



HAL
open science

Normalizing flows as an enhanced sampling method for atomistic supercooled liquids

Gerhard Jung, Giulio Biroli, Ludovic Berthier

► **To cite this version:**

Gerhard Jung, Giulio Biroli, Ludovic Berthier. Normalizing flows as an enhanced sampling method for atomistic supercooled liquids. *Machine Learning: Science and Technology*, 2024, 5 (3), pp.035053. 10.1088/2632-2153/ad6ca0 . hal-04699130

HAL Id: hal-04699130

<https://hal.science/hal-04699130v1>

Submitted on 16 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Normalizing flows as an enhanced sampling method for atomistic supercooled liquids

Gerhard Jung,^{1,2,*} Giulio Biroli,³ and Ludovic Berthier^{1,4}

¹*Laboratoire Charles Coulomb (L2C), Université de Montpellier, CNRS, 34095 Montpellier, France*

²*Laboratoire Interdisciplinaire de Physique (LIPhy),*

Université Grenoble Alpes, 38402 Saint-Martin-d'Hères, France

³*Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France*

⁴*Gulliver, UMR CNRS 7083, ESPCI Paris, PSL Research University, 75005 Paris, France*

(Dated: September 16, 2024)

Normalizing flows can transform a simple prior probability distribution into a more complex target distribution. Here, we evaluate the ability and efficiency of generative machine learning methods to sample the Boltzmann distribution of an atomistic model for glass-forming liquids. This is a notoriously difficult task, as it amounts to ergodically exploring the complex free energy landscape of a disordered and frustrated many-body system. We optimize a normalizing flow model to successfully transform high-temperature configurations of a dense liquid into low-temperature ones, near the glass transition. We perform a detailed comparative analysis with established enhanced sampling techniques developed in the physics literature to assess and rank the performance of normalizing flows against state-of-the-art algorithms. We demonstrate that machine learning methods are very promising, showing a large speedup over conventional molecular dynamics. Normalizing flows show performances comparable to parallel tempering and population annealing, while still falling far behind the swap Monte Carlo algorithm. Our study highlights the potential of generative machine learning models in scientific computing for complex systems, but also points to some of its current limitations and the need for further improvement.

I. INTRODUCTION

One of the most important methodological revolution in science in the last century is scientific computing [1]. Numerical simulations represent a way, complementary to experiments, to study physical systems, thus providing a unique lens on the microscopic mechanisms underpinning macroscopic physical phenomena [1–4].

A major application of numerical simulations, from the very beginning, has been sampling physical configurations at thermal equilibrium [5]. This was initially done by using either Monte Carlo Markov chains [5, 6] or molecular dynamics [7, 8]. Both methods can be viewed as ways to implement some physical dynamics to ergodically explore the configuration space, just as the physical system does. The basic challenge is to run those dynamics long enough to be able to generate a large set of uncorrelated configurations to perform accurate ensemble averages of physical observables [2].

When the system is characterized by large relaxation times, for instance near phase transitions or in disordered media, sampling can become so challenging that conventional methods may fail [2, 4]. In such cases, the only solution, so far, consists in devising alternative dynamics that ensure equilibrium sampling while being characterised by substantially smaller decorrelation times. Such strategies are described in many standard textbooks on computer simulations and statistical physics [3, 4, 9, 10]. In the context of off-lattice molecular simulations, we can

mention non-local [11–13], lifting [14], or collective [15–17] Monte Carlo algorithms, parallel tempering [18–21], population annealing [22–24], or irreversible Langevin dynamics [25].

However, there exist physical systems in which equilibrium sampling remains a difficult open problem even with enhanced sampling methods. Among them, glassy systems [26] stand as one of the most difficult to simulate in condensed and soft-matter physics. Molecular, colloidal and spin glasses are in fact known to display an extremely slow physical dynamics, which creates a major challenge to standard simulation algorithms [27, 28]. Glassy systems can in fact serve as a severe test of any newly proposed method, and can be seen as a paradigm for complex systems.

The recent discovery of generative models in artificial intelligence (AI) able to generate large structured data such as images, sound, 3D-video, and text has the potential to induce a second revolution in scientific computing [29–33]. These AI models are not only able to accurately produce complex data, but are also very fast. Speed is a central requirement in the realm of scientific computing. Several applications appeared already. In 2019, Noé *et al.* proposed the usage of normalizing flows (NF) [34], and independently Wu *et al.* variational autoregressive models [35], for Boltzmann sampling in statistical and condensed matter physics. These works have found numerous applications for sampling [36–40] and free energy calculations [41, 42]. However, despite these interesting premises, a clear view on when, where and how these methods work, and in particular their limitations and efficiency against known algorithms, is currently lacking. For standard phase transitions, promis-

* Corresponding author: gerhard.jung.physics@gmail.com

ing results have been obtained in [43, 44]. However, for hard computational problems such as complex and glassy systems, it is unclear whether they can circumvent the problem of large relaxation times. Positive results have been reported in [45, 46] for spin-glasses. On the other hand, theoretical and numerical analysis of mean-field models for structural glasses [47, 48], related to other hard problems in computer science [49], has shown that several generative models do not, and sometimes can not, have good performances. Worse, they sometimes perform less efficiently than conventional algorithms, such as local Monte Carlo [47]. These results provide a rather pessimistic view of the potentiality of machine learning (ML) techniques in the field of glassy systems. A last difficulty is that the performance of generative models is generically expected to scale very badly with the size of the system, so that applications to study phase transitions and collective effects in many-body systems appear out of sight.

Given the rapid progress made in ML studies [50–54], we feel that there is room for hope and progress. At the moment, there is a clear need of further studies to develop and test generative models in hard computational problems, and benchmark their performances against the ones of existing algorithms. The aim of this work is to perform such analysis for atomistic models of glass-forming liquids [27]. Contrary to [47, 48], we study an off-lattice, finite-dimensional glassy model which displays an extremely slow dynamics, associated to a super-Arrhenius evolution of relaxation and sampling times. It is challenging to perform numerical simulations in realistic experimental conditions for this model, as the physical relaxation time increases by more than 14 orders of magnitude towards the experimental glass transition temperature [55].

We focus on a specific two-dimensional glass-forming model that shows conventional signatures of glassy dynamics [56, 57], and represents therefore a relevant and challenging test bench for enhanced sampling methods. At low temperatures, molecular dynamics becomes rapidly unable to perform an equilibrium sampling of the configuration space, even for modest system sizes. Given the challenges mentioned above, we intentionally study a relatively small system size to separate the capabilities of NF to tackle complex landscapes from its potentially problematic scaling with system size.

We optimise a ML technique based on normalizing flows [33, 34], which we carefully benchmark against several advanced techniques introduced in the physics literature, such as parallel tempering [19], population annealing [23], and swap Monte Carlo [58]. By studying and comparing their abilities to produce an ensemble of equilibrated low-temperature configurations, we provide the first quantitative analysis of the performance of ML methods to sample realistic supercooled liquids at low temperatures. Surprisingly, our results demonstrate the great potential of such method which turns out to be much more efficient than conventional molecular dy-

namics and achieves performances comparable to parallel tempering and population annealing. Finally, we assess current limitations of these new methods and provide guidelines for further studies, in particular to improve the parametrization of the normalizing flow and to extend this technique to larger system sizes.

The paper is organised as follows. In Sec. II we define the numerical glass-forming model and explain how to assess the performance of sampling algorithms. In Sec. III we benchmark various known algorithms: molecular dynamics, swap Monte Carlo, parallel tempering and population annealing. In Sec. IV we introduce, optimize, and study the performance of a NF model. Finally, in Sec. V we collect our results, and discuss the implications for future research.

II. SETTING THE STAGE: MODEL AND SAMPLING TASK

A. A two-dimensional ternary Lennard-Jones mixture

We study a two-dimensional ($d = 2$) model introduced and developed in Ref. [56]. This is a variation of the binary Lennard-Jones mixture introduced long ago by Kob and Andersen [59] in which a third component is introduced to both improve the glass-forming ability (i.e. to prevent easy crystallization) and enable a more efficient use of the swap Monte Carlo algorithm, a strategy proposed in [58, 60]. We refer to [56, 57] for all details regarding the model parameters and simulation details.




FIG. 1. Snapshot of a typical amorphous glassy configuration of the two-dimensional model at temperature $T = 0.205$. Colors indicate different types of particles. The goal of this study is to produce a large number of independent configurations drawn from the Boltzmann distribution in Eq. (1).

We investigate systems with $N = 43$ particles using periodic boundary conditions with box length $L = 6.0$, in reduced units (see Fig. 1 for a snapshot). The unit of length is σ , which corresponds to the diameter of the large particles. When using molecular dynamics, the unit of time is the Lennard-Jones timescale $\tau = \sqrt{m\sigma^2/\epsilon}$ where m is the particle mass, and ϵ the interaction strength between large particles. For different algorithms, we express times in units of τ , in order to carefully reflect the actual computational cost of each method.

The relatively small system size is actually comparable to previous studies using normalizing flows for sampling in complex systems [34, 36, 39]. The total dimensionality of the problem is $D = Nd = 86$. The main goal of this work is to benchmark the efficiency of normalizing flows in sampling such small glassy systems according to the Boltzmann distribution,

$$\rho_*(x) = Z_*^{-1} \exp(-\beta_* U(x)). \quad (1)$$

Here, Z_* is a proportionality constant, β_* is the target inverse temperature and $U(x) = E_{\text{pot}}(x)$ is the total (potential) energy. For each configuration, we measure the total potential energy, $E_{\text{pot}} = \sum_{i \neq j} V(r_{ij})$, where $V(r)$ denotes the short-range repulsive pair interaction potential, as defined in the Supp. Mat. I of Ref. [56], and r_{ij} the relative distance between particles i and j .

Results for a larger system size, $N = 172$, are presented in Appendix C. Scaling to larger systems introduces additional challenges that have to be considered separately and are left for future work. While $N = 43$ seems a small number of particles, we emphasize that this is large enough to produce glassy dynamics and local structure for this dense fluid that are nearly equivalent to those of much larger systems [61] (see also Fig. 1). This implies in particular that equilibrium sampling even for this modest system size is already a difficult computational challenge.

B. The specific heat as a sampling task

In order to assess and compare the properties of various algorithms, we first need to define a specific sampling task to be able to test how well and how fast that task is achieved by the various algorithms.

In supercooled liquids approaching their glass transition, changes in many structural quantities are typically very modest, and deciding whether or not a given configuration is equilibrated is not straightforward. The standard solution is to measure time correlation functions, as glassy dynamics is extremely sensitive to small temperature changes, so that lack of equilibration, insufficient sampling, or small drifts are more easily detected using dynamic quantities. This approach is however not available to parallel tempering, population annealing and normalizing flows which output a set of low-temperature configurations that are not connected by any obvious dynamics.

To solve this problem we analyse the statistics of energy fluctuations measured in an ensemble of configurations. We define the average potential energy over this ensemble, $\langle E_{\text{pot}} \rangle$, and the variance of its fluctuations, which is directly connected to the specific heat as [3]

$$c_V = \frac{C_V}{N} = \frac{\langle E_{\text{pot}}^2 \rangle - \langle E_{\text{pot}} \rangle^2}{Nk_B T^2}. \quad (2)$$

Here, T denotes the temperature of the system and k_B the Boltzmann constant ($k_B = 1$ in our units). A correct estimate of the specific heat at a given temperature thus requires the production of several independent equilibrium configurations in order to correctly assess the fluctuations around the mean $\langle E_{\text{pot}} \rangle$. We conclude that the determination of the specific heat represents a well-defined task that is able to probe the capability of a given algorithm to (i) reach thermal equilibrium, (ii) sample a large number of independent configurations x representative of the Boltzmann distribution. Another advantage is that c_V does not require knowledge of a dynamics between configurations, and is thus broadly applicable to any sampling technique. Due to the generality of this sampling task we therefore believe that it is similarly suited to benchmark enhanced sampling techniques for various different complex systems.

In practice, we additionally define a convergence timescale which quantifies the computational time it takes for a given algorithm to correctly approach the equilibrium value of the energy at a given temperature. This timescale will thus allow us to rank the different algorithms by their efficiency to accomplish the requested sampling task.

We also studied alternative, previously-proposed determination of equilibration, such as different definitions for the specific heat related by a fluctuation-dissipation relation, the radial distribution function, histograms of potential energies, and density of states. See Appendix B for more details on these other approaches. We found that none of these measures can reliably be used, as they often overestimate the degree of equilibration and are blind to small deviations from equilibrium. Therefore, we focus on c_V as our main observable.

III. BENCHMARKING KNOWN SAMPLING ALGORITHMS

In this section we analyse the sampling performances of four distinct algorithms: swap Monte Carlo (SMC), molecular dynamics (MD), parallel tempering (PT) and population annealing (PA).

A. Swap Monte Carlo (SMC)

A major problem for benchmarking enhanced sampling techniques is usually the absence of a reference solution

and therefore of a clear performance measure. Here, this issue is easily settled by using swap Monte Carlo (SMC) [13, 58, 62]. The algorithm adds non-local Monte Carlo moves on top of conventional molecular dynamics simulations [63]. The Monte Carlo moves are swap moves, in which a pair of particles with different types are randomly selected and their radii are swapped. This swap move is accepted according to a Metropolis scheme. This algorithm is extremely efficient and can be used to equilibrate supercooled liquids at extremely low temperatures, including below the experimental glass transition temperature [56, 58, 60]. A detailed introduction and discussion of this algorithm can be found in Ref. [63].

Anticipating that SMC is the most efficient sampling method, we therefore perform SMC simulations to provide a benchmark for the following analysis of other sampling methods. In detail, we perform 105 swap attempts every 50 MD steps. This is the highest swap frequency that we could use without inducing small, but noticeable energy shifts, which were then affecting the quality of the benchmarking performed below. Since swap moves are not frequent, it is pertinent to use the Lennard-Jones MD time unit τ as the time unit also for the SMC method.

We first equilibrate the system for roughly $10^5 - 10^7 \tau$ starting from temperature $T_{\text{init}} = 0.5$. During equilibration, we monitor the non-equilibrium potential energy, $\langle E_{\text{pot}}(t) \rangle_{\text{neq}}$. Here, the average, $\langle \dots \rangle_{\text{neq}}$, is taken over $N_s = 64$ independent simulations, all starting from different equilibrium configurations sampled at T_{init} at $t = 0$. Afterwards, we perform SMC sampling for another $t_{\text{samp}} = 10^7 - 10^8 \tau$ (depending on the temperature) to extract the ensemble average as introduced in Sec. II B. During sampling, we also extract the time-dependent average $\langle X \rangle_t = \sum_{t_s < t} X(t_s)$, from which we obtain the time-dependent specific heat $c_V(t) = (\langle E_{\text{pot}}^2 \rangle_t - \langle E_{\text{pot}} \rangle_t^2) / (NT^2)$. In the limit $t \rightarrow t_{\text{samp}}$ we then recover the long-time average $\langle X \rangle_t \rightarrow \langle X \rangle$. We have ensured that the measured timescale reflects the actual computational cost for these SMC simulations to enable quantitative comparison of equilibration and sampling timescales.

Results for the time dependence of different observables during equilibration and then during sampling are shown in Fig. 2. The potential energy $\langle E_{\text{pot}}(t) \rangle_{\text{neq}}$ decays strongly during equilibration until it reaches a plateau. Only for temperatures significantly below the estimated glass transition temperature ($T_g \approx 0.15$) we observe that the potential energy continues to decay even beyond $t > 10^7$, suggesting that SMC falls out of equilibrium at these temperatures.

We also investigate the time dependence of the specific heat measured after the long equilibration run. Here and in the following, error bars are calculated from the variance over several independent runs. Starting from a small value at short time (when a single configuration has been probed), $c_V(t)$ rapidly accumulates on short time scales contributions from vibrations within one state (leading to $c_V \approx 1$, since we work in $d = 2$ space dimensions). At




FIG. 2. **Sampling with swap Monte Carlo (SMC).** (a) Equilibration of the potential energy $\langle E_{\text{pot}}(t) \rangle_{\text{neq}}$ ($T_{\text{init}} = 0.5$ at $t = 0$). (b) Sampling of the specific heat $c_V(t)$, which no longer reach a plateau for temperatures $T < 0.12 < T_g$. Horizontal dashed lines show the long-time averages. (c) Long-time average of the potential energy $\langle E_{\text{pot}} \rangle$ and (d) the specific heat c_V . The vertical line marks $T_{\text{SMC}} = 0.12$, which is the temperature below which SMC sampling fails.

much later times, the system visits a manifold of different states to eventually correctly sample the Boltzmann distribution. Different from the equilibration discussed above, reaching a plateau in $c_V(t)$ requires much longer times for c_V than for $\langle E_{\text{pot}} \rangle$: it takes longer to explore enough configurations to estimate c_V than simply reach-

ing an energy value close to the equilibrium one.

Explicitly, we can deduce from Fig. 2(b) that the system does not reach a plateau anymore for temperatures $T < T_{\text{SMC}} = 0.12$ during sampling of $c_V(t)$. We therefore identify T_{SMC} as the temperature below which SMC is no longer able to perform the assigned sampling task. Notice that near T_{SMC} the energy can reach a plateau, indicating that the system is very close to equilibrium, but the simulations are nevertheless not sufficiently long to sample a large enough number of independent configurations to provide the correct estimate of the specific heat.

The corresponding long-time averages of $\langle E_{\text{pot}} \rangle$ and c_V are shown in Fig. 2. We observe that the potential energy decays monotonically for decreasing temperature. Furthermore it can be seen that above $T > T_{\text{SMC}}$ the results for c_V are monotonically increasing with decreasing temperature. Once the system falls out of equilibrium near $T_{\text{SMC}} \approx 0.12$, there is a strong decrease of c_V , as found previously in many experiments and simulations in cases where the system falls out of equilibrium [64].

In the following, we will consider the SMC results down to T_{SMC} as reflecting the correct equilibrium behavior, in order to test and benchmark the performance of the alternative techniques.

B. Molecular dynamics (MD)

Molecular dynamics (MD) simulations consist in solving Newton's equations of motion with an added thermostat to control the temperature [3]. Consequently, the dynamic relaxation proceeds through realistic dynamics. Our simulations use a Nose-Hoover thermostat with relaxation time $\tau_{\text{NH}} = 1.0$ and time step of $\Delta t = 0.005$. Identical to SMC we create high temperature configurations at $T_{\text{init}} = 0.5$ and then quench the temperature to the desired value T to monitor the relaxation of the potential energy towards equilibrium. After thermalization for $t > 10^7 \tau$, we investigate $\langle E_{\text{pot}} \rangle_t$ and its fluctuations to measure $c_V(t)$. In order to perform a fair comparison, we use the same range of timescales and report the same averaged quantities for MD, SMC and all other sampling methods.

Results for MD dynamics are shown in Fig. 3. Since SMC provides equilibrium measurements down to low temperatures, we can study the difference $\Delta E_{\text{pot}} = \langle E_{\text{pot}}^{\text{MD}} \rangle - \langle E_{\text{pot}}^{\text{SMC}} \rangle$ to better quantify differences to the established SMC results which are equilibrated down to T_{SMC} . For MD, the potential energy starts to systematically deviate from the expected SMC result already for temperatures $T < 0.3$, see Fig. 3(a). This is confirmed in Fig. 3(b) which shows that the specific heat measured by MD simulations shows a peak near $T_{\text{MD}} = 0.3$, indicating lack of sampling for lower temperatures.

While equilibrium dynamics can easily be measured for MD using for instance time correlation functions, we show instead how the energy decay after a quench from




FIG. 3. **Benchmarking MD dynamics.** (a) Difference to the SMC potential energy, $\Delta E_{\text{pot}} = \langle E_{\text{pot}}^{\text{MD}} \rangle - \langle E_{\text{pot}}^{\text{SMC}} \rangle$, where $\langle E_{\text{pot}}^{\text{MD}} \rangle$ is the long-time average of the MD dynamics. (b) Long-time average of the specific heat c_V . The vertical lines mark $T_{\text{MD}} = 0.3$, below which MD sampling fails. (c) Equilibration of the potential energy characterized by the quantity $\Delta E_{\text{pot}}(t) = \langle E_{\text{pot}}^{\text{MD}}(t) \rangle_{\text{neq}} - \langle E_{\text{pot}}^{\text{SMC}} \rangle$. (d) Sampling of the specific heat $c_V(t)$. $c_V(t)$ does not reach a plateau anymore for temperatures $T < T_{\text{MD}}$. Dashed horizontal lines mark the long-time SMC results. Color code in (c, d) as in Fig. 2.

$T_{\text{init}} = 0.5$ and the time dependence of the specific heat measured during equilibrium sampling in Figs. 3(c,d). Compared to SMC, we show a narrower regime of temperatures down to $T = 0.205$. In Figs. 3(c) we observe an increasing time scale to reach the correct value of the energy which becomes impossible for $T < T_{\text{MD}}$ over the simulated time window. The lack of sampling becomes more severe when considering the specific heat which can only reach its plateau value for $T = 0.32$ but not below. Our data confirm that MD is much less efficient than SMC, as expected. More importantly perhaps, since MD follows the physical dynamics of the system, the data in Fig. 3 in fact serve as a benchmark in order to assess how much gain over the physical dynamics any enhanced algorithm can achieve [17].

C. Monte Carlo in temperature space: Parallel tempering (PT)

Parallel tempering [20], also known as replica exchange [19, 21], is a popular enhanced sampling technique applied in a wide range of fields, including spin glasses [19, 21], protein folding [65, 66], polymer melts [67], and solid state physics [68]. In the field of glass-forming liquids, it has been used to create equilibrium structures of deeply supercooled liquids [69, 70], characterize point-to-set length scales [71, 72], and analyze the physics of randomly pinned systems [73].

The key idea is to perform several MD simulations in parallel, each using the same MD parameters as explained in Sec. III B but running with a set of n different temperatures T_0, \dots, T_{n-1} . Each of these simulations is called a replica. In addition to MD steps, every N_{ex} MD steps we attempt to exchange the configurations between two replicas with adjacent temperatures. The exchange of the configuration in replica j with the one in the neighbor $j \pm 1$ is accepted according to a Metropolis scheme,

$$P_{\text{acc}}(j \leftrightarrow j \pm 1) = \exp(-(\beta_j - \beta_{j \pm 1})\Delta U), \quad (3)$$

with energy difference $\Delta U = E_{\text{pot}}(x_j) - E_{\text{pot}}(x_{j \pm 1})$ between the two configurations and inverse temperatures $\beta_j = (k_B T_j)^{-1}$. An extended derivation of this equation and efficient implementation can be found in Chapter 14.1 of Ref. [2]. Using this algorithm configurations evolve both by the physical MD dynamics but also by performing a random walk in temperature space. Low-temperature configurations can therefore follow an “easy” relaxation path by being exchanged with replicas from higher temperatures, then evolving faster at these high temperatures and subsequently being exchanged back to the low temperature. This may avoid configurations being stuck for extremely long times in deep minima at low temperatures. Physically, if the basins relevant at low temperatures are also sampled at high temperature, PT can become a very efficient method [27].

The most important factor to optimize PT simulations is the choice of the replica temperatures. On the one hand, large temperature differences will significantly reduce the acceptance rate for exchange events and therefore slow down the exploration of the temperature space. On the other hand, a large number of replicas implies a larger computational effort. After trial and error, we finally use $n = 8$ replicas with temperatures $T = 0.4, 0.359, 0.32, 0.287, 0.256, 0.229, 0.205, 0.183$. To arrive at this choice, we have started with a small number of replicas and systematically increased their number until we found the optimal result in the given computational time. This choice is rationalized by a significant overlap between energy distributions at neighboring temperatures and therefore a large acceptance rate $\langle P_{\text{acc}} \rangle > 0.25$ for all replicas. We also checked that smaller temperature steps do not improve the results. We optimized the maximal and minimal temperatures in the above range to finally settle on this list of $n = 8$ replicas. We also




FIG. 4. **Benchmarking parallel tempering (PT) simulations.** The description is the same as for Fig. 3. The vertical dashed line in (b) represents $T_{\text{PT}} = 0.23$ below which PT sampling fails.

choose $N_{\text{ex}} = 5000$ which is a reasonable choice between too frequent or too infrequent temperature swaps, but the results are not very sensitive to this specific choice.

The results for the benchmarking of PT are shown in Fig. 4 using the same organization as for MD. Comparing to Fig. 3 shows that PT is superior to MD. Within errorbars, PT predicts the correct potential energy in a temperature regime in which MD is already substantially out of equilibrium. Using the specific heat as a sharper test for sampling in Fig. 4(b), we conclude that PT succeeds in the sampling task down to $T_{\text{PT}} = 0.23$, below which c_V decreases as a result of insufficient sampling. This temperature is considerably lower than $T_{\text{MD}} = 0.3$, but much higher than $T_{\text{SMC}} = 0.12$.

To understand better the efficiency and limits of the PT sampling we turn to the energy decay after a quench from an initial condition where all n replicas are initialized at a high temperature $T_{\text{init}} = 0.5$, see Fig. 4(c). Interestingly, the relaxation time of $\Delta E_{\text{pot}}(t)$ is only weakly dependent on the temperature. This behavior is qualitatively different from the MD results. Therefore, while PT

accelerates dynamics at low temperatures, it also slows down the dynamics at higher temperatures compared to MD. This is the direct result of the nature of PT exchange events: since replicas travel across the entire temperature spectrum, there is a nearly unique emerging timescale controlling the approach to equilibrium of the entire simulation composed of n replicas. In other words, different temperatures are no longer independent when using PT, and the relaxation is in effect slaved to the slowest replica. This conclusion also explains why we did not include replicas with even lower temperatures into the PT dynamics as it negatively impacts the performance of the PT simulations.

The time dependence of the specific heat, $c_V(t)$, is more interesting, see Fig. 4(d). Here, the equilibrium plateau in c_V is reached faster for higher temperatures, and the corresponding timescale becomes longer than our simulation time for $T < T_{PT}$, explaining the spurious peak in the measured specific heat in Fig. 4(b). The time dependent relaxation can thus be used to quantify the speedup offered by PT simulations over MD, and this will be discussed in Sec. V A.

D. Population annealing (PA) and reweighting (RW)

Population annealing (PA) is deeply rooted in the reweighting (RW) technique known from statistical mechanics, which we recap first.

Given a set of R configurations, $\{x_i\}$, with $i = 1, \dots, R$ taken from the Boltzmann NVT ensemble at temperature T_1 , it is possible to reweight these configurations to perform an equilibrium average at a different temperature T_2 [74]. To this end, one assigns a new Boltzmann weight $W_i = \exp(-(\beta_1 - \beta_2)U(x_i))$ to each configuration i . The ensemble average of an observable $A(x)$ at T_2 is given by

$$\langle A \rangle_{T_2} = \frac{\sum_i W_i A(x_i)}{\sum_i W_i}. \quad (4)$$

Reweighting is used extensively for free energy calculations in molecular simulations [75, 76]. The method, however, only works efficiently for small enough temperature steps so that the weights W_i remain meaningful.

In PA, a large set of configurations is used to perform small temperature steps to gradually anneal the temperature to the target low-temperature, while reweighting their Boltzmann weights at each step [22, 23, 77–80]. In practice, the PA algorithm works as follows. We first create R configurations at an initial, high temperature $\beta_1 = (k_B T_1)^{-1}$ and evaluate the Boltzmann weight of each configuration x_i as $W_i = \exp(-(\beta_1 - \beta_2)U(x_i))$. Here, T_2 should be slightly smaller than T_1 . We then create on average τ_i copies of each configuration i , where τ_i is given by

$$\tau_i = R \frac{W_i}{\sum_{k=1}^R W_k}. \quad (5)$$

Recently, different schemes were compared to numerically implement Eq. (5) [80]. We apply the “systematic resampling” scheme, which was the most efficient for a constant population size R . Following resampling, we finally perform M MD steps on each copy at temperature T_2 to help thermalize the configurations at the new temperature T_2 . This ends the annealing from T_1 to T_2 . This annealing step $T_1 \rightarrow T_2$ is then repeated several times until the final target temperature T is reached. Each annealing step consists in (i) resampling the population (ii) a small number of M MD steps for each configuration. More extended derivations of the technique and algorithms can be found in Ref. [23].

For the choice of annealing temperatures in PA, we use the same series T_1, \dots, T_{n-1} used for PT in Sec. III C. This is reasonable since the PT temperatures were optimized to provide good overlaps between the probability distributions of potential energy, which also controls the quality of the reweighting in Eq. (4). Contrary to PT, the annealing procedure in PA is unidirectional as the population flows from T_1 to T_{n-1} but no information is carried backwards. As a result, including lower temperatures is harmless (at worse, PA sampling fails) and so we include two lower temperatures $T_n = 0.164$ and $T_{n+1} = 0.148$. We perform high temperature MD simulations at $T_1 = 0.359$ and save configurations every $10^4 \tau$, which corresponds roughly to the MD structural relaxation time at this T_1 . This choice ensures that within a similar computational effort invested into PT we can create an initial set of $R = 2 \times 10^5$ statistically independent configurations. In addition, this comparison enables us to assign a computational time $t = R \times 10^4 \tau$ to the PA task, as the annealing steps themselves can be efficiently performed. None of the above choices critically affects the result when reasonably changed. The most critical parameter is the number M of MD relaxation steps. A too small number $M < 10^3$ leads to tiny but systematic differences in the observed $\langle E_{pot} \rangle$ and c_V . We therefore choose $M = 5 \times 10^3$. Since the creation of the initial set of R configurations is the computational bottleneck, such a large M value does not significantly increase the computational effort.

It is instructive to compare the gradual population annealing from T_1 to a given target temperature T with a direct reweighting performed in a single step $T_1 \rightarrow T$ directly using Eq. (4) applied to the entire initial population of configurations created at T_1 , see Fig. 5. We observe that RW is already much more efficient than MD dynamics with correct energy and specific heat obtained down to $T_{RW} = 0.25$. The effect of the gradual annealing and resampling performed within PA improve the RW results dramatically, and nearly-correct energy values are predicted down to the lowest temperature. A more careful inspection of the c_V data shows however that PA sampling fails below $T_{PA} = 0.19$.

Different from MD and PT, there are no separate equilibration and sampling procedures within PA. Nevertheless, it is possible to provide an equivalent time depen-




FIG. 5. **Benchmarking population annealing (PA) and reweighting (RW).** The description is the same as for Fig. 3. The time dependence shown in panels (c) and (d) is obtained by using different numbers of initial configurations R . The vertical dashed line in (b) represents $T_{\text{PA}} = 0.19$ below which PA sampling fails.

dence to both $\Delta E_{\text{pot}}(t(R))$ and $c_V(t(R))$ by following the evolution of the PA performance as a function of the population size R produced at the initial temperature T_1 , because by definition correct sampling is obtained in the limit $R \rightarrow \infty$, just as correct sampling is performed in the large time limit for any of the other algorithms. Convergence in the large time or population limit is obvious, since system size and thus energy barriers are finite. However, our problem is to reach convergence in a tractable computational timescale. We can then convert the population size R into a computational timescale using the dictionary $t = 10^4 R$, which corresponds to the effective computational time invested into creating the set of R configurations.

The results are shown in Figs. 5(c,d) which illustrate the convergence of the energy and its fluctuations to the correct values as R is increased. Differently from PT, we see that the equilibration of the potential energy slows down with temperature, see Fig. 5(c). We can still observe the relaxation of $c_V(t)$ towards its equilibrium plateau value. In fact the R dependence of c_V ,

and its eventual convergence to a plateau at large R , serves as a stringent test of the quality of sampling with PA. In particular, we confirm that PA sampling fails below $T_{\text{PA}} = 0.19$. This result shows that PA performs slightly better than the two previous sampling methods: $T_{\text{PA}} < T_{\text{PT}} < T_{\text{MD}}$, a result that could not be anticipated based on previous efforts. An additional advantage of PA is that the task of sampling an initial set of R independent samples at the high temperature T_1 can be easily parallelized, by running several independent simulations in parallel.

IV. SAMPLING BY NORMALIZING FLOWS

In Sec. III we established benchmarks for enhanced sampling techniques known from physics. This sets the stage for a thorough analysis of the performance of the ML method of normalizing flows (NF) [33, 81]. Since NF are relatively new methods, we first provide a general introduction before describing our implementation and the main results.

A. Continuous normalizing flows (NF)

The general idea of normalizing flows is to learn an invertible mapping between two probability distributions: a prior distribution $\rho_P(x)$, from which we can sample easily (Gaussian random numbers, high temperature liquids), and a target distribution, $\rho_*(x)$, in our case the Boltzmann distribution [34]. The mapping is in general only approximate, so one needs to reweight the configurations obtained by the NF. A more accurate mapping leads to lower rejection. Normalizing flows found applications in computer vision [82], sampling via Markov Chain Monte Carlo [37, 83, 84], lattice field theories [85] and condensed matter physics [34, 40].

Boltzmann generators are the first application of normalizing flows for sampling of complex systems in condensed matter [34]. Compared to applications such as image generation, a specific property of Boltzmann generators is that the target distribution is known and corresponds to the Boltzmann distribution (see Eq. 1). The challenge for Boltzmann generators in statistical physics is to efficiently sample from this distribution using the learned NFs. For a general introduction to Boltzmann generators see Refs. [34, 39].

Here, we use equivariant, continuous normalizing flows [86]. This specific NF has the advantage of being equivariant to translations, rotations and permutations and thus mirrors the fundamental symmetries of the underlying physical system. In Ref. [86] detailed benchmarking compared to discrete NF layers and gradient flows has been performed on a similar problem showing the superiority of the equivariant continuous NF for particle systems over discrete ones. Continuous NFs transform the prior distribution into the target distribution

by learning a time- and space-dependent vector field, $v(x_v(t), t)$, $t \in [0, 1]$, which can be interpreted as force field,

$$\frac{d}{dt}x_v(t) = v(x_v(t), t), \quad x_v(0) = x_0 \text{ drawn from } \rho_P(x). \quad (6)$$

We define the invertible transformation F , as $x_v(t=1) \equiv Fx_0$ [87]. Importantly for the calculation of the loss function for training, we can evaluate the transformation of the prior probability distribution [86],

$$\log \rho_F(Fx_0) = \log \rho_P(x_0) - \int_0^1 dt \operatorname{div} v(x_v(t), t). \quad (7)$$

For a given transformation F , the NF thus produces a “push-forward” probability distribution $\rho_F(x)$ given by Eq. (7), which is different from the target $\rho_*(x)$ if the transformation is not perfect. Similarly, $\rho_{\bar{F}}(x)$ emerges from the inverse transformation \bar{F} when applied on the true distribution $\rho_*(x)$.

The model used for the force field is a sum of pairwise potentials which depend on the distance between particle pairs [86],

$$v(x(t), t) = \nabla_x \Phi(x(t), t), \quad (8)$$

$$\Phi(x(t), t) = \sum_{ij} \tilde{\Phi}(d_{ij}(t), t), \quad (9)$$

with $d_{ij}(t) = |x_i(t) - x_j(t)|$. The learnable weights $\{w\}$ of the NF are the parameters of the potential field $\tilde{\Phi}(d, t)$, which is parameterized using Gaussian radial basis functions in both distance d and time t . The calculation of the divergence terms is numerically exact and stable, as detailed in Ref. [86]. Our implementation is based on the public code `bgflow` provided by the authors of this publication. In the following, our goal is to transform an easy-to-sample high-temperature distribution, ρ_P , at inverse temperature β_P into a low-temperature target distribution.

B. Loss function and training

Normalizing flows can be trained using the Kullback-Leibler divergence as minimizable loss function L , which quantifies the similarity between the target distributions and the transformed NF distributions, $L = \alpha D_{\text{KL}}(\rho_F || \rho_*) + (1 - \alpha) D_{\text{KL}}(\rho_* || \rho_{\bar{F}})$. We differentiate in L between two different training contributions. The first term [86],

$$D_{\text{KL}}(\rho_F || \rho_*) = \int_{\Omega} [\beta_* U(x) + \log \rho_F(x)] \rho_F(x) dx, \quad (10)$$

is based on having provided a set of high-temperature configurations, $\{x_0\}$, which are transformed using F , $x = Fx_0$. A loss based on $D_{\text{KL}}(\rho_F || \rho_*)$ is called “variational” or “energy-based” training. This equation can

be discretized as a sum over individual configurations i ,

$$D_{\text{KL}}(\rho_F || \rho_*) = \sum_i \left[\beta_* U(x^i) - \int_0^1 dt \operatorname{div} v(x_v^i(t), t) \right], \quad (11)$$

where we have dropped $\beta_P U(x_0^i)$ from Eq. (7) since it is a constant that does not influence the loss function.

It has been found in many studies that the training can be improved using a small set of low temperature configurations, $\{x_*\}$ sampled from $\rho_*(x)$. The second term [86],

$$D_{\text{KL}}(\rho_* || \rho_{\bar{F}}) = \int_{\Omega} [-\beta_* U(x_*) - \log \rho_{\bar{F}}(x_*)] \rho_*(x_*) dx_* \quad (12)$$

quantifies the similarity between the transformed low-temperature configurations, $x = \bar{F}x_*$, and the prior distribution. This second contribution is also known as “maximum likelihood” training. We also discretize this equation to use it for training the NF.,

$$D_{\text{KL}}(\rho_* || \rho_{\bar{F}}) = \sum_i \left[\beta_P U(\bar{F}x_*^i) - \int_0^1 dt \operatorname{div} v(x_v^i(t), t) \right]. \quad (13)$$

We have systematically analyzed the optimal value for α , as discussed in detail in App. A. We find that $\alpha = 0.5$ leads to the most stable training procedure and the best final result (see Fig. 9). In particular, this figure also highlights the importance of including the maximum-likelihood training. We also tested protocols in which α changes during the training procedure as suggested in Ref. [86], but this did not affect the results qualitatively.

Using a set of high- and low temperature configurations, $\{x_0\}$ and $\{x_*\}$, enables the optimization of the learnable parameters $\{w\}$, which quantify the strength of the force field $v(x(t), t)$. This step is achieved using the above-defined loss function and stochastic gradient descent. The errors are backpropagated using automatic gradient differentiation of the discretized NF in Eq. (6) implemented by PyTorch, which replaces the typical backpropagation known from artificial neural networks [86].

In our case, however, we face the problem that we do not have access to any low temperature configurations, because generating them is the whole purpose of the NF. We bypass this contradiction by using an iterative procedure. We use reweighting to generate some approximate low-temperature configurations. Using these configurations, we train a NF in the first iteration and apply it to create an improved set of low-temperature samples. In a second iteration we then utilize this improved set for the training of a second NF, which finally produces the low-temperature configurations that are analysed below. We can iteratively improve the performance of the NF with this iterative procedure but using more than two iterations did not lead to significant changes. In fact, we have tested that using low-temperature configurations prepared with SMC for training only marginally improved the performance of the NF, which thus confirms the efficiency of the proposed iterative procedure.

In fact, our approach is similar in spirit to Ref. [88] which also avoids the usage of target configurations by applying annealed importance sampling.

There are many hyperparameters that can be tuned to parameterize the NF and optimize the training procedure, including the number, location and time-discretization of the radial basis functions, training parameters and batch sizes. However, we found that the results are not very sensitive to the explored choices of these parameters. See Appendix A for more details.

C. Unbiasing the NF distribution

After training the NF, we obtain a transformation F , which is used to transform all available high-temperature configurations, $x^i = Fx_0^i$. Because the mapping performed by the NF is only approximate, the resulting set of configurations are biased, i.e. they do not exactly sample the Boltzmann distribution at temperature T . This can be corrected by performing an unbiasing step. We calculate the statistical weight of each transformed configuration as

$$W_i = \exp \left[\beta_P U(x_0^i) - \beta_* U(x^i) + \int_0^1 dt \operatorname{div} v(x_v^i(t), t) \right]. \quad (14)$$

Similar to RW, we then use the weights W_i to create a set of low temperature configurations that can sample the Boltzmann distribution. The NF therefore not only provides transformed configurations, but also their statistical weights W_i , which describe the effective weights of each configuration i at the low temperature β_* . In direct RW, only the second step is performed, and thus NF has the potential to provide a large improvement over RW by first transforming the original configurations, resulting in larger statistical weights in Eq. (14).

D. Results: Sampling efficiency of NF

We now show the performance of the NF in Fig. 6 using the same metrics introduced for the other sampling methods, comparing long-time averages for the energy and the specific heat to SMC results. Regarding efficiency and timescales, we can analyze NF in much the same way as we did for PA in Sec. III D. In particular, NF inherits the computational time $t = R \times 10^4 \tau$ of PA since it uses the same initial samples. Just as for PA, the sampling part of NF is computationally significantly more expensive than the subsequent transformation and unbiasing steps.

We first compare NF to conventional MD results. From Fig. 6(a,b), we conclude that the results for NF are much closer to the SMC groundtruth than what is achieved by MD simulations (see Fig. 3). From Fig. 6(b) we conclude that NF produces an equilibrium ensemble down to $T_{\text{NF}} = 0.2$, which is significantly smaller than $T_{\text{MD}} = 0.3$.




FIG. 6. **Benchmarking machine-learned normalizing flows (NF)**. Description is the same as for Fig. 5. For long-time averages in (a, b) we also show results for the combination of population annealing and normalizing flow (NF/PA). The vertical dashed line in (b) represents $T_{\text{NF}} = 0.2$ below which NF sampling fails.

Thus, the NF generative modeling approach is indeed an enhanced sampling method, in the sense that it works better than the physical dynamics in sampling low temperature configurations of the glassy system under study. Given published results regarding generative models for atomistic [39] or glass [47] models, this is an interesting result.

It is interesting to compare NF also with the direct RW approach studied in Sec. III D as both methods use the same high-temperature configurations to predict low-temperature properties. The key step distinguishing the two methods is the NF transformation in Eq. (6) itself. The fact that NF performs much better than RW implies that the NF is able to efficiently transform the high-temperature configurations so that the transformed configurations are much closer to equilibrium than the original ones.

As for PA, we can follow the approach to equilibrium of the potential energy, see Fig. 6(c), and the specific heat, see Fig. 6(d) when the size of the initial population R is increased, which can be translated into timescale.

These data allow us to define and measure a growing timescale for equilibration and sampling which becomes longer than the simulated time for $T < T_{\text{NF}}$.

Since NF outperforms MD simulations, it is pertinent to compare its performances with known enhanced techniques, which justifies our efforts to carefully benchmark various methods in Sec. III. Broadly speaking we find that all techniques (PA, PT, NF) perform nearly similarly, with NF and PA being slightly better than PT with the rough hierarchy, $T_{\text{PA}} \lesssim T_{\text{NF}} < T_{\text{PT}}$. This detailed comparison and ranking of several techniques is one of the main results of this work: it provides evidences of the usefulness of NF for the difficult sampling problem of finite dimensional glassy systems.

Given the success of NF over direct RW, it is tempting to combine the NF method with the successful PA approach in Sec. III D, in order to possibly improve the performance of both these methods. In this combined approach, we use the global framework of PA, but we replace the second step in the PA algorithm (where copies are created from weights W_i calculated using the Boltzmann distribution) by the usage of a trained NF to transform the configuration and calculate the new weights W_i . We refer to this mixed method as “NF/PA”.

The results shown in Fig. 6 are however disappointing. Although they are slightly better than NF, showing that multiple small steps are better handled than a large one, they are not better than PA. This is surprising, since NF clearly performs better than RW in the one-shot annealing procedure and PA is based on consecutive RW steps. Our interpretation is that for very small temperature steps, RW becomes actually superior to NF, presumably because it uses the exact expression of the Boltzmann distribution.

E. Analysis of the effective sample size

Different from population annealing and parallel tempering, NFs have the capacity to produce low-temperature configurations in one shot, without the introduction of a large number of intermediate temperature steps.

One practical consequence of such one-shot annealing is the possibility to define an interpretable effective sample size based on Kish’s formula [89],

$$R_{\text{eff}} = \frac{\left(\sum_{i=1}^R W_i\right)^2}{\sum_{i=1}^R W_i^2}, \quad (15)$$

using the statistical weights W_i introduced in Eq. (14). The effective sample size describes roughly how many independent configurations have been produced during the annealing procedure. The physical idea behind Eq. (15) is clear: the effective sample size $R_{\text{eff}} \approx R$ if all weights W_i are comparable, whereas $R_{\text{eff}} \ll R$ when a few samples have a much larger weight than all others, indicating poor sampling.




FIG. 7. Effective sample size R_{eff} , as defined in Eq. (15), for the PA and NF results shown in Fig. 5 and Fig. 6, respectively. The corresponding symbols on the x-axis mark the temperatures at which the sampling of each given method starts to fail.

Starting from an initial set of $R = 2 \times 10^5$ samples we observe for NFs an exponential decay of the effective sample size with temperature T in Fig. 7. At the temperature $T_{\text{NF}} = 0.2$, identified before as the temperature at which NF sampling starts to fail, the effective sample size is $10^2 < R_{\text{eff}} < 10^3$. This order of magnitude is consistent with our empirical findings that at least 100 independent samples are required for a proper representation of the equilibrium ensemble at temperature T . This analysis suggests that the effective sample size R_{eff} can be used as an independent and easy tool to check for equilibration when using NF as an enhanced sampling technique.

In Fig. 7, we show the evolution of R_{eff} evaluated during the gradual annealing employed for PA in Sec. III D. We observe a much slower decrease when temperature is reduced, with an effective size that remains quite large, $R_{\text{eff}} \sim 10^4$ when crossing the temperature T_{PA} , indicating that the sample size is a poor indicator of adequate sampling in that case. This presumably results from the resampling of the algorithm whereby the samples that eventually dominate the low temperature behavior are replicated more often than others, which introduces strong correlations in the population. These correlations make the use of Kish formula inefficient. Other independent measures have been suggested for PA to test equilibration, but they require more involved analysis [90].

V. DISCUSSION: WHAT IS THE MOST EFFICIENT SAMPLER?

A. Quantitative comparison between techniques

For each technique, we have obtained a temperature below which the assigned sampling task, i.e. measuring

the specific heat, starts to fail. This allowed us to rank the various techniques. For the particular glass model studied here, we find that SMC is by far the best technique, with $T_{\text{SMC}} = 0.12$. Then come the three enhanced algorithms, PA, NF and PT with $T_{\text{PA}} = 0.19$, $T_{\text{NF}} = 0.2$, $T_{\text{PT}} = 0.23$, which all perform much better than conventional MD with $T_{\text{MD}} = 0.3$. For comparison, we recall that the mode-coupling crossover is near $T = 0.3$ and the experimental glass transition temperature near $T = 0.15$ [56].

This ranking does not easily translate into an actual computational speedup, or efficiency gain, which may depend on the temperature. For each algorithm and each temperature, we showed that the approach to equilibrium of the energy or the convergence timescale for the specific heat can both be recorded to assign a representative sampling timescale. In practice, we use the former and define a timescale τ_c as $\langle \Delta E_{\text{pot}}(\tau_c) \rangle_{\text{neq}} = 0.5$. As a rule of thumb, a smaller τ_c implies a smaller computational cost and thus improved performance of the technique.

We collect the results for the evolution of τ_c for all algorithms in Fig. 8. This provides a more detailed comparison between algorithms. Starting with very high temperatures, we observe in Fig. 8 that MD is more efficient than the three enhanced sampling techniques, PT, PA and NF. In Sec. III C we explained this finding for PT by the coupling between low and high temperatures through the temperature swap exchanges. The explanation is different for PA and NF which are less efficient due to the quite coarse sampling performed at high temperatures with a time $10^4\tau$ between each stored configurations, see Sec. III D. This time scale was the best compromise we found empirically between the time invested into the annealing and efficient sampling for the lowest temperatures. This could clearly be reduced if the focus was on higher temperatures. Given that MD is very efficient in this regime, this is not a crucial endeavour.

When temperature decreases, Fig. 8 shows that the MD timescale increases more rapidly than any other technique, and MD sampling is therefore the first to fail. The relaxation times of the three enhanced sampling methods, PT, PA and NF seem to roughly follow the same temperature dependence, with minor differences between them. Their behavior appears to be approximately Arrhenius, but the apparent energy barrier is much smaller than for MD. Notice that for PT the timescale τ_c does not take into account the fact that n replicas need to be simulated in parallel. In the same vein, we note that the computational time for PA and NF is mostly due to the preparation of a large population of independent configurations at relatively high temperatures. This task can trivially be parallelized by running a large number of independent simulations, thereby making PA and NF potentially much more efficient than PT where no additional parallelization can be implemented.

Interestingly, SMC seems to follow the same Arrhenius dependence of the three enhances methods, at least in this temperature regime, but with a prefactor that is




FIG. 8. **Temperature evolution of the efficiency timescale for all algorithms.** In practice, τ_c quantifies the approach to equilibrium of the potential energy. We compare swap Monte Carlo (SMC), molecular dynamics (MD), parallel tempering (PT), population annealing (PA), and normalizing flows (NF). The corresponding symbols on the x-axis mark the temperatures at which the sampling of each given method starts to fail.

considerably smaller by about four orders of magnitude. This large difference quantitatively explains why SMC is the most efficient sampling technique for this system.

Despite the success of SMC, it is encouraging that NF can truly compete with state-of-the-art sampling techniques such as PT and PA, with a significant speedup over MD dynamics. At the lowest temperature where NF still operates, $T_{\text{NF}} = 0.2$, the speedup over MD dynamics is about four orders of magnitude in relaxation time.

B. Perspectives

In this work we compared state-of-the-art enhanced sampling techniques for equilibrating supercooled liquids with a new method based on the machine learning technique using normalizing flows. Our results demonstrate the potentiality of ML methods to equilibrate model supercooled liquids at low temperature. In fact, the NF method applied to small systems at very low temperatures has a performance comparable to the sampling methods developed for complex systems, such as parallel tempering and population annealing. This very good result is obtained despite the fact that NF does not introduce a large set of replica (as in PT) or intermediate annealing temperatures (as in PA) and directly targets low temperatures in one shot. This positive conclusion suggests that all important modes of the low temperature states are already present, although affected by thermal fluctuations, in the high-temperature regime. However, NF are, like PT and PA, suboptimal with respect to the

swap Monte Carlo technique.

We have focused on small systems with $N = 43$ at very low temperatures in $d = 2$. As demonstrated, this provides a challenging setting for all sampling methods for an atomistic model with realistic interactions. Applying the sampling methods to larger system sizes introduces new challenges for all of them. SMC and MD methods do not suffer too much with larger N , since the computational time increases linearly with N while their performances do not degrade. The situation is different for PT, PA and NF, for different reasons. We provide in Appendix C results with a system that is four times larger with $N = 172$, showing poorer performances. Addressing the challenge of scaling these algorithms to larger system sizes should be the focus of a dedicated future work. In fact, even a very accurate NF method would eventually lead to an increasing level of rejection in the reweighting step for large sizes, as the statistical weight should scale as $\exp(-cN)$ with c a finite constant, which account for a small difference between the generated distribution and the Boltzmann target. This generic argument does not take into account the complex nature of the glassy configuration space, which may very well lead to additional sampling issues at larger system sizes.

Still, the observed performance of NF should encourage further work towards the development of improved techniques. For instance, it would be interesting to study more complex parametrization of the flow than the one we used. Possible candidates are: equivariant coupling flows [91], which combine the efficiency of coupling flows while maintaining equivariance, equivariant flow matching [92, 93], which uses alternative loss functions for training [94], annealed flow transport Monte Carlo [95], or approaches based on diffusion models [96–98]. Additionally, it might be possible to combine NF layers with intermittent periods of SMC dynamics to create a stochastic normalizing flow as in Ref. [99].

The benchmarks outlined in this manuscript aim to accelerate and simplify the development of such sophisticated machine learning methods for sampling of complex systems. It is anticipated that any enhancement will manifest directly in the resulting relaxation time. Similar benchmarking for other complex system would be very valuable. We therefore believe that this manuscript marks an important step on the quest of finding methods that outperform traditional enhanced sampling techniques and, potentially, even the swap Monte Carlo technique.

ACKNOWLEDGMENTS

We thank the Noé group for publicly providing their library `bgflow`, and S. Ciarella, M. Gabrié, and F. Zamponi for discussions. This work was supported by a grant from the Simons Foundation (#454933, L. Berthier, #454935 G. Biroli).

Appendix A: Details on the NF method

We provide more information on the hyperparameters used in the NF method. Our implementation is based on the `bgflow` library [34, 86] (<https://github.com/noegroup/bgflow>), which was extended to include periodic boundary conditions and multiple particle types. Thus any detail provided in Ref. [86] similarly applies to the present manuscript.

1. Hyperparameter

Most notably, we discretize the differential equation flow in Eq. (6) using just $N_t = 1$ (first iteration) or $N_t = 3$ timesteps (second iteration) in a multi-step fourth order Runge-Kutta scheme [86]. This is a strong reduction of the complexity of the normalizing flows, but we have empirically found that larger values of N_t do not improve the results. A consequence of this choice is that the transformation of the configurations amounts to quite small displacements $\Delta x \ll \sigma$ within the particle cages, rather than large-scale rearrangements. We have tried intensively to learn more general models, starting from $T \rightarrow \infty$ (uniformly distributed particles), but none of these models was able to propose acceptable configurations for low temperatures and reach accuracies comparable to the results presented in the main text.

We include 80 independent Gaussian radial basis functions centered non-uniformly at distances d in the range $0.65 \leq d \leq 2.8$. In total this adds up to 966 learnable parameters (1938 for $N_t = 3$). The gradient decent is based on an Adam optimizer with accuracy 10^{-4} . For the first iteration we train on 512 different structures using only one epoch, the second iteration uses four epochs and 4096 different structures. The batch size is always 64 structures.

2. Ablation study for mixing parameter α

We have introduced in Sec. IV B the parameter α in the loss function which interpolates continuously between energy-based training ($\alpha = 1$) and maximum-likelihood training ($\alpha = 0$). Which is the best choice for α ?

To answer this question we have performed different training procedures for different values of α . We report in Fig. 9 the effective sample size R_{eff} which we have identified in the main manuscript as an important factor to quantify the performance of the NF. The figure highlights a maximum near $\alpha \approx 0.5$, which is our final choice. For all other choices, the effective sample size is significantly lower.

In particular, Fig. 9 rules out the possibility to perform pure energy-based training ($\alpha = 1$) which would avoid the iterative procedure of finding low temperature training configurations described in Sec. IV B. In fact, we




FIG. 9. Effective sample size R_{eff} , as defined in Eq. (15), for different values of α at $T = 0.205$. The result for normalizing flows (NF) is compared to reweighing (RW), see Sec. III D.

find that the problem with $\alpha = 1$ is not mode-collapse as in other studies in the field of computer vision. For example, we have attempted the β -NF approach in which the entropy term (i.e., the second term in Eq. (11)) is scaled by a factor $\beta > 1$ [100, 101] and we did not find any improvement. Instead the only solution we found to increase R_{eff} for $\alpha \neq 0.5$ is early stopping, which hints to some instabilities in the learning. Nevertheless, even with early stopping $R_{\text{eff}}(\alpha)$ never reaches the value $R_{\text{eff}}(\alpha = 0.5)$.

In conclusion, this analysis shows that $\alpha = 0.5$ is the optimal choice for the mixing parameter.

Appendix B: Additional criteria for equilibration

In the main text we state that the best way to validate sampling is by verifying whether $c_V(t)$ attains a plateau. During our research, we have tested several different possibilities, which we briefly describe.

1. Fluctuation-dissipation theorem for specific heat

A popular way to validate equilibrium sampling is by calculating the specific heat using two different formula. The first one is used throughout this work and corresponds to the variance of fluctuations in potential energy, $c_V^{\text{VAR}} = (\langle E_{\text{pot}}^2 \rangle - \langle E_{\text{pot}} \rangle^2) / NT^2$. A second definition is based on the temperature derivative of the average potential energy, $c_V^{\text{DER}} = N^{-1} \partial \langle E_{\text{pot}}(T) \rangle / \partial T$. For equilibrium samples, these two definitions yield the identical result by virtue of the fluctuation-dissipation theorem. Any difference between these two quantities can therefore reveal a departure from equilibrium.

The results in Fig. 10 show that this indicator does not clearly signal departure from equilibrium. In fact, it




FIG. 10. Specific heat c_V calculated for various sampling techniques using two different definitions, c_V^{VAR} from fluctuations and c_V^{DER} from derivative. In equilibrium both are related by a fluctuation-dissipation theorem. The data indicate however that when the system falls out of equilibrium the fluctuation-dissipation theorem remains valid and both expressions similarly depart from equilibrium.

seems that when a given sampling technique departs from the SMC solution both definitions of the specific heat depart similarly at the same temperature, but remain consistent with each other within the error bar. A slightly better indicator of departure from equilibrium is the notable increase of the error bars, which indicate increasing correlations between configurations, indirectly revealing lack of ergodicity. It is however difficult to transform this observation into a clear-cut criterion for equilibration.

2. Probability distribution of potential energy

We have analyzed in detail the average potential energy, $\langle E_{\text{pot}} \rangle$ and its variance in the form of the specific heat, c_V . Here we investigate whether the full probability distribution of potential energies gives additional information, in particular on whether equilibrium sampling has been achieved.

We observe that histograms do not yield much more information compared to the first two moments, see Fig. 11. At $T = 0.256$ where equilibrium sampling already fails for MD, the energy histograms remain quite close, with small deviations only visible in the left tail at low energy values. At $T = 0.205$, the MD dynamics are completely out-of-equilibrium which can be observed by a clear shift compared to the SMC result. However, by rescaling the first and second moment of the SMC distribution (dashed line) we observe nearly perfect overlap with the MD results (blue squares).

We further exploit these data and evaluate the density of state $G(E_{\text{pot}}) \propto P(E_{\text{pot}}, T) \exp(\beta E_{\text{pot}})$. The interest




FIG. 11. Probability distribution function of the potential energy, $P(E_{\text{pot}}, T)$, calculated for various sampling techniques at different temperatures T .

of the density of state is that it is a temperature independent quantity which is only accurately obtained if proper equilibrium sampling of energy fluctuations is performed. As such it has been used as a tool to assess the degree of equilibration [69].

Our results are shown in Fig. 12. Since the density of states is only known up to a prefactor, each set of curves is arbitrarily shifted to maximize the overlap between estimates of the density of states stemming from different temperatures for a given algorithm. In addition, the result for each method is shifted independently for better visualization.

The excellent data collapse for the SMC data confirms that equilibrium sampling is achieved down to very low temperatures. The expected temperature-independent mastercurve is obtained when stitching together the data from $P(E_{\text{pot}}, T)$ obtained at different temperatures.

Interestingly, the MD data indeed reveals, that the ensemble falls out of equilibrium since no perfect overlap can be achieved. This shows that the low-energy tails of the energy distribution are not properly sampled, in a way that is perhaps clearer than in Fig. 11.

In contrast, even at $T = 0.148$, the data extracted from PA sampling shows perfect overlap although we know that they do not perfectly represent the equilibrium ensemble, as identified above. The reason for the qualitative difference between MD and PA is two-fold: (i) MD falls out of equilibrium much more violently, in particular when investigating the potential energy, while in PA the differences are much smaller even when the system is out of equilibrium. (ii) The number of samples used to create the histograms is much smaller in PA, since each sample needs to be treated independently, making it impossible to maintain a huge set. The very subtle difference in $G(E_{\text{pot}})$ observed for MD is therefore nearly invisible for PA.

We conclude that $G(E_{\text{pot}})$ can detect non-equilibrium




FIG. 12. Density of states, $G(E_{\text{pot}})$, calculated for various sampling techniques. For each technique, estimates of the $G(E_{\text{pot}})$ obtained at different temperatures are stitched together to form a mastercurve. Each mastercurve is vertically shifted, for clarity.

properties, but it requires significant departure from equilibrium and huge datasets. In other words, this is not a very sensitive test of equilibrium sampling.

3. Radial distribution function

There are also two different ways to calculate the radial distribution function, $g(r)$, in particle systems at thermal equilibrium. The first traditional approach is based on histograms, and measures the density profile around a tagged particle. The second is based on forces, as recently proposed in Ref. [102]. The identity between both methods is based on the assumption that the system is in thermal equilibrium. Therefore, any difference between the two expressions can be taken as the sign that the system is not equilibrated, but this approach has not been tested before.

Overall, we find that the relative difference between the two expressions for the pair correlation are extremely small, typically smaller than 1 %, see Fig. 13. A small systematic signal is observed when calculating the difference between both techniques. However, this signal depends on the discretization and binning of the histograms and is thus observable independently of the temperature. Apart from this signal, no systematic difference between the two computation methods can be observed. This method therefore cannot be used to detect non-equilibrium properties.

This result is reminiscent of similar findings for the configurational temperature, which is shown to decay instantaneously to the thermal temperature during equilibration [103]. The relationship between the histogram and the force methods for $g(r)$ corresponds roughly to a space-dependent generalization of the global configura-




FIG. 13. Relative difference between two expressions of the radial distribution function $g(r)$, calculated using either the traditional histogram method and the force method [102], for various temperatures. The small systematic differences are due to discretisation issues and do not depend on the degree of equilibration.

tional temperature.

Appendix C: Scaling with system size

In the main text, we concentrate on a small system size, $N = 43$, and we only briefly mention how the results may change with system size.

We repeated sampling with SMC, MD, PT, PA and NF for a larger system size, $N = 172$. As in the main text, we then use the SMC results as a benchmark and report in Fig. 14(a) deviations of the average energy with respect to SMC. In Fig. 14(b) we show results for the specific heat for the various sampling methods.

Compared to $N = 43$, the performance of the MD approach are essentially the same, with deviations appearing near $T_{\text{MD}} = 0.3$ in both quantities. However, we observe that the efficiency of the three enhanced sampling techniques (PT, PA, and NF) is significantly reduced in larger systems, as expected [20, 68]. In detail, we see that PT and PA now have a comparable performance, with a speedup compared to MD that is much less impressive than for $N = 43$ particles. This strong decrease in performance for both techniques stems from the complexity of sampling multiple low-energy states in glassy systems.

We also conclude that the normalizing flows suffer from the same reduction in performance with increasing system size. Therefore, our current implementation of NF does not get more efficient in larger systems compared to traditional enhanced sampling techniques such as PT and PA. Scaling the NF method to large sizes is clearly a challenging problem, which therefore deserves further attention in future work.




FIG. 14. Scaling the results to a larger system size, $N = 172$. (a) Difference in potential energy from the SMC result $\Delta E_{\text{pot}} = \langle E_{\text{pot}} \rangle - \langle E_{\text{pot}}^{\text{SMC}} \rangle$. (b) Specific heat c_V . The open symbols show the SMC results for $N = 43$, for comparison.

-
- [1] G. Battimelli, G. Battimelli, G. Ciccotti, P. Greco, and Scalone, *Computer Meets Theoretical Physics* (Springer, 2020).
- [2] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, Vol. 1 (Elsevier, 2001).
- [3] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids* (Oxford university press, 2017).
- [4] M. E. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics* (Clarendon Press, 1999).
- [5] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of State Calculations by Fast Computing Machines, *The Journal of Chemical Physics* **21**, 1087 (1953).
- [6] W. K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications* (Oxford University Press, 1970).
- [7] B. J. Alder, T. E. Wainwright, *et al.*, Phase transition for a hard sphere system, *The Journal of chemical physics* **27**, 1208 (1957).
- [8] A. Rahman and F. H. Stillinger, Molecular dynamics study of liquid water, *The Journal of Chemical Physics* **55**, 3336 (1971).
- [9] W. Krauth, *Statistical mechanics: algorithms and computations*, Vol. 13 (OUP Oxford, 2006).
- [10] D. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics* (Cambridge university press, 2021).
- [11] D. Gazzillo and G. Pastore, Equation of state for symmetric non-additive hard-sphere fluids: An approximate analytic expression and new monte carlo results, *Chemical Physics Letters* **159**, 388 (1989).
- [12] W. Kranendonk and D. Frenkel, Computer simulation of solid-liquid coexistence in binary hard sphere mixtures, *Molecular physics* **72**, 679 (1991).
- [13] T. S. Grigera and G. Parisi, Fast monte carlo algorithm for supercooled soft spheres, *Physical Review E* **63**, 045102 (2001).
- [14] M. Vucelja, Lifting—a nonreversible markov chain monte carlo algorithm, *American Journal of Physics* **84**, 958 (2016).
- [15] E. P. Bernard, W. Krauth, and D. B. Wilson, Event-chain monte carlo algorithms for hard-sphere systems, *Phys. Rev. E* **80**, 056704 (2009).
- [16] W. Krauth, Event-chain monte carlo: foundations, applications, and prospects, *Frontiers in Physics* **9**, 663457 (2021).
- [17] F. Ghimenti, L. Berthier, and F. van Wijland, Irreversible monte carlo algorithms for hard disk glasses: from event-chain to collective swaps, *arXiv preprint arXiv:2402.06585* (2024).
- [18] E. Marinari and G. Parisi, Simulated tempering: a new monte carlo scheme, *Europhysics letters* **19**, 451 (1992).
- [19] K. Hukushima and K. Nemoto, Exchange monte carlo method and application to spin glass simulations, *Journal of the Physical Society of Japan* **65**, 1604 (1996).
- [20] D. J. Earl and M. W. Deem, Parallel tempering: Theory, applications, and new perspectives, *Physical Chemistry Chemical Physics* **7**, 3910 (2005).
- [21] R. H. Swendsen and J.-S. Wang, Replica monte carlo simulation of spin-glasses, *Physical review letters* **57**, 2607 (1986).
- [22] K. Hukushima and Y. Iba, Population annealing and its application to a spin glass, in *AIP Conference Proceedings*, Vol. 690 (American Institute of Physics, 2003) pp. 200–206.
- [23] J. Machta, Population annealing with weighted averages: A monte carlo method for rough free-energy landscapes, *Physical Review E* **82**, 026704 (2010).
- [24] C. Amey and J. Machta, Analysis and optimization of population annealing, *Physical Review E* **97**, 033301 (2018).
- [25] F. Ghimenti, L. Berthier, G. Szamel, and F. van Wijland, Sampling efficiency of transverse forces in dense liquids, *Phys. Rev. Lett.* **131**, 257101 (2023).
- [26] L. Berthier and G. Biroli, Theoretical perspective on the glass transition and amorphous materials, *Reviews of modern physics* **83**, 587 (2011).
- [27] L. Berthier and D. R. Reichman, Modern computational studies of the glass transition, *Nature Reviews Physics* **5**, 102 (2023).
- [28] J.-L. Barrat and L. Berthier, Computer simulations of the glass transition and glassy materials, *Comptes Rendus. Physique* **24**, 1 (2023).
- [29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, *Advances in neural information processing systems* **27** (2014).
- [30] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in *International conference on machine learning* (PMLR, 2015) pp. 2256–2265.
- [31] J. Ho, A. Jain, and P. Abbeel, Denoising diffusion probabilistic models, *Advances in neural information processing systems* **33**, 6840 (2020).
- [32] D. P. Kingma and M. Welling, Auto-encoding variational bayes, *arXiv preprint arXiv:1312.6114* (2013).
- [33] D. Rezende and S. Mohamed, Variational inference with normalizing flows, in *International conference on machine learning* (PMLR, 2015) pp. 1530–1538.
- [34] F. Noé, S. Olsson, J. Köhler, and H. Wu, Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning, *Science* **365**, eaaw1147 (2019).
- [35] D. Wu, L. Wang, and P. Zhang, Solving statistical mechanics using variational autoregressive networks, *Physical review letters* **122**, 080602 (2019).
- [36] M. Invernizzi, A. Krämer, C. Clementi, and F. Noé, Skipping the replica exchange ladder with normalizing flows, *The Journal of Physical Chemistry Letters* **13**, 11643 (2022).
- [37] M. Gabrié, G. M. Rotskoff, and E. Vanden-Eijnden, Adaptive monte carlo augmented with normalizing flows, *Proceedings of the National Academy of Sciences* **119**, e2109420119 (2022).
- [38] S. Falkner, A. Coretti, S. Romano, P. L. Geissler, and C. Dellago, Conditioning boltzmann generators for rare event sampling, *Machine Learning: Science and Technology* **4**, 035050 (2023).
- [39] A. Coretti, S. Falkner, P. Geissler, and C. Dellago, Learning mappings between equilibrium states of liquid systems using normalizing flows (2022),

- arXiv:2208.10420.
- [40] S. van Leeuwen, A. P. de Alba Ortíz, and M. Dijkstra, A boltzmann generator for the isobaric-isothermal ensemble (2023), arXiv:2305.08483.
- [41] X. Ding and B. Zhang, Computing absolute free energy with deep generative models, *The Journal of Physical Chemistry B* **124**, 10166 (2020), pMID: 33143418.
- [42] P. Wirnsberger, B. Ibarz, and G. Papamakarios, Estimating gibbs free energies via isobaric-isothermal flows, *Machine Learning: Science and Technology* **4**, 035039 (2023).
- [43] T. Marchand, M. Ozawa, G. Biroli, and S. Mallat, Wavelet conditional renormalization group (2022), arXiv:2207.04941.
- [44] A. Singha, D. Chakrabarti, and V. Arora, Conditional normalizing flow for markov chain monte carlo sampling in the critical region of lattice field theory, *Physical Review D* **107**, 014512 (2023).
- [45] B. McNaughton, M. V. Milošević, A. Perali, and S. Pilati, Boosting monte carlo simulations of spin glasses using autoregressive neural networks, *Phys. Rev. E* **101**, 053312 (2020).
- [46] G. Scriva, E. Costa, B. McNaughton, and S. Pilati, Accelerating equilibrium spin-glass simulations using quantum annealers via generative deep learning, *SciPost Phys.* **15**, 018 (2023).
- [47] S. Ciarella, J. Trinquier, M. Weigt, and F. Zamponi, Machine-learning-assisted monte carlo fails at sampling computationally hard problems, *Machine Learning: Science and Technology* **4**, 010501 (2023).
- [48] D. Ghio, Y. Dandi, F. Krzakala, and L. Zdeborová, Sampling with flows, diffusion and autoregressive neural networks: A spin-glass perspective (2023), arXiv:2308.14085.
- [49] M. Mezard and A. Montanari, *Information, physics, and computation* (Oxford University Press, 2009).
- [50] P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro, and Z. Nussinov, Detection of hidden structures for arbitrary scales in complex physical systems, *Scientific reports* **2**, 329 (2012).
- [51] S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu, A structural approach to relaxation in glassy liquids, *Nature Physics* **12**, 469 (2016).
- [52] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. Nelson, T. Back, D. Hassabis, *et al.*, Unveiling the predictive power of static structure in glassy systems, *Nature physics* **16**, 448 (2020).
- [53] J. Paret, R. L. Jack, and D. Coslovich, Assessing the structural heterogeneity of supercooled liquids through community inference, *The Journal of chemical physics* **152** (2020).
- [54] G. Jung, R. M. Alkemade, V. Bapst, D. Coslovich, L. Filion, F. P. Landes, A. Liu, F. S. Pezzicoli, H. Shiba, G. Volpe, *et al.*, Roadmap on machine learning glassy liquids, arXiv preprint arXiv:2311.14752 (2023).
- [55] C. Scalliet, B. Guiselin, and L. Berthier, Thirty milliseconds in the life of a supercooled liquid, *Phys. Rev. X* **12**, 041028 (2022).
- [56] G. Jung, G. Biroli, and L. Berthier, Predicting dynamic heterogeneity in glass-forming liquids by physics-inspired machine learning, *Phys. Rev. Lett.* **130**, 238202 (2023).
- [57] G. Jung, G. Biroli, and L. Berthier, Dynamic heterogeneity at the experimental glass transition predicted by transferable machine learning, *Phys. Rev. B* **109**, 064205 (2024).
- [58] A. Ninarello, L. Berthier, and D. Coslovich, Models and algorithms for the next generation of glass transition studies, *Phys. Rev. X* **7**, 021039 (2017).
- [59] W. Kob and H. C. Andersen, Testing mode-coupling theory for a supercooled binary lennard-jones mixture i: The van hove correlation function, *Phys. Rev. E* **51**, 4626 (1995).
- [60] A. D. S. Parmar, M. Ozawa, and L. Berthier, Ultrastable metallic glasses in silico, *Phys. Rev. Lett.* **125**, 085505 (2020).
- [61] A. Heuer, Exploring the potential energy landscape of glass-forming systems: from inherent structures via metabasins to macroscopic transport, *Journal of Physics: Condensed Matter* **20**, 373101 (2008).
- [62] L. Berthier, D. Coslovich, A. Ninarello, and M. Ozawa, Equilibrium sampling of hard spheres up to the jamming density and beyond, *Physical review letters* **116**, 238002 (2016).
- [63] L. Berthier, E. Flenner, C. J. Fullerton, C. Scalliet, and M. Singh, Efficient swap algorithms for molecular dynamics simulations of equilibrium supercooled liquids, *Journal of Statistical Mechanics: Theory and Experiment* **2019**, 064004 (2019).
- [64] E. Flenner and G. Szamel, Hybrid monte carlo simulation of a glass-forming binary mixture, *Physical Review E* **73**, 061505 (2006).
- [65] Y. Sugita and Y. Okamoto, Replica-exchange molecular dynamics method for protein folding, *Chemical physics letters* **314**, 141 (1999).
- [66] G. Bussi, F. L. Gervasio, A. Laio, and M. Parrinello, Free-energy landscape for β hairpin folding from combined parallel tempering and metadynamics, *Journal of the American Chemical Society* **128**, 13435 (2006).
- [67] A. Bunker and B. Dünweg, Parallel excluded volume tempering for polymer melts, *Phys. Rev. E* **63**, 016701 (2000).
- [68] M. Falcioni and M. W. Deem, A biased monte carlo scheme for zeolite structure solution, *The Journal of chemical physics* **110**, 1754 (1999).
- [69] R. Yamamoto and W. Kob, Replica-exchange molecular dynamics simulation for supercooled liquids, *Phys. Rev. E* **61**, 5473 (2000).
- [70] C. De Michele and F. Sciortino, Equilibration times in numerical simulation of structural glasses: Comparing parallel tempering and conventional molecular dynamics, *Physical Review E* **65**, 051202 (2002).
- [71] S. Yaida, L. Berthier, P. Charbonneau, and G. Tarjus, Point-to-set lengths, local structure, and glassiness, *Physical Review E* **94**, 032605 (2016).
- [72] L. Berthier, P. Charbonneau, and S. Yaida, Efficient measurement of point-to-set correlations and overlap fluctuations in glass-forming liquids, *The Journal of chemical physics* **144** (2016).
- [73] W. Kob and L. Berthier, Probing a liquid to glass transition in equilibrium, *Phys. Rev. Lett.* **110**, 245702 (2013).
- [74] A. M. Ferrenberg and R. H. Swendsen, New monte carlo technique for studying phase transitions, *Phys. Rev. Lett.* **61**, 2635 (1988).
- [75] T. Shen and D. Hamelberg, A statistical analysis of the precision of reweighting-based simulations, *The Journal*

- of chemical physics **129** (2008).
- [76] Y. Miao, W. Sinko, L. Pierce, D. Bucher, R. C. Walker, and J. A. McCammon, Improved reweighting of accelerated molecular dynamics simulations for free energy calculation, *Journal of chemical theory and computation* **10**, 2677 (2014).
- [77] S. T. Tokdar and R. E. Kass, Importance sampling: a review, *Wiley Interdisciplinary Reviews: Computational Statistics* **2**, 54 (2010).
- [78] W. Wang, J. Machta, and H. G. Katzgraber, Population annealing: Theory and application in spin glasses, *Phys. Rev. E* **92**, 063307 (2015).
- [79] C. Amey and J. Machta, Analysis and optimization of population annealing, *Phys. Rev. E* **97**, 033301 (2018).
- [80] D. Gessert, W. Janke, and M. Weigel, Resampling schemes in population annealing—numerical and theoretical results, arXiv preprint arXiv:2305.19994 (2023).
- [81] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *Journal of Machine Learning Research* **22**, 1 (2021).
- [82] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density estimation using real nvp, arXiv preprint arXiv:1605.08803 (2016).
- [83] J. Song, S. Zhao, and S. Ermon, A-nice-mc: Adversarial training for mcmc, *Advances in neural information processing systems* **30** (2017).
- [84] L. Klein, A. Foong, T. Fjelde, B. Mlodozieniec, M. Brockschmidt, S. Nowozin, F. Noé, and R. Tomioka, Timewarp: Transferable acceleration of molecular dynamics by learning time-coarsened dynamics, *Advances in Neural Information Processing Systems* **36** (2024).
- [85] M. S. Albergo, D. Boyda, D. C. Hackett, G. Kanwar, K. Cranmer, S. Racaniere, D. J. Rezende, and P. E. Shanahan, Introduction to normalizing flows for lattice field theory, arXiv preprint arXiv:2101.08176 (2021).
- [86] J. Köhler, L. Klein, and F. Noé, Equivariant flows: exact likelihood generative learning for symmetric densities, in *International conference on machine learning* (PMLR, 2020) pp. 5361–5370.
- [87] Usually this transformation is defined as T , see Ref. [86], which we avoid due to the importance of the temperature T in the present study.
- [88] L. I. Midgley, V. Stimper, G. N. C. Simm, B. Schölkopf, and J. M. Hernández-Lobato, Flow annealed importance sampling bootstrap, in *The Eleventh International Conference on Learning Representations* (2023).
- [89] L. Kish, *Survey sampling* (Wiley, 1965).
- [90] C. Amey and J. Machta, Measuring glass entropies with population annealing, arXiv preprint arXiv:2103.13837 (2021).
- [91] L. I. Midgley, V. Stimper, J. Antorán, E. Mathieu, B. Schölkopf, and J. M. Hernández-Lobato, Se(3) equivariant augmented coupling flows (2023), arXiv:2308.10364.
- [92] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, Flow matching for generative modeling (2023), arXiv:2210.02747.
- [93] L. Klein, A. Krämer, and F. Noé, Equivariant flow matching (2023), arXiv:2306.15030.
- [94] L. Felardos, J. Hénin, and G. Charpiat, Designing losses for data-free training of normalizing flows on boltzmann distributions (2023), arXiv:2301.05475.
- [95] M. Arbel, A. Matthews, and A. Doucet, Annealed flow transport monte carlo, in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 139, edited by M. Meila and T. Zhang (PMLR, 2021) pp. 318–330.
- [96] M. Xu, L. Yu, Y. Song, C. Shi, S. Ermon, and J. Tang, Geodiff: a geometric diffusion model for molecular conformation generation (2022), arXiv:2203.02923.
- [97] S. Zheng, J. He, C. Liu, Y. Shi, Z. Lu, W. Feng, F. Ju, J. Wang, J. Zhu, Y. Min, H. Zhang, S. Tang, H. Hao, P. Jin, C. Chen, F. Noé, H. Liu, and T.-Y. Liu, Towards predicting equilibrium distributions for molecular systems with deep learning (2023), arXiv:2306.05445.
- [98] D. Shu, Z. Li, and A. Barati Farimani, A physics-informed diffusion model for high-fidelity flow field reconstruction, *Journal of Computational Physics* **478**, 111972 (2023).
- [99] H. Wu, J. Köhler, and F. Noe, Stochastic normalizing flows, in *Advances in Neural Information Processing Systems*, Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Curran Associates, Inc., 2020) pp. 5933–5944.
- [100] H. Sun and K. L. Bouman, Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging, *Proceedings of the AAAI Conference on Artificial Intelligence* **35**, 2628 (2021).
- [101] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, beta-VAE: Learning basic visual concepts with a constrained variational framework, in *International Conference on Learning Representations* (2017).
- [102] B. Rotenberg, Use the force! Reduced variance estimators for densities, radial distribution functions, and local mobilities in molecular simulations, *The Journal of Chemical Physics* **153**, 150902 (2020).
- [103] S. Mehri, T. S. Ingebrigtsen, and J. C. Dyre, Single-parameter aging in a binary Lennard-Jones system, *The Journal of Chemical Physics* **154**, 094504 (2021).