



From data to noise to data for mixing physics across temperatures with generative artificial intelligence

Yihang Wang^{a,b}, Lukas Herron^{a,b}, and Pratyush Tiwary^{b,c,1}

Edited by Sharon Glotzer, University of Michigan, Ann Arbor, MI; received March 2, 2022; accepted July 7, 2022

Using simulations or experiments performed at some set of temperatures to learn about the physics or chemistry at some other arbitrary temperature is a problem of immense practical and theoretical relevance. Here we develop a framework based on statistical mechanics and generative artificial intelligence that allows solving this problem. Specifically, we work with denoising diffusion probabilistic models and show how these models in combination with replica exchange molecular dynamics achieve superior sampling of the biomolecular energy landscape at temperatures that were never simulated without assuming any particular slow degrees of freedom. The key idea is to treat the temperature as a fluctuating random variable and not a control parameter as is usually done. This allows us to directly sample from the joint probability distribution in configuration and temperature space. The results here are demonstrated for a chirally symmetric peptide and single-strand RNA undergoing conformational transitions in all-atom water. We demonstrate how we can discover transition states and metastable states that were previously unseen at the temperature of interest and even bypass the need to perform further simulations for a wide range of temperatures. At the same time, any unphysical states are easily identifiable through very low Boltzmann weights. The procedure while shown here for a class of molecular simulations should be more generally applicable to mixing information across simulations and experiments with varying control parameters.

molecular simulations | generative artificial intelligence | enhanced sampling

How does one learn physics and chemistry at a certain temperature on the basis of experiment or experiments performed at certain other temperatures? The Arrhenius equation (1) provides the simplest way to do it, if one is willing to make significant simplifying assumptions about the energy landscape of the system. Often arbitrarily complex systems indeed display kinetics conforming to the Arrhenius picture (2), but frequently one also observes its violations (3). Even when the Arrhenius picture is applicable, it allows only extrapolation of a single number, i.e., reaction rate across the temperatures. It thus remains desirable to develop techniques that allow inferring much more detailed information across a temperature range. These could be thermodynamic observables such as relative probabilities of a molecule's different conformations or of a crystal's different polymorphs or more detailed kinetic observables than just an overall rate constant. Even more generally, given observations of the positions and velocities of all constituents of an N -body system at some set of temperatures, we seek to estimate what they will be at a temperature where the experiment or simulation was never performed. The problem is hard because the probability distribution connecting configurational coordinates across temperatures is hard to sample from, especially since N is extremely large for most systems of interest, equaling at least tens of thousands.

Here we propose a framework using generative artificial intelligence (AI) that learns to efficiently sample such high-dimensional, complex probability distributions for N -body molecular systems valid across temperatures. There are two central ideas guiding our framework. First, we do not treat the temperature T as just a control parameter. By noting that the temperature is a measure of the average kinetic energy, we instead work directly with the fluctuating kinetic energy, using the equipartition theorem to define an instantaneous, effective temperature that we call $\mathcal{T}$. For finite N not yet in the thermodynamic limit, $\mathcal{T}$ so defined will display significant relative fluctuations proportional to $1/\sqrt{N}$. Second, if we have experiments or simulations performed with heat bath temperatures $T_1, T_2, \dots, T_K$, we can view them all together as being sampled from the same, but unknown, joint probability $p(\mathbf{x}, \mathcal{T})$ for N -body configurations $\mathbf{x}$. We then use a generative AI method, specifically denoising diffusion probabilistic models (DDPMs) (4, 5) to generate many more samples from $p(\mathbf{x}, \mathcal{T})$, given the sparse, high-dimensional dataset.

To learn a $p(\mathbf{x}, \mathcal{T})$ from such a high-dimensional and sparse dataset, we use the DDPM (4, 5), a generative AI framework inspired by nonequilibrium dynamics to approximate

Significance

While it is tempting to use high-temperature simulations to infer observations about low temperature, it is not always clear how to do so. Here we demonstrate how using generative artificial intelligence we can mix information from simulations conducted at a set of temperatures and generate molecular configurations at any temperature of interest including temperatures at which simulations were never performed. The configurations we generate carry correct Boltzmann weights, and our model minimizes the generation of spurious unphysical configurations. We demonstrate its use here through combining with replica exchange molecular dynamics in a postprocessing framework for sampling peptide and ribonucleic acid. We believe the framework is extensible to generic simulations and experiments for mixing control parameters other than temperature.

Author affiliations: ^aBiophysics Program, University of Maryland, College Park, MD 20742; ^bInstitute for Physical Science and Technology, University of Maryland, College Park, MD 20742; and ^cDepartment of Chemistry and Biochemistry, University of Maryland, College Park, MD 20742

Author contributions: Y.W. and P.T. designed research; Y.W. and L.H. performed research; and Y.W., L.H., and P.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: ptiwary@umd.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2203656119/-DCSupplemental>.

Published August 4, 2022.

and sample from very high-dimensional distributions. By learning $p(\mathbf{x}, T)$ we are referring to the ability to generate many more samples from $p(\mathbf{x}, T)$ given the dataset we have from simulations or experiments performed so far. The DDPM has been shown to possess the ability to infer and learn the underlying relationship from complicated data, i.e., high dimensional and with complex correlations, but also noisy and sparse (4, 5). DDPMs learn two diffusion processes called noising and denoising. The forward diffusion or noising process converts the samples from high-dimensional, structured, and unknown probability distribution into simpler and analytically tractable white noise. The backward diffusion or denoising process learns the mapping back from noise to meaningful data. The central idea is that it is easy to generate numerous samples from the noisy distribution, which can then be denoised back to structured data. The process has been demonstrated to be comparable to, or even outperform, other generative AI models for generating high-quality samples and in its ability to model the underlying semantics behind meaningful data in an unsupervised manner (4, 5).

Despite their promising potential, to the best of our knowledge DDPMs have not yet been used in the context of mixing information across processing conditions as done here. We are able to achieve this through recognizing that for finite-size systems, as one has in molecular simulations, the temperature can be associated with a random variable that has significant fluctuations rather than just a control parameter. The protocol developed here should be applicable quite generally to data coming from simulations or experiments at different temperatures. It should also be extensible to mixing data from control parameters other than temperatures—possibly including concentration, pressure, and volume. Here we focus on a specific class of molecular simulations known as replica exchange molecular dynamics (REMD) (6, 7) that use information generated at a ladder of temperatures to swap configurations across temperatures. REMD has been extremely powerful over the decades for the study of molecular systems with rough energy landscapes for fundamental science and practical applications (8, 9). Numerous advances have been introduced over this basic idea to make it more efficient computationally (10–17) and it continues to be an area of very active research. We demonstrate how the DDPM applied to all-atom, femtosecond resolution REMD simulations of a small peptide and an RNA nucleotide very significantly improves the quality of data generated across temperatures, including temperatures at which the simulation was never performed and even at temperatures outside the ladder of temperatures. DDPM significantly improves the estimates of free energies made through REMD and provides accurate sampling of configurations that were not previously visited in the lowest-temperature replica during REMD. This could be metastable states or transition states. We also show how this can be used to generate samples at temperatures not included in the ladder of replicas. The samples we generate have thermodynamic relevance and correspond to correct Boltzmann weights as opposed to being just wild hallucinations (5, 18). Due to its generalizable framework and simple postprocessing nature of application, we thus believe that this work should be extensible to a wide range of studies to supplement and extend the range of simulations and experiments without actually performing them at all conditions of interest.

Results

Defining a Fluctuating Effective Temperature. We work with the equivalent $p_T(\mathbf{x})$ where T is the thermodynamics temperature. Molecular dynamics (MD) or Monte Carlo (MC) methods

allow sampling configurations $\mathbf{x}$ as per the equilibrium probability $p_T(\mathbf{x}) \equiv e^{-\frac{1}{k_B T} U(\mathbf{x})} / Z$, where $U(\mathbf{x})$ is the potential energy of the N -body system and $Z = \int d\mathbf{x} e^{-\frac{1}{k_B T} U(\mathbf{x})}$ is the partition function. For systems of practical interest in biology, chemistry, and materials science, if T is not large enough, it becomes nearly impossible to sample reliably from $p_T(\mathbf{x})$, as many regions of interest in configuration space such as transition states will have extremely low probability compared with low free-energy states. To deal with this problem, in REMD one simulates $K + 1$ replicas of the system at temperatures $T = T_0 < T_1 < \dots < T_K$. For high enough T_K , the sampling from $p_{T_K}(\mathbf{x}) \equiv e^{-\frac{1}{k_B T_K} U(\mathbf{x})} / Z_K$ is expected to be more ergodic. One then periodically exchanges conformations between consecutive pairs of replicas with a Metropolis-type acceptance probability that depends on the potential energies of the two replicas and their temperatures. This way even the low-temperature T_0 replica can explore configurations that would not have been visited otherwise.

We now argue that by treating the temperature just as a control parameter REMD is not making full use of the information gathered across the ladder of temperatures. In most current incarnations of REMD (excluding exceptions such as ref. 19), all that the higher temperatures do is to help the lower-temperature replicas discontinuously appear in different locations of the configuration space. We take a slightly different view of the temperature in REMD in this article, but this idea should be easily generalizable to other thermodynamic variables that show fluctuations for finite system size. Therefore, we are able to treat the approximate, effective temperature as a random variable rather than a control parameter. Additionally, the effective temperature of the system is associated with its kinetic energy instead of the temperature of the heat bath that we expect the thermostat to enforce on average. More rigorously, we are sampling the joint $p(\mathbf{x}, T)$, where the instantaneous temperature T is associated with the per-particle kinetic energy κ through $\frac{3}{2} k_B T = \kappa$ and is related to the thermodynamics or heat bath temperature through its ensemble average; i.e., $\langle T \rangle = T$. The instantaneous temperature can be calculated at each frame of a simulation.

Our motivation in doing so is that all replicas across different temperatures can be viewed as being sampled from the same joint probability $p(\mathbf{x}, T)$ —as opposed to different replicas sampling from their respective $p_T(\mathbf{x})$. This change of perspective allows us to combine together the data collected from different temperatures as having arisen from the same, although intractable, probability distribution $p(\mathbf{x}, T)$.

Challenge in Sampling the Intractable Joint Probability across Configurations and Temperatures. Our task at hand now is to learn the joint probability $p(\mathbf{x}, T)$ given sampling that has already been performed through REMD across $\mathbf{x}$ space and temperatures $T = T_0 < T_1 < \dots < T_K$. There are two main challenges in this. The first challenge is due to the curse of dimensionality: The memory or computational resources needed to track a very high-dimensional distribution function increase exponentially with the number of degrees of freedom. For instance, for REMD of a small nine-residue peptide in explicit water, which we study here, we exchange all 749 atomic coordinates between the replicas, but for the purpose of analysis we set $\mathbf{x}$ as 18 Ramachandran dihedral angles. This means we already have a 19-dimensional space where binning procedures are out of the question. The second challenge comes from the sparsity of the data. Most of our samples come from high-probability regions of $p(\mathbf{x}, T)$ and we have very few samples of low-probability states or even high-probability states at low temperature due to inefficient exchange between replicas.




Fig. 1. (Lower panel) The neural network architecture used in this work. It has the basic structure of the U-Net model (20). During the diffusion process, which is indicated by the direction of blue arrows in *Upper* panel, noise is gradually added to the sampled data, in this case the picture of a very good dog, through a diffusion process labeled with diffusion step t_i . This changes the sampled distribution (for example, more pictures of dogs) to a simpler isotropic Gaussian distribution from which one can easily generate more samples. An AI model is then trained to reverse such a diffusion process and starting from sampled noise, to learn to generate images similar to the input image by following the direction indicated by the orange arrows. It takes a one-dimensional (1-d) array $\mathbf{s}'$, which is the noisy sample at diffusion step t_i , as input and outputs the parameterization of a Gaussian distribution to get the reversed-transition kernel $q_\theta(\mathbf{s}, t_{i+1} | \mathbf{s}', t_i)$. Each residue block consists of three components: two 1-d convolution operations with kernel size 3 and a group normalization (21) between them. The diffusion step t_i is added to each convolutional block after being transformed by the sinusoidal position embedding (22). The final conv label denotes the convolution operation with kernel size equal to 1. Max-pooling reduces size of the features by half, while up-sampling uses the transposed convolution to expand the size of features.

In summary, we have sparse sampling of data points in very high-dimensional $(\mathbf{x}, \mathcal{T})$ space and wish to construct $p(\mathbf{x}, \mathcal{T})$ from this information so that we can create many more samples at any temperature of interest.

DDPMs Can Generate Many More Samples from $p(\mathbf{x}, \mathcal{T})$. The main idea behind the use of DDPM here is to learn a simple and easy-to-sample-from distribution $q_\theta(\mathbf{x}, \mathcal{T})$ that approximates the true $p(\mathbf{x}, \mathcal{T})$. For notational simplicity, we denote $\mathbf{s} = \{\mathbf{x}, \mathcal{T}\}$ and refer to $q_\theta(\mathbf{s})$ and $p(\mathbf{s})$ henceforth. DDPM does this by learning to reverse a gradual, multistep noising process that starts with the relatively limited number of samples generated from the distribution $p(\mathbf{s})$ and diffuses to the simpler distribution $p_{\text{simple}}(\mathbf{s})$ that is easy to sample. For instance, $p_{\text{simple}}(\mathbf{s})$ could be an isotropic Gaussian. In addition to learning this noising process, DDPMs also learn the reverse denoising process that allows us to go back from samples generated using $p_{\text{simple}}(\mathbf{s})$ to samples that would have been generated from the underlying $p(\mathbf{s})$. Both the noising and denoising processes are modeled using diffusion processes that convert probability distributions to one another and are implemented using the architecture shown in Fig. 1, which is based on the standard architecture for DDPMs as described in ref. 5.

Noising and Denoising Processes. The noising diffusion process carried out in the space $\mathbf{s}$ that converts $p(\mathbf{s})$ to the simpler $p_{\text{simple}}(\mathbf{s})$ can be decomposed into M discrete steps denoted by corresponding transition probabilities $p(\mathbf{s}', t_{i+1} | \mathbf{s}, t_i)$, where $i \in [0, M]$, $p(\mathbf{s}, t_0) \equiv p(\mathbf{x}, \mathcal{T})$, and $p(\mathbf{s}, t_M) \equiv p_{\text{simple}}(\mathbf{x}, \mathcal{T})$. In DDPM, this noising diffusion process that converts sampled data to essentially noise is set to be an Ornstein–Uhlenbeck (OU) process in which the transition probability follows a simple Gaussian form. One can then easily generate samples from $p(\mathbf{s}, t_0) \equiv p(\mathbf{x}, \mathcal{T})$. The tricky bit is now to convert these samples back to the original distribution. In ref. 4, it was shown that this transition-reversed diffusion kernel $p(\mathbf{s}, t_i | \mathbf{s}', t_{i+1})$ can also be written in a Gaussian form. A deep neural network (Fig. 1) is then trained using variational inference to learn the approximate reversed transition kernel $q_\theta(\mathbf{s}, t_i | \mathbf{s}', t_{i+1}) \approx p(\mathbf{s}, t_i | \mathbf{s}', t_{i+1})$. Thus by generating samples from a normal Gaussian distribution, which

we can easily do in large numbers, and then passing these through the reversed-transition kernel, we can generate samples that follow the target distribution $p(\mathbf{x}, \mathcal{T})$ as desired. Note that instead of learning the joint probabilities $p(\mathbf{s}_1, \mathbf{s}_2)$, it can be advantageous to learn the conditional probability $p(\mathbf{s}_1 | \mathbf{s}_2)$. This can be done through the protocol in ref. 4 by adding a delta function to allow only a subset of $\mathbf{s}$ to change during the noising and denoising process, i.e., $\delta(\mathbf{s}_2 - \mathbf{s}_2') p(\mathbf{s}_1, \mathbf{s}_2, t_i | \mathbf{s}_1', \mathbf{s}_2', t_{i+1})$. This is very useful when for instance we are interested in generating samples only at a certain temperature or only in certain regions of the configuration space. In the most general form of diffusion probabilistic models (DPMs) (4), the reversed-transition kernel $q_\theta(\mathbf{s}, t_i | \mathbf{s}', t_{i+1})$ is considered as $q_\theta(\mathbf{s}, t - 1 | \mathbf{s}', t) = \mathcal{N}(\mathbf{s}; \hat{\boldsymbol{\mu}}_\theta(\mathbf{s}', t), \hat{\boldsymbol{\sigma}}_\theta(\mathbf{s}', t))$ and the neural network is trained to learn the mean $\hat{\boldsymbol{\mu}}_\theta(\mathbf{s}', t)$ and variance $\hat{\boldsymbol{\sigma}}_\theta(\mathbf{s}', t)$. In practice, however, there are many different ways to choose the Gaussian distribution parameterization. In ref. 5, DDPM was introduced with a different parameterization approach to reduce the complexity of the training task. DDPM got its name because such a design makes the learning task resemble a denoising score-matching procedure (23). In ref. 5, it was shown that with such a design, DDPM can generate samples of a quality that is comparable to or even better than that in other generative models.

DDPM Applied to Replica Exchange MD: Desirables. We now demonstrate how the above protocol can be applied to mix data collected from different temperatures and configurations in REMD and significantly improve the quality of sampling. Specifically, we consider the following two challenging tasks:

- 1) Can we improve the sampling quality for the lowest-temperature replica with more accurate probability estimates than directly seen after REMD? This includes being able to generate samples in low-probability regions such as transition and metastable states and reliably estimate their free energies.
- 2) Can we generate samples at temperatures that are not included in the replica ladder? This would include temperatures within the range of the replica ladder and also extrapolation to temperatures outside the range.




Fig. 2. (A) AIB₉ in left-handed (L) and right-handed (R) conformations in explicit TIP3P water. (B) Free-energy profile of residues (with residue 5 as an example) with two stable states labeled as L, R and two excited state labeled as La, Ra. When all residues are in L states, the peptide chain is in the left-handed state and similarly for the right-handed state. The samples used to estimate this free energy are generated by 4- μ s-long MD simulation at 400 K. Free energies in B here as well as throughout the article and *SI Appendix* are provided in units of $k_B T$.

Peptide Conformational Transitions. To demonstrate the performance of DDPM, we first study a small peptide chain Aib₉ in explicit water with transferable intermolecular potential with 3 points potential (TIP3P water) (24) using the CHARMM36m force field (25). This nine-residue system (Fig. 2A) displays rich and complex conformational dynamics (26) including the transitions between fully left-handed and right-handed helices. However, even in 4- μ s unbiased MD at 400 K, one can see only around two to three transitions between these two dominant equiprobable conformations and even fewer, if any, transitions to the higher-energy metastable states. To improve the sampling of this system, we perform 100-ns REMD with 10 replicas at geometrically spaced temperatures ranging from 400 to 518 K, with temperature increased by 3% for each replica. The attempt to exchange configurations was made every 20 ps, with acceptance rate around 1% $\sim$ 2% between neighboring replicas, which is intentionally kept lower than what one usually has in REMD. This is because we want to show that even in the extreme cases where the number of atoms N is so large that replicas do not have enough overlap, or if one wants to reduce the number of replicas to save computational resources, our DDPM can still do a decent job of complementing REMD and reconstructing the true probability distribution at any temperature of interest. To benchmark our results, we ran unbiased MD at 400 K for 4 μ s and at 500 K for 0.6 μ s. These simulations are long enough such that there are multiple transitions between the left-handed state and the right-handed state. Therefore, we can use the free-energy profiles estimated from these unbiased simulations to benchmark the performance of DDPM.

To quantify the quality of sampling, we focus on the 18 dihedral angles corresponding to all nine residues ($\Phi_1, \Psi_1, \Phi_2, \Psi_2, \dots, \Phi_9, \Psi_9$). As shown in the Ramachandran plot in Fig. 2B, the free-energy surface along any pair of dihedrals is mainly characterized by four metastable states: two equiprobable low-energy ground states (L, R) and two excited states (La, Ra). The Ramachandran plots for all nine residues in the system look qualitatively similar but the middle residues of the peptide are known to be less flexible with higher-energy barriers (27, 28). We thus focus on the sampling for residues 5 and 8 as shown in Fig. 3. We train our DDPM on the 100-ns-long REMD trajectories. At the lowest temperature of interest (400 K), this trajectory has not yet achieved sufficient sampling. As shown in Fig. 3, DDPM successfully mixes information from all temperatures and

configurations and generates samples in states that are not present in the training dataset for both residues, indicated with thick black arrows in Fig. 3, *Right*. Specifically, for residue 5 we can see that the transition states between state R and state La, which were not being sampled in the 400-K replica, are populated in samples from DDPM. For residue 8 the improvement is even more striking as the state Ra that was simply not sampled in REMD now gets populated after DDPM. To further quantify the improvement gained due to DDPM, we compare the free-energy differences between different configurations from both REMD and DDPM against much longer reference unbiased MD at 400 K. As shown in *SI Appendix, Fig. S2*, the free-energy differences calculated from samples generated by DDPM are much closer to that from the reference MD. Thus, DDPMs are able to accomplish the first task highlighted above. We also highlight that while some spurious samples are generated, seen through dots outside the free-energy contours in Fig. 3, their Boltzmann weights are very low. Thus, our model “dreams” additional configurations without hallucinating spurious configurations.

We now move to the second task from the list above and test DDPM’s ability to generate samples by interpolating or even extrapolating across temperatures not considered in the ladder of replicas. In the first example, we use it to generate samples at 500 K, as in the training set with 10 replicas at geometrically spaced temperatures between 400 and 518 K, there is no replica with temperature 500 K. As shown in *SI Appendix, Fig. S2B*, the ΔG calculated from samples of DDPM is in good agreement with that from the reference MD at 500 K. In the second example, we completely remove the samples from the 400-K replica in the training set and use it to train a new model, which is then used to generate samples at 400 K. Even though in the training set, the lowest temperature is 412 K, the model can make accurate prediction of the free-energy difference between states as shown in Fig. 4. We are particularly encouraged by this last finding as sampling at lower temperatures generally tends to be more difficult than at higher temperatures.

To further push the limit of this model, we also investigated the extreme case of when there is no exchange of configurations between replicas. We ran the simulations with the same temperature ladder and simulation steps as in the previous experiment, but without any exchange of configurations between replicas. We trained the DDPM model with these trajectories and repeated the previous tests. As shown in *SI Appendix, Figs. S4–S9*, even




Fig. 3. Samples (dots) from REMD (Left) and DDPM (Right) at 400 K. The Boltzmann weights for different samples are indicated through their free energy (contour lines, separated every $0.74 k_b T$). (A) Samples generated from DDPM and free-energy profile projected on dihedral angles of residue 5. (B) Samples generated from DDPM and free-energy profile projected on dihedral angles of residue 8. DDPM was able to generate samples in states that are not present in the training dataset for both residues, indicated with thick black arrows in A and B, Right.

without any exchange of configurations between replicas, DDPM still successfully generated samples of metastable states that are not sampled through parallel unbiased simulation. Furthermore, by comparing the free-energy profiles calculated from DDPM samples and estimated by the multistate Bennett acceptance ratio estimator MBAR (29), we observed that both methods yield similar free-energy profiles when the target temperature is within the ladder of simulated temperatures. However, DDPM significantly outperforms MBAR in estimating free-energy profiles at temperatures that are not included in the MD simulations (see *SI Appendix* for more details.)

RNA Conformational Transitions. As a second example to illustrate the general applicability of our DDPM+REMD approach,

we turn our attention toward sampling an RNA conformational ensemble. Rare RNA structures have been previously shown to be biologically relevant (30–32), but estimating the conformational ensemble still remains computationally intractable using traditional sampling techniques (33). RNA dynamics may occupy a wide range of timescales—from several hours for conformational changes that require breaking base pairs to picoseconds for more continuous deformations (32). As a consequence, identifying rare transient structures and estimating their contribution to the RNA ensemble have proved to be difficult. In this example, we consider a GACC tetranucleotide as our model system (Fig. 5A). As a single-stranded RNA consisting of four nucleotides labeled G1, A2, C3, and C4, GACC has been previously used as a challenging test system for REMD-based sampling methods for its conforma-




Fig. 4. Free-energy differences between different metastable states ΔG calculated from samples from DDPM by extrapolating to a temperature at 400 K, which is lower than any temperature included in the dataset. The DDPM model was trained with REMD trajectories with nine replicas starting at the lowest temperature 412 K. Green open crosses and purple open triangles are reference free-energy differences ΔG from 4- μ s-long unbiased MD at 400 K. Green solid plus signs and purple diamonds show the energy differences from DDPM. ΔG_{R-La} is the free-energy difference between ground-state L ($-2.2 < \Phi < -0.1$, $-1.6 < \Psi < 0.9$) and excited-state Ra ($-1.6 < \Phi < -0.2$, $1.4 < \Psi < 3.14$); and $\Delta G_{L-\gamma S}$ is the free-energy difference between the ground-state L ($-2.2 < \Phi < -0.1$, $-1.6 < \Psi < 0.9$) and a transition state ($-1.8 < \Phi < -0.8$, $0.8 < \Psi < 1.5$).




Fig. 5. Projection of samples and free-energy profile on dihedral angles ζ and γ of A2 in GACC at 325 K. The Boltzmann weights for different samples are indicated through their free energy (contour lines, separated every $0.74 k_B T$). (A) The structure of GACC. (B) Samples from REMD. (C) Samples from the benchmark $5\mu\text{s}$ -long unbiased MD. (D) Samples from DDPM. The metastable states that are not present in the training dataset are indicated with thick black arrows.

tional flexibility (34). Despite the fact that GACC has been widely studied, it still remains challenging to effectively sample possible configurations and is a good system to test new methods (33). For example, in a previous study, to get the converged structural populations, a multidimensional replica exchange molecular dynamics (M-REMD) simulation was performed with 192 replicas with around $1\mu\text{s}$ of simulation time per replica, thus totaling almost 192 ms of all-atom simulations (34). Here we show that with DDPM, we can better estimate the free-energy landscape using fewer computational resources, totaling only 12 ms of all-atom simulations adding all replicas.

We trained our DDPM model on 250-ns REMD trajectories from 48 replicas with temperatures ranging from 277 to 408 K (see *MD Simulations Setup* for details). In each frame of the trajectory, the structure of GACC is characterized by six dihedral angles for each nucleotide. In these REMD simulations, the sampling is not sufficient, especially for replicas at very low temperatures. Fig. 5B and *SI Appendix, Figs. S4 and S5* show the free-energy profiles of GACC projected on the γ and ζ angles of A2 and C3 at different temperatures. We can see that expected high-energy metastable states indicated by a black arrow in Fig. 5B were not sampled in REMD. In contrast, for different target temperatures, DDPM successfully generates the ensemble of such high free-energy states that were never visited at low temperatures. Here as well, any spurious configurations can be seen through dots outside the free-energy contours in Fig. 5 with negligible Boltzmann weights.

Discussion

We have presented a generative AI-based approach that combines physics from simulations performed at different temperatures to generate reliable molecular configurations and accurate ther-

modynamic estimates at any arbitrary temperature even if no actual simulation was performed there. The central idea is to not work with the thermodynamic temperature T of the system as a parameter set by the heat bath, but instead work with an effective temperature, calculated from the instantaneous kinetic energy of the system. This effective temperature shows nontrivial fluctuations for a finite-size system and on an average equals the thermodynamic temperature. Given sparse sampling from the high-dimensional space comprising configurational space coordinates and the effective temperature, we train a generative AI model that generates countless more samples of configurations at any temperature of interest. Here we demonstrate its usefulness in the context of the widely used REMD framework to improve the sampling of REMD through a postprocessing framework. We show how this significantly improves the quality of sampling at low temperatures and even generates samples of states that have not been visited in the replicas and at temperatures not considered in the ladder of replicas. Since the kinetic energy fluctuations are an important component of this method, we believe DDPM will benefit from algorithms that precisely simulate these fluctuations, such as the Gaussian moment thermostat (35), or enhanced sampling methods that can generate more continuous temperature distributions, such as the extended Hamiltonian approach proposed in ref. 36. On the other hand, based on our experiments on systems simulated with different thermostats and under the condition of extremely low exchange rates, we can see the advantage of this method learning from data with limited fluctuations and its potential for application to a very wide range of simulation scenarios, including those with limited thermal fluctuations. It is worth mentioning here that a recent application (37) of normalizing flows also attempts to enhance REMD sampling through somewhat similar ideas to ours. However, in that work the machinery is used to directly affect the acceptance protocol

in REMD while ours is a purely postprocessing scheme. Our work also bears similarity to the T-WHAM approach of ref. 19 and the MBAR approach of ref. 29, which intend to estimate a similar joint probability over configuration and energy space. However, as shown in *SI Appendix, Figs. S4–S9*, compared with MBAR, numerically our approach is far superior in extrapolating free energy at temperatures outside the ladder of simulated temperatures.

The generative AI framework of DDPM used here belongs to the broad class of flow-type methods, which have been shown to have the ability to generate samples from high-dimensional space with many interdependent degrees of freedom. Compared with other flow-type models such as normalizing flow (38, 39) that use deterministic functions to map from an easy-to-sample distribution to target distribution, the stochastic nature of DDPM avoids the restriction of preserving the topology of configuration space and thus allows the learning of significantly more complicated distributions. At the same time the design of the transition kernels in DDPM reduces the learning task to just learning means of Gaussian kernels. This makes the training easier compared to other methods (40) while at the same time being able to learn more complicated transition kernels.

We also believe that our generative AI model, while “dreaming” thermodynamically relevant structures at different temperatures, avoids the so-called hallucinations suffered by other generative AI models; i.e., we do not generate meaningless, unphysical structures with significant thermodynamic weights (18). We believe this is through the use of relatively simple transition kernels, which avoids overparameterization of the model, and through utilizing molecular basis functions instead of all-atom coordinates, which reduces the space that needs to be sampled. The issue of generating out-of-distribution samples that has been problematic in other methods attempting to generate molecular structure, such as the Boltzmann generator, is usually avoided by discarding the translational and rotational degrees of freedom and reweighting the samples (39, 40). However, in the Boltzmann generator, the reweighting factors for each sample are functions of samples’ potential energy, which is well defined only when all coordinates of a system are given. Therefore, in principle, the Boltzmann generator can only be applied to samples generated in the full configuration space. In the case that a system contains explicit water molecules, the configuration space of the samples will again become too sparse to be effectively sampled. For a subset of the degrees of freedom of the system where the potential is not known, normalizing flows are a category of methods that may also be used to learn the probability distribution of states. Nevertheless, the method of normalizing flows, which also underlies the Boltzmann generator, has been observed to have limited expressiveness compared to stochastic methods such as DDPM (38, 40).

We finally point out that this work shows the possibility of learning generative models in the space of generic thermodynamic ensembles, by following the simple recipe that control parameters can also be viewed as fluctuating variables. As long as one is not in the thermodynamic limit—something we do not have to worry about in molecular simulations—this should be thus a practical and useful procedure for problems far beyond replica exchange molecular dynamics.

MD Simulations Setup.

AIB₉. The simulation of AIB₉ was set up by following a previous study (28). The Protein Data Bank (PDB) file was taken from the authors with permission, and the simulations are done with the CHARMM36m all-atom force field (25) using TIP3P water

molecules (24), a Parrinello–Rahman barostat (41), and a Nose–Hoover thermostat (42, 43) under the constant number, pressure, and temperature (NPT) ensemble. Simulations were performed using GROMACS 2016 (44). The structures of AIB₉ were saved every 0.2 ps for all simulations and the dihedral angles were calculated using PLUMED 2.4 (45). These datasets with 0.2-ps temporal resolution were used for both AI model training and free energy estimation.

GACC. The simulations of GACC were done with the AMBER ff12 all-atom force field, using TIP3P water molecules (24), a Parrinello–Rahman barostat (41), and a Bussi–Parrinello velocity rescaling thermostat (46) under the NPT ensemble. The simulations were performed using GROMACS 2016 (44). The AMBER force field was chosen because it exhibits more conformational variability compared with the CHARMM force field in a previous study (34). The structures of GACC were also saved every 0.2 ps and the dihedral angles defined in *SI Appendix, Table S1* were calculated using PLUMED 2.4 (45). Similar to AIB₉, the datasets with 0.2-ps temporal resolution were used for both AI model training and free-energy estimation.

The PDB file of the GACC structure that served as the starting point for our simulation was generated using PyMOL. The initial structure was assumed to be at a temperature of 10 K, and the system’s energy was minimized with positional restraints of $25 \text{ kcal} \cdot \text{mol}^{-1} \cdot \text{\AA}^{-2}$ in a two-step process—first, the steepest-descent algorithm was applied for 1,000 steps followed by the conjugate gradient algorithm for another 1,000 steps. Next, the system was equilibrated to 150 K over 100 ps under the constant number, volume, and temperature (NVT) ensemble with positional restraints of $25 \text{ kcal} \cdot \text{mol}^{-1} \cdot \text{\AA}^{-2}$. Then, the GACC was equilibrated from 150 to 277 K at 1 atm under the NPT ensemble for 100 ps with positional restraints of $5 \text{ kcal} \cdot \text{mol}^{-1} \cdot \text{\AA}^{-2}$. Finally, a long 5-ns equilibration was performed over 5 ns at 298 K under the NPT ensemble with positional restraints of $0.5 \text{ kcal} \cdot \text{mol}^{-1} \cdot \text{\AA}^{-2}$, after which the system was copied to 48 replicas, and each replica was equilibrated to its target temperature under the NPT ensemble with positional restraints of $0.5 \text{ kcal} \cdot \text{mol}^{-1} \cdot \text{\AA}^{-2}$.

The replica temperatures were geometrically spaced temperatures ranging from 277 to 408 K, with temperature increased by 1% for each replica. The attempt to exchange configurations was made every 10 ps, which is determined by checking the time correlation function of the potential energy.

Network Structure and Hyperparameters. The structure of the network is shown in Fig. 1. It has a U-net structure (20), where the input is down-sampled by four residue blocks and then up-sampled by another four residue blocks. The diffusion process is divided into 1,000 discrete steps. The network parameters are optimized by the Adam algorithm (47) with a learning rate equal to 2×10^{-5} . Exponential moving average (EMA) with a decrease rate 0.995 is used to stabilize the stochastic gradient descent.

Data Availability. Code for implementing DDPM on REMD data is available at https://github.com/tiwarilab/DDPM_REMD (48). All PLUMED input files required to calculate the observables of two example systems are available on PLUMED-NEST (<https://www.plumed-nest.org/>), the public repository of the PLUMED consortium (45), as plumedID:22.030, (49).

ACKNOWLEDGMENTS. We thank Alexander A. Alemi for initially giving us the idea of considering diffusion probabilistic models. We thank Shams Mehdi for helping set up the AIB₉ simulation. We also thank Zachary Smith, Ziyue Zou,

Eric Beyerle, Bodhi Vani, and Dedi Wang for proofreading the manuscript. This work was supported by NSF Grant CHE-2044165 and used the Extreme Science and Engineering Discovery Environment (XSEDE) Bridges through allocation

TG-CHE180053, which is supported by NSF Grant ACI-1548562. Y.W. thanks the National Cancer Institute (NCI)-University of Maryland (UMD) Partnership for Integrative Cancer Research for financial support.

1. K. J. Laidler, The development of the Arrhenius equation. *J. Chem. Educ.* **61**, 494 (1984).
2. M. L. Scalley, D. Baker, Protein folding kinetics exhibit an Arrhenius temperature dependence when corrected for the temperature dependence of protein stability. *Proc. Natl. Acad. Sci. U.S.A.* **94**, 10636–10640 (1997).
3. H. S. Chan, K. A. Dill, Protein folding in the landscape perspective: Chevron plots and non-Arrhenius kinetics. *Proteins* **30**, 2–33 (1998).
4. J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics" in *Proceedings of the 32nd International Conference on Machine Learning*, F. Bach, D. Blei, Eds. (Proceedings of Machine Learning Research, Lille, France, 2015), vol. 37, pp. 2256–2265.
5. J. Ho, A. Jain, P. Abbeel, "Denosing diffusion probabilistic models in Advances" in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin, Eds. (Curran Associates, Inc.), vol. 33, pp. 6840–6851 (2020).
6. Y. Sugita, Y. Okamoto, Replica-exchange molecular dynamics method for protein folding. *Chem. Phys. Lett.* **314**, 141–151 (1999).
7. U. H. Hansmann, Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.* **281**, 140–150 (1997).
8. C. Abrams, G. Bussi, Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. *Entropy (Basel)* **16**, 163–199 (2014).
9. R. Abel, L. Wang, E. D. Harder, B. J. Berne, R. A. Friesner, Advancing drug discovery through enhanced free energy calculations. *Acc. Chem. Res.* **50**, 1625–1632 (2017).
10. P. Liu, B. Kim, R. A. Friesner, B. J. Berne, Replica exchange with solute tempering: A method for sampling biological systems in explicit water. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 13749–13754 (2005).
11. L. Wang, R. A. Friesner, B. J. Berne, Replica exchange with solute scaling: A more efficient version of replica exchange with solute tempering (REST2). *J. Phys. Chem. B* **115**, 9431–9438 (2011).
12. A. J. Ballard, C. Jarzynski, Replica exchange with nonequilibrium switches. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 12224–12229 (2009).
13. S. Trebst, M. Troyer, U. H. Hansmann, Optimized parallel tempering simulations of proteins. *J. Chem. Phys.* **124**, 174903 (2006).
14. W. Nadler, U. H. Hansmann, Optimizing replica exchange moves for molecular dynamics. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **76**, 057102 (2007).
15. J. Kim, T. Keyes, J. E. Straub, Generalized replica exchange method. *J. Chem. Phys.* **132**, 224107 (2010).
16. J. D. Chodera, M. R. Shirts, Replica exchange and expanded ensemble simulations as Gibbs sampling: Simple improvements for enhanced mixing. *J. Chem. Phys.* **135**, 194110 (2011).
17. A. Gil-Ley, G. Bussi, Enhanced conformational sampling using replica exchange with collective-variable tempering. *J. Chem. Theory Comput.* **11**, 1077–1085 (2015).
18. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, 2016).
19. E. Gallicchio, M. Andrec, A. K. Felts, R. M. Levy, Temperature weighted histogram analysis method, replica exchange, and transition paths. *J. Phys. Chem. B* **109**, 6722–6731 (2005).
20. O. Ronneberger, P. Fischer, T. Brox, "U-net: Convolutional networks for biomedical image segmentation" in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, A. F. Frangi, Eds. (Springer International Publishing, Cham, 2015), pp. 234–241.
21. Y. Wu, K. He, "Group normalization" in *Proceedings of the European Conference on Computer Vision (ECCV)* (Springer, Munich, Germany, 2018) pp. 3–19.
22. A. Vaswani et al., "Attention is all you need" in *Advances in Neural Information Processing Systems*, I. Guyon, et al., Eds. (Curran Associates, Inc., 2017), vol. 30, pp. 5998–6008.
23. P. Vincent, A connection between score matching and denoising autoencoders. *Neural Comput.* **23**, 1661–1674 (2011).
24. W. L. Jorgensen, J. Tirado-Rives, Potential energy functions for atomic-level simulations of water and organic and biomolecular systems. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 6665–6670 (2005).
25. J. Huang et al., CHARMM36m: An improved force field for folded and intrinsically disordered proteins. *Nat. Methods* **14**, 71–73 (2017).
26. V. Botan et al., Energy transport in peptide helices. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 12749–12754 (2007).
27. M. Biswas, B. Lickert, G. Stock, Metadynamics enhanced Markov modeling of protein dynamics. *J. Phys. Chem. B* **122**, 5508–5514 (2018).
28. S. Mehdi, D. Wang, S. Pant, P. Tiwary, Accelerating all-atom simulations and gaining mechanistic understanding of biophysical systems through state predictive information bottleneck. *J. Chem. Theory Comput.* **18**, 3231–3238 (2022).
29. M. R. Shirts, J. D. Chodera, Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **129**, 124105 (2008).
30. S. Bannwarth, A. Gaignol, HIV-1 TAR RNA: The target of molecular interactions between the virus and its host. *Curr. HIV Res.* **3**, 61–71 (2005).
31. U. Schulze-Gahmen, J. H. Hurley, Structural mechanism for HIV-1 TAR loop recognition by Tat and the super elongation complex. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 12973–12978 (2018).
32. L. R. Ganser, M. L. Kelly, D. Herschlag, H. M. Al-Hashimi, The roles of structural dynamics in the cellular functions of RNAs. *Nat. Rev. Mol. Cell Biol.* **20**, 474–489 (2019).
33. J. Šponer et al., RNA structural dynamics as captured by molecular simulations: A comprehensive overview. *Chem. Rev.* **118**, 4177–4338 (2018).
34. C. Bergonzo, N. M. Henriksen, D. R. Roe, T. E. Cheatham 3rd, Highly sampled tetranucleotide and tetraloop motifs enable evaluation of common RNA force fields. *RNA* **21**, 1578–1590 (2015).
35. Y. Liu, M. E. Tuckerman, Generalized Gaussian moment thermostating: A new continuous dynamical approach to the canonical ensemble. *J. Chem. Phys.* **112**, 1685–1700 (2000).
36. G. Gobbo, B. J. Leimkuhler, Extended Hamiltonian approach to continuous tempering. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **91**, 061301 (2015).
37. M. Dibak, L. Klein, F. Noé, Temperature-steerable flows. *arXiv [Preprint]* (2020). <https://doi.org/10.48550/arXiv.2012.00429>. Accessed 1 December 2020.
38. D. Rezende, S. Mohamed, "Variational inference with normalizing flows" in *Proceedings of the 32nd International Conference on Machine Learning*, F. Bach, D. Blei, Eds. (Proceedings of Machine Learning Research, Lille, France, 2015), vol. 37, pp. 1530–1538.
39. F. Noé, S. Olsson, J. Köhler, H. Wu, Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science* **365**, eaaw1147 (2019).
40. H. Wu, J. Köhler, F. Noé, "Stochastic normalizing flows" in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin, Eds. (Curran Associates, Inc., 2020), vol. 33, pp. 5933–5944.
41. M. Parrinello, A. Rahman, Crystal structure and pair potentials: A molecular-dynamics study. *Phys. Rev. Lett.* **45**, 1196–1199 (1980).
42. W. G. Hoover, Canonical dynamics: Equilibrium phase-space distributions. *Phys. Rev. A Gen. Phys.* **31**, 1695–1697 (1985).
43. D. J. Evans, B. L. Holian, The Nose-Hoover thermostat. *J. Chem. Phys.* **83**, 4069–4074 (1985).
44. M. Abraham, B. Hess, D. van der Spoel, E. Lindahl, *Gromacs Reference Manual, Version 2016.4* (Department of Biophysical Chemistry, University of Groningen, Groningen, The Netherlands, 2016).
45. M. Bonomi; PLUMED Consortium, Promoting transparency and reproducibility in enhanced molecular simulations. *Nat. Methods* **16**, 670–673 (2019).
46. G. Bussi, M. Parrinello, Accurate sampling using Langevin dynamics. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **75**, 056707 (2007).
47. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization. *arXiv [Preprint]*, arXiv:1412.6980 (2014).
48. Y. Wang, L. Herron, P. Tiwary, DDPM-REMD, Denoising diffusion probabilistic models for replica exchange MD simulation. GitHub. <https://github.com/tiwarylab/DDPM-REMD>. Deposited 14 August 2021.
49. Y. Wang, plumed:22.030, Mixing physics across temperatures with generative artificial intelligence. PLUMED-NEST. <https://www.plumed-nest.org/eggs/22/030/>. Deposited 21 July 2022.