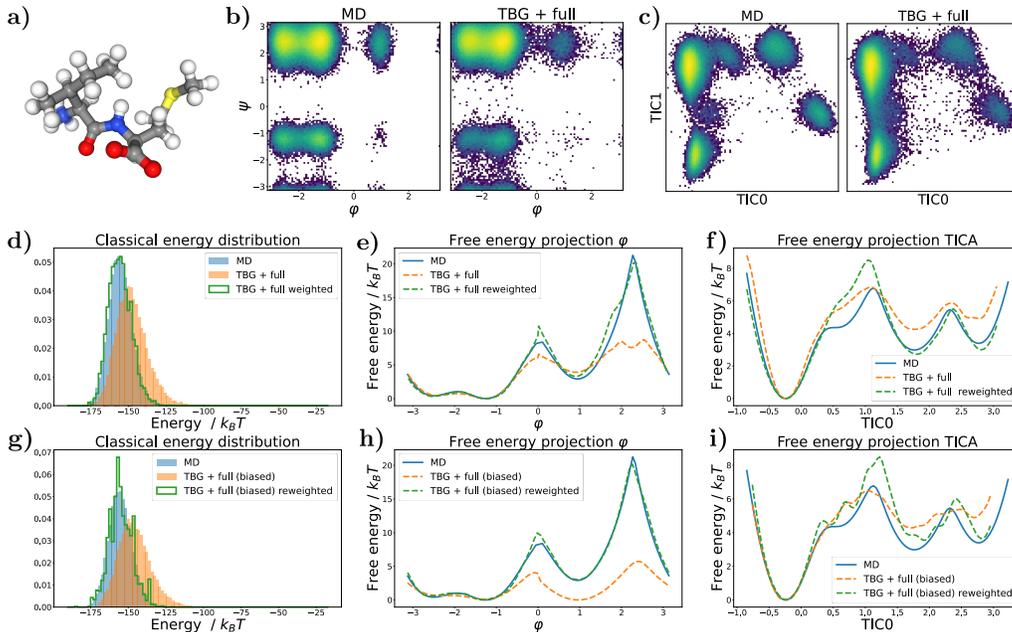


Figure 10: Results for the IM dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of samples generated with the TBG + full model. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0). (g) Energies of samples generated with the TBG + full (biased) model. (h) Free energy projection along the φ angle for the TBG + full (biased) model. (i) Free energy projection along the slowest transition (TIC0) for the TBG + full (biased) model.

586 B Technical details

587 B.1 Code libraries

588 We primarily use the following code libraries: *PyTorch* (BSD-3) [66], *bgflow* (MIT license) [7, 35],
 589 *torchdyn* (Apache License 2.0) [67], and *NetworkX* (BSD-3) [68] for validating graph isomorphisms.
 590 Additionally, we use the code from [37] (MIT license) for EGNNs, as well as the code from [10]
 591 (MIT license) and [22] (MIT license) for datasets and related evaluation methods.

592 Our code is available here: [https://osf.io/n8vz3/?view_only=](https://osf.io/n8vz3/?view_only=1052300a21bd43c08f700016728aa96e)
 593 [1052300a21bd43c08f700016728aa96e](https://osf.io/n8vz3/?view_only=1052300a21bd43c08f700016728aa96e). We will make our code public upon publication.

594 B.2 Benchmark systems

595 The investigated benchmark systems were created in prior studies [22, 10].

596 **Alanine dipeptide** The alanine dipeptide datasets were created in [22] (CC BY 4.0), we refer
 597 to them for detailed simulation details. The classical trajectory was created at $T = 300\text{K}$ with
 598 the classical *Amber ff99SBildn* force-field. The subsequent relaxation was performed with the
 599 semi-empirical *GFN2-xTB* force-field [59].

600 **Dipeptides (2AA dataset)** The original dipeptide dataset as introduced in [10] (MIT License) is
 601 available here: <https://huggingface.co/datasets/microsoft/timewarp>. As this includes
 602 a lot of intermediate saved states and quantities, like energies, we create a smaller version with is
 603 available here: https://osf.io/n8vz3/?view_only=1052300a21bd43c08f700016728aa96e.
 604 For a comprehensive overview of the simulation details, refer to [10]. All dipeptides were simulated

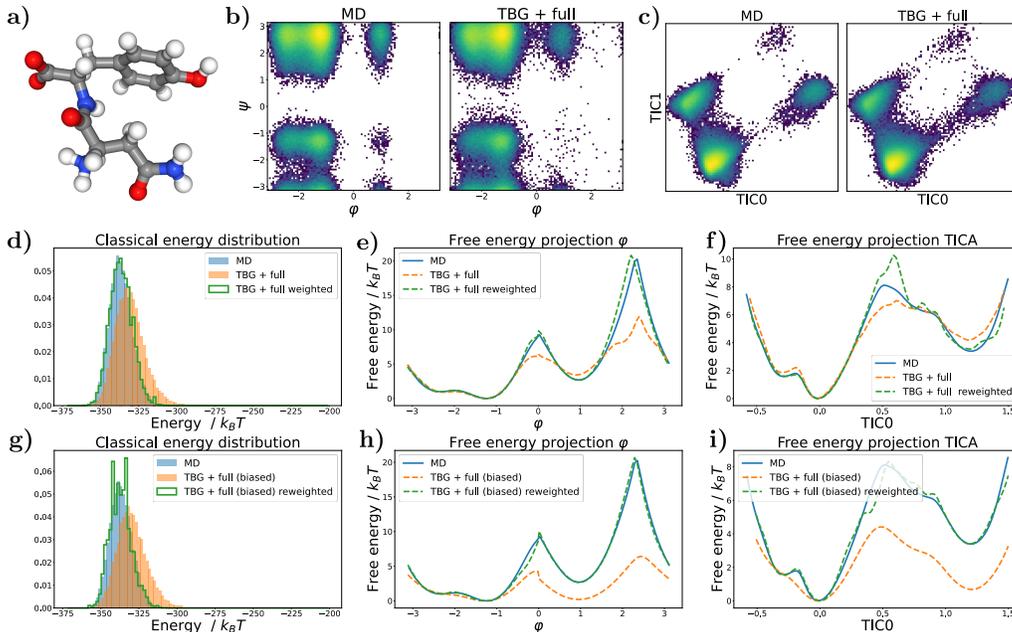


Figure 11: Results for the NY dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of samples generated with the TBG + full model. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0). (g) Energies of samples generated with the TBG + full (biased) model. (h) Free energy projection along the φ angle for the TBG + full (biased) model. (i) Free energy projection along the slowest transition (TIC0) for the TBG + full (biased) model.

605 with a classical *amber-14* force-field at $T = 310\text{K}$. The simulation of the training peptides were run
 606 for 50ns, while the test set peptides were run for $1\mu\text{s}$.

607 **Choice of test set peptides** Inference is a costly process with CNFs (see Appendix B.7), and
 608 extensive sampling is necessary to obtain reliable estimates for the relative effective sample size
 609 (ESS). Therefore, we only evaluate the transferable models on a subset of the test set. The dipeptides
 610 are randomly selected, but it is ensured that all amino acids are represented at least once.

611 B.3 Hyperparameters

612 We report the model hyperparameters for the different model architectures as describes in Section 4.1
 613 in Table 5. As in [22] all neural networks ϕ_α have one hidden layer with n_{hidden} neurons and *SiLU*
 614 activation functions. The input size of the embedding $n_{\text{embedding}}$ depends on the model architecture.

615 We report training hyperparameters for the different model architectures in Table 6. It should be
 616 noted that all TBG models are trained in an identical manner if the training set is identical. We use
 617 the ADAM optimizer for all experiments [69]. For the dipeptide training, each batch consists of three
 618 samples for each peptide.

619 B.4 Biasing target samples

620 As introduced in [22], it can be beneficial to bias the training data in such a way that unlikely states
 621 are more prominent. For alanine dipeptide and many dipeptides, the positive φ states at $\varphi = 1$ are
 622 often the unlikely ones and transition the positive and negative φ states are slow. For the
 623 alanine dipeptide dataset, the biasing methodology proposed in [22] is employed. Similarly, we bias
 624 the dipeptides based on the von Mises distribution f_{VM} . The weights ω are computed along the φ

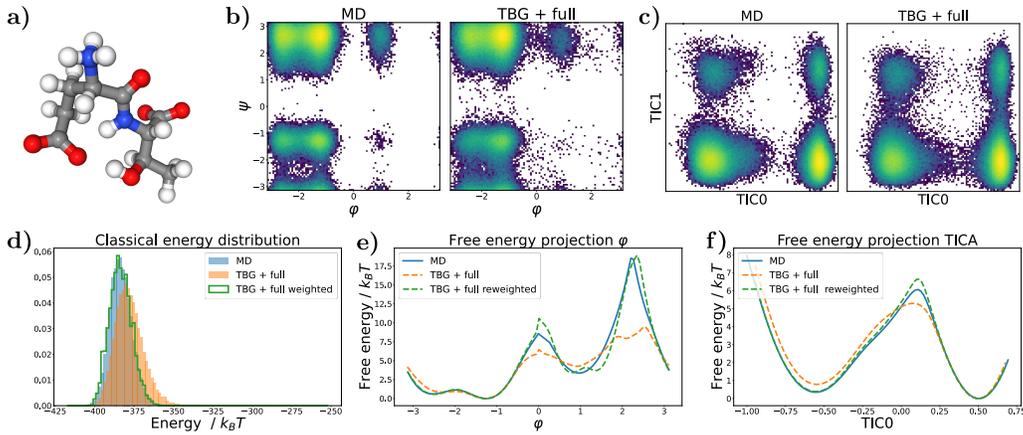


Figure 12: Results for the ET dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

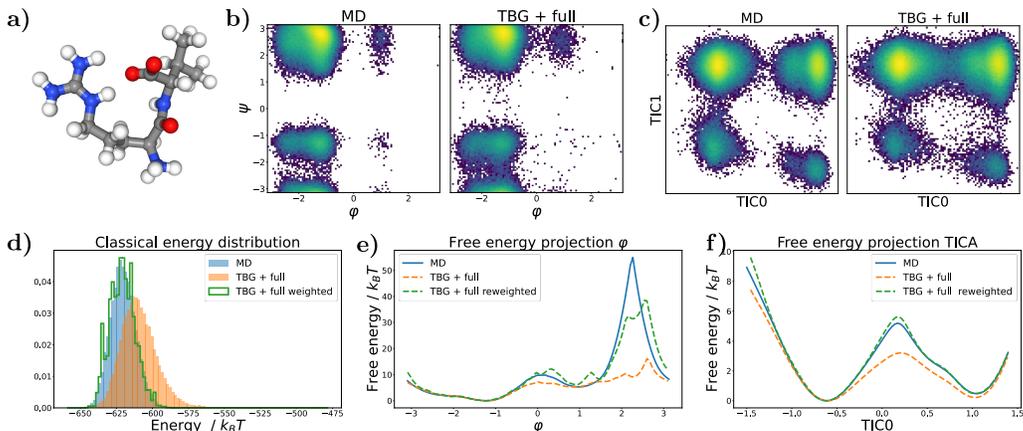


Figure 13: Results for the RV dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

625 dihedral angle as

$$\omega(\varphi) = r \cdot f_{\text{VM}}(\varphi|\mu = 1, \kappa = 10) + 1, \quad (12)$$

626 where r is computed based on the ratio of positive and negative φ states, such that both have nearly
627 the same weight after the biasing.

628 B.5 Encoding of atom types

629 The atom type embedding a_i is a one-hot vector of 54 classes. The classes are mostly defined by the
630 atom type in the topology for a classical force field. Therefore, only a few atoms are indistinguishable,
631 such as hydrogen atoms that are bound to the same carbon or nitrogen atom. Moreover, we also treat
632 oxygen atoms bound to the same carbon atom as indistinguishable, unless they are in the carboxyl
633 group. Notably, we never treat particle groups as indistinguishable, such as two CH3 groups bound to
634 the same carbon atom.

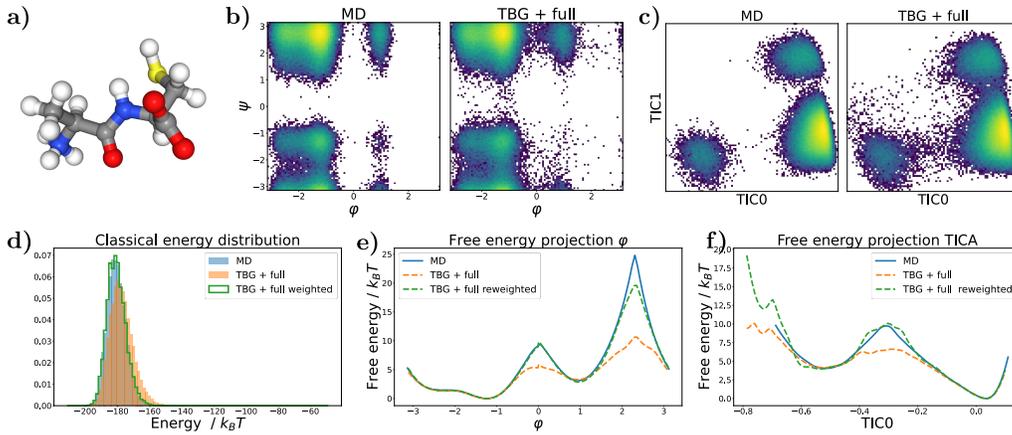


Figure 14: Results for the AC dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

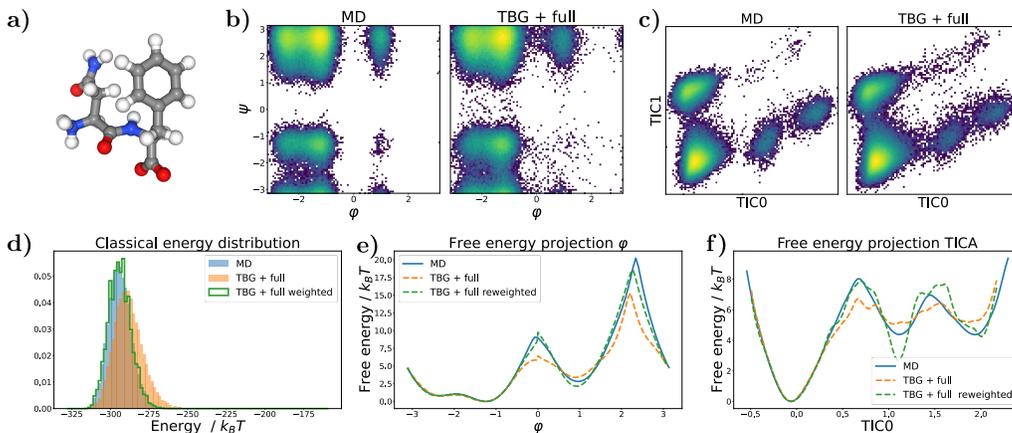


Figure 15: Results for the NF dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

635 B.6 Effective samples sizes

636 The relative effective sample sizes (ESS) are computed with Kish’s equation [51] as in prior work.
 637 For the alanine dipeptide experiments we use 2×10^5 samples for the forward ESS and 1×10^4 for the
 638 negative log likelihood computation. A total of 3×10^4 samples were used for each dipeptide in the
 639 forward ESS, while 4.5×10^3 samples were employed for the negative log likelihood computation.

640 B.7 Computing resources

641 All training and inference was performed on single *NVIDIA A100 GPUs* with 80GB of RAM.

642 The training time for the models is reported in Appendix B.3, although it should be noted that a
 643 significant amount of time was required for hyperparameter tuning. It is estimated that at least ten
 644 times the compute time reported in Appendix B.3 was necessary to identify suitable hyperparam-
 645 eters. Furthermore, inference with CNFs is expensive, especially if one requires the reweighting

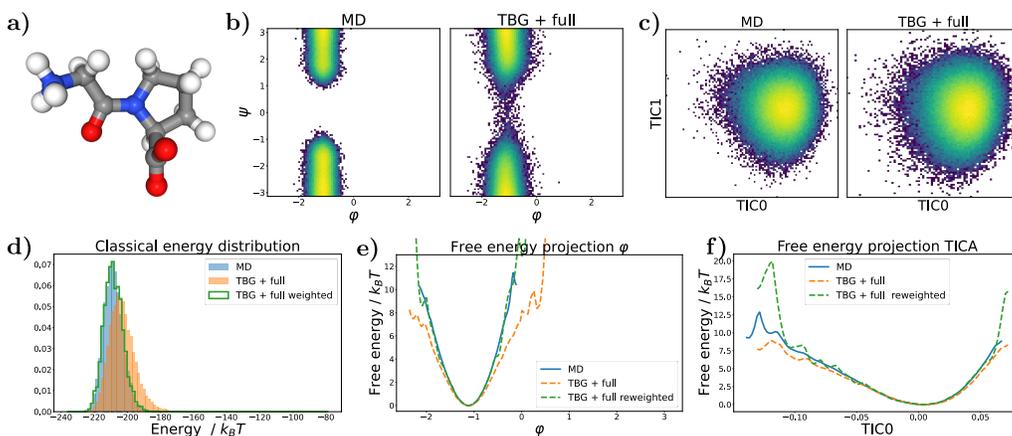


Figure 16: Results for the GP dipeptide (a) Sample generated with the TBG + full model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

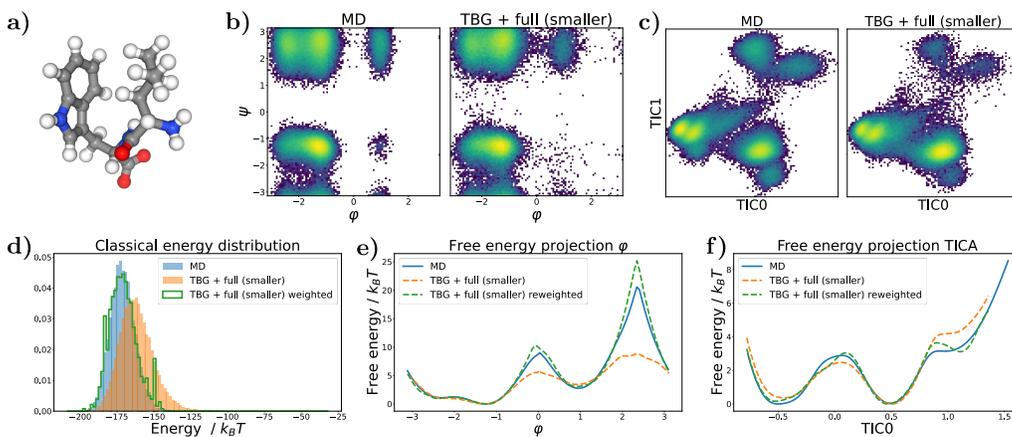


Figure 17: Results for the LW dipeptide for the TBG + full (smaller) model, which is trained on tenfold smaller trajectories than the TBG + full model. (a) Sample generated with the TBG + full (smaller) model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full (smaller) model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full (smaller) model (right). (d) Energies of generated samples. (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

646 weights. Generating 3×10^4 samples with the large transferable models for the dipeptides requires
 647 approximately four days, whereas generating 2×10^5 samples for the alanine dipeptide experiments
 648 takes less than one day. However, generating samples without corresponding weights significantly
 649 accelerates the sampling process. In the case of the dipeptides, the generation of 2×10^5 samples can
 650 be completed in less than one day. However, it should be noted that sampling can be done fully in
 651 parallel, as Boltzmann Generators generate independent samples.

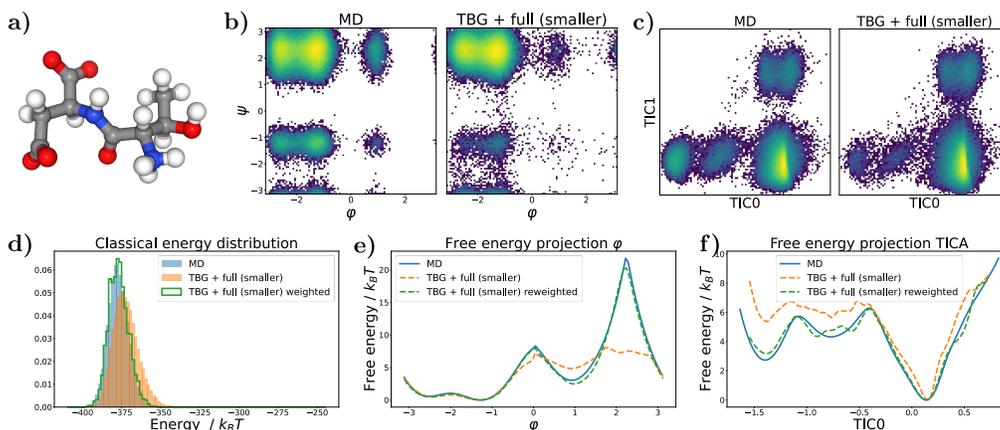


Figure 18: Results for the TD dipeptide for the TBG + full (smaller) model, which is trained on tenfold smaller trajectories than the TBG + full model. (a) Sample generated with the TBG + full (smaller) model (b) Ramachandran plot for the weighted MD distribution (left) and for samples generate with the TBG + full (smaller) model (right). (c) TICA plot for the weighted MD distribution (left) and for samples generate with the TBG + full (smaller) model (right). (d) Energies of generated samples (e) Free energy projection along the φ angle. (f) Free energy projection along the slowest transition (TIC0).

Table 5: Model hyperparameters

Model	L	n_{hidden}	$n_{\text{embedding}}$	Num. of parameters
alanine dipeptide				
BG + backbone	5	64	8	147599
TBG + full encoding	5	64	15	149147
Dipeptides (2AA)				
TBG	9	128	5	1044239
TBG + backbone	9	128	13	1046295
TBG + full encoding	9	128	76	1062486

Table 6: Training hyperparameters

Mdoel	Batch size	Learning rate	Epochs	Training time
Alanine dipeptide				
BG + backbone	256	$5e-4/5e-5$	500/500	3.5h
TBG + full	256	$5e-4/5e-5$	500/500	3.5h
Dipeptides (2AA)				
TBG	600	$5e-4/5e-5/5e-6$	4/4/4	3d
TBG + backbone	600	$5e-4/5e-5/5e-6$	4/4/4	3d
TBG + full	600	$5e-4/5e-5/5e-6$	4/4/4	3d
TBG + full (smaller)	600	$5e-4/5e-5/5e-6$	30/30/30	2.5d

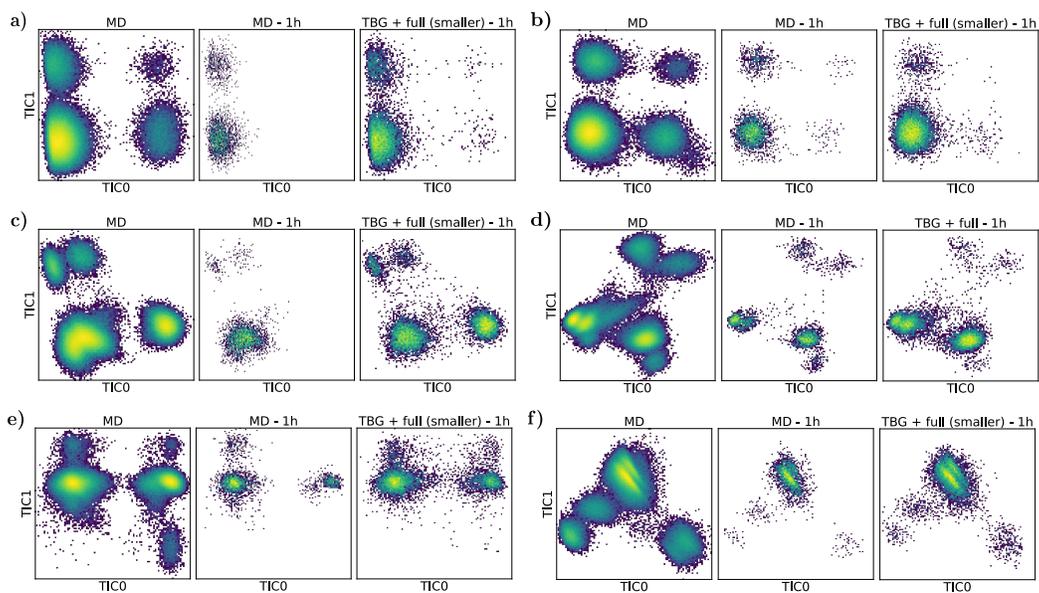


Figure 19: Comparison of classical MD runs for 1 hour (MD - 1h) and the sampling with the TBG + full (smaller) model without weight computation for 1 hour (TBG + full (smaller) - 1h). The TICA plots of different peptides from the test set are shown. It is important to note that the TICA projection is always computed with respect to the long MD trajectory (MD). All peptides stem from the test set. (a) CS dipeptide (b) EK dipeptide (c) KI dipeptide (d) LW dipeptide (e) RL dipeptide (f) TF dipeptide.

652 **NeurIPS Paper Checklist**

653 **1. Claims**

654 Question: Do the main claims made in the abstract and introduction accurately reflect the
655 paper's contributions and scope?

656 Answer: [Yes]

657 Justification: The claims made in the abstract and introduction are all based on the results of
658 our work, as shown in Section 5 and Appendix A .

659 Guidelines:

- 660 • The answer NA means that the abstract and introduction do not include the claims
661 made in the paper.
- 662 • The abstract and/or introduction should clearly state the claims made, including the
663 contributions made in the paper and important assumptions and limitations. A No or
664 NA answer to this question will not be perceived well by the reviewers.
- 665 • The claims made should match theoretical and experimental results, and reflect how
666 much the results can be expected to generalize to other settings.
- 667 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
668 are not attained by the paper.

669 **2. Limitations**

670 Question: Does the paper discuss the limitations of the work performed by the authors?

671 Answer: [Yes]

672 Justification: We discuss the limitations in Section 7.

673 Guidelines:

- 674 • The answer NA means that the paper has no limitation while the answer No means that
675 the paper has limitations, but those are not discussed in the paper.
- 676 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 677 • The paper should point out any strong assumptions and how robust the results are to
678 violations of these assumptions (e.g., independence assumptions, noiseless settings,
679 model well-specification, asymptotic approximations only holding locally). The authors
680 should reflect on how these assumptions might be violated in practice and what the
681 implications would be.
- 682 • The authors should reflect on the scope of the claims made, e.g., if the approach was
683 only tested on a few datasets or with a few runs. In general, empirical results often
684 depend on implicit assumptions, which should be articulated.
- 685 • The authors should reflect on the factors that influence the performance of the approach.
686 For example, a facial recognition algorithm may perform poorly when image resolution
687 is low or images are taken in low lighting. Or a speech-to-text system might not be
688 used reliably to provide closed captions for online lectures because it fails to handle
689 technical jargon.
- 690 • The authors should discuss the computational efficiency of the proposed algorithms
691 and how they scale with dataset size.
- 692 • If applicable, the authors should discuss possible limitations of their approach to
693 address problems of privacy and fairness.
- 694 • While the authors might fear that complete honesty about limitations might be used by
695 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
696 limitations that aren't acknowledged in the paper. The authors should use their best
697 judgment and recognize that individual actions in favor of transparency play an impor-
698 tant role in developing norms that preserve the integrity of the community. Reviewers
699 will be specifically instructed to not penalize honesty concerning limitations.

700 **3. Theory Assumptions and Proofs**

701 Question: For each theoretical result, does the paper provide the full set of assumptions and
702 a complete (and correct) proof?

703 Answer: [NA]

704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757

Justification: We only utilize theorems of prior work, which we reference accordingly.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We explain how we performed our experiments in Section 5 and Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

Answer: [Yes]

Justification: Code, data, model checkpoints as well as detailed instructions are available here: https://osf.io/n8vz3/?view_only=1052300a21bd43c08f700016728aa96e, as stated in Appendix B.1. Nevertheless, high level training, evaluation and implementation details are described in Appendix B and Section 4.1.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We discuss the dataset details in Appendix B and refer to related work for datasets introduced in prior work.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide error bars for the alanine dipeptide experiments. In contrast, for the much more expensive transferable experiments, we utilize our computational resources to sample a multitude of different peptides from the test set, rather than training and sampling distinct instances of the same architecture for the same peptide. Consequently, we obtain error bounds by averaging results over different test dipeptides rather than runs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- 810 • The factors of variability that the error bars are capturing should be clearly stated (for
811 example, train/test split, initialization, random drawing of some parameter, or overall
812 run with given experimental conditions).
- 813 • The method for calculating the error bars should be explained (closed form formula,
814 call to a library function, bootstrap, etc.)
- 815 • The assumptions made should be given (e.g., Normally distributed errors).
- 816 • It should be clear whether the error bar is the standard deviation or the standard error
817 of the mean.
- 818 • It is OK to report 1-sigma error bars, but one should state it. The authors should
819 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
820 of Normality of errors is not verified.
- 821 • For asymmetric distributions, the authors should be careful not to show in tables or
822 figures symmetric error bars that would yield results that are out of range (e.g. negative
823 error rates).
- 824 • If error bars are reported in tables or plots, The authors should explain in the text how
825 they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

827 Question: For each experiment, does the paper provide sufficient information on the com-
828 puter resources (type of compute workers, memory, time of execution) needed to reproduce
829 the experiments?

830 Answer: [Yes]

831 Justification: We discuss the required computational resources for this work for the training
832 and inference in Appendix B.3 and Appendix B.7.

833 Guidelines:

- 834 • The answer NA means that the paper does not include experiments.
- 835 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
836 or cloud provider, including relevant memory and storage.
- 837 • The paper should provide the amount of compute required for each of the individual
838 experimental runs as well as estimate the total compute.
- 839 • The paper should disclose whether the full research project required more compute
840 than the experiments reported in the paper (e.g., preliminary or failed experiments that
841 didn't make it into the paper).

9. Code Of Ethics

843 Question: Does the research conducted in the paper conform, in every respect, with the
844 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

845 Answer: [Yes]

846 Justification: The research conducted conforms with the NeurIPS Code of Ethics. All
847 authors have read the NeurIPS Code of Ethics.

848 Guidelines:

- 849 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 850 • If the authors answer No, they should explain the special circumstances that require a
851 deviation from the Code of Ethics.
- 852 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
853 eration due to laws or regulations in their jurisdiction).

10. Broader Impacts

855 Question: Does the paper discuss both potential positive societal impacts and negative
856 societal impacts of the work performed?

857 Answer: [Yes]

858 Justification: We discuss the broader impact of our work in Section 8.

859 Guidelines:

- 860 • The answer NA means that there is no societal impact of the work performed.

- 861 • If the authors answer NA or No, they should explain why their work has no societal
862 impact or why the paper does not address societal impact.
- 863 • Examples of negative societal impacts include potential malicious or unintended uses
864 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
865 (e.g., deployment of technologies that could make decisions that unfairly impact specific
866 groups), privacy considerations, and security considerations.
- 867 • The conference expects that many papers will be foundational research and not tied
868 to particular applications, let alone deployments. However, if there is a direct path to
869 any negative applications, the authors should point it out. For example, it is legitimate
870 to point out that an improvement in the quality of generative models could be used to
871 generate deepfakes for disinformation. On the other hand, it is not needed to point out
872 that a generic algorithm for optimizing neural networks could enable people to train
873 models that generate Deepfakes faster.
- 874 • The authors should consider possible harms that could arise when the technology is
875 being used as intended and functioning correctly, harms that could arise when the
876 technology is being used as intended but gives incorrect results, and harms following
877 from (intentional or unintentional) misuse of the technology.
- 878 • If there are negative societal impacts, the authors could also discuss possible mitigation
879 strategies (e.g., gated release of models, providing defenses in addition to attacks,
880 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
881 feedback over time, improving the efficiency and accessibility of ML).

882 11. Safeguards

883 Question: Does the paper describe safeguards that have been put in place for responsible
884 release of data or models that have a high risk for misuse (e.g., pretrained language models,
885 image generators, or scraped datasets)?

886 Answer: [NA]

887 Justification: Our models do not pose such a risk, as they are for molecular data.

888 Guidelines:

- 889 • The answer NA means that the paper poses no such risks.
- 890 • Released models that have a high risk for misuse or dual-use should be released with
891 necessary safeguards to allow for controlled use of the model, for example by requiring
892 that users adhere to usage guidelines or restrictions to access the model or implementing
893 safety filters.
- 894 • Datasets that have been scraped from the Internet could pose safety risks. The authors
895 should describe how they avoided releasing unsafe images.
- 896 • We recognize that providing effective safeguards is challenging, and many papers do
897 not require this, but we encourage authors to take this into account and make a best
898 faith effort.

899 12. Licenses for existing assets

900 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
901 the paper, properly credited and are the license and terms of use explicitly mentioned and
902 properly respected?

903 Answer: [Yes]

904 Justification: We give credit to used assets in this work in Appendix B.1 and Appendix B.2.
905 Our assets will be available under the MIT / CC BY 4.0 license.

906 Guidelines:

- 907 • The answer NA means that the paper does not use existing assets.
- 908 • The authors should cite the original paper that produced the code package or dataset.
- 909 • The authors should state which version of the asset is used and, if possible, include a
910 URL.
- 911 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 912 • For scraped data from a particular source (e.g., website), the copyright and terms of
913 service of that source should be provided.

- 914
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- 915
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- 916
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.
- 917
- 918
- 919
- 920
- 921

922 13. **New Assets**

923 Question: Are new assets introduced in the paper well documented and is the documentation
924 provided alongside the assets?

925 Answer: [\[Yes\]](#)

926 Justification: We document our models in Appendix B.

927 Guidelines:

- The answer NA means that the paper does not release new assets.
 - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
 - The paper should discuss whether and how consent was obtained from people whose asset is used.
 - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.
- 928
- 929
- 930
- 931
- 932
- 933
- 934
- 935

936 14. **Crowdsourcing and Research with Human Subjects**

937 Question: For crowdsourcing experiments and research with human subjects, does the paper
938 include the full text of instructions given to participants and screenshots, if applicable, as
939 well as details about compensation (if any)?

940 Answer: [\[NA\]](#)

941 Justification: No crowdsourcing nor human subjects

942 Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
 - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
 - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.
- 943
- 944
- 945
- 946
- 947
- 948
- 949
- 950

951 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human 952 Subjects**

953 Question: Does the paper describe potential risks incurred by study participants, whether
954 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
955 approvals (or an equivalent approval/review based on the requirements of your country or
956 institution) were obtained?

957 Answer: [\[NA\]](#)

958 Justification: No crowdsourcing nor human subjects

959 Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
 - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- 960
- 961
- 962
- 963
- 964

965
966
967
968
969

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.