



Inferring phase transitions and critical exponents from limited observations with thermodynamic maps

Lukas Herron^{a,b}, Kinjal Mondal^a, John S. Schneekloth Jr.^c, and Pratyush Tiwary^{b,d,1}

Affiliations are included on p. 8.

Edited by Frederick C. MacKintosh, Rice University, Houston, TX; received December 13, 2023; accepted October 15, 2024, by Editorial Board Member Paul Chaikin

Phase transitions are ubiquitous across life, yet hard to quantify and describe accurately. In this work, we develop an approach for characterizing generic attributes of phase transitions from very limited observations made deep within different phases' domains of stability. Our approach is called thermodynamic maps (TM), which combines statistical mechanics and molecular simulations with score-based generative models. TM enable learning the temperature dependence of arbitrary thermodynamic observables across a wide range of temperatures. We show its usefulness by calculating phase transition attributes such as melting temperature, temperature-dependent heat capacities, and critical exponents. For instance, we demonstrate the ability of TM to infer the ferromagnetic phase transition of the Ising model, including temperature-dependent heat capacity and critical exponents, despite never having seen samples from the transition region. In addition, we efficiently characterize the temperature-dependent conformational ensemble and compute melting curves of the two RNA systems: a GCAA tetraloop and the HIV-TAR RNA, which are notoriously hard to sample due to glassy-like energy landscapes.

phase transitions | statistical mechanics | AI | complex systems | RNA

Phase transitions are widely observed in biological, material, and social sciences. Across these disciplines, phase transitions can be defined as the emergence of higher-level, large-scale organization from the coordinated, short-range interactions between many individual constituents. Classical examples include the ferromagnetic to paramagnetic transition, boiling of water, and conformational transitions in biomolecules like proteins and nucleic acids.

Statistical mechanics, and especially the framework of energy landscapes, provides a simple and unifying way of studying phase transitions in these diverse systems (1). In this work, we are specifically interested in phase transitions in systems that stay in equilibrium throughout. For these, the Boltzmann distribution relates the probability of finding a system in a particular microscopic configuration \mathbf{x} to its energy $U(\mathbf{x})$ and the system's inverse temperature β as

$$\mu(\mathbf{x}, \beta) = \frac{e^{-\beta U(\mathbf{x})}}{Z(\beta)} \quad \text{with} \quad [1]$$

$$Z(\beta) = \int e^{-\beta U(\mathbf{x})} d\mathbf{x}. \quad [2]$$

$Z(\beta)$ is a normalization constant known as the partition function, whose behavior is often associated with phase transitions. Exploration of the energy landscape is guided by competition between energy and entropy, which is encapsulated by a temperature-dependent free energy $F(\beta)$ which may be computed from the partition function as

$$F(\beta) = -\beta^{-1} \ln Z(\beta). \quad [3]$$

At a glance, Eqs. 1–3 suggest that the relationship between temperature, energy, microscopic probability, and macroscopic free energy is simple and tractable. It appears as if with these equations, one has the machinery to directly calculate the free energy across temperatures. By doing so for different macroscopic phases, one could then obtain various thermodynamic attributes of phase transitions, including transition temperatures and phase diagrams. By calculating appropriate fluctuations, one could directly obtain response functions such as heat capacities.

In reality, however, the situation is quite complex. Studying phase transitions and their characteristics computationally is made difficult by Eq. 2, which requires integrating over

Significance

Phase transitions play a pivotal role in diverse scientific domains, from metal solidification to water boiling and biological macromolecule folding. However, quantifying these transitions, especially their temperature-dependent attributes, presents challenges due to complex and sparse data. This study introduces "thermodynamic maps" (TM), an approach merging statistical mechanics, molecular simulations, and generative AI. TM efficiently characterizes temperature-dependent attributes, even with limited data from stable phases. Demonstrating its efficacy, TM precisely infers critical exponents of the Ising model and deciphers complex behaviors in challenging RNA systems. Beyond theoretical implications, TM's computational efficiency suggests its potential for studying multifaceted behaviors in data-scarce scientific realms, contributing to a deeper understanding of complex systems and their coordinated constituents.

The authors declare no competing interest.

This article is a PNAS Direct Submission. F.C.M. is a guest editor invited by the Editorial Board.

Copyright © 2024 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.2321971121/-DCSupplemental>.

Published December 16, 2024.

a (usually) intractably large number of dimensions. Numerous elegant theoretical and computational schemes have been proposed over the decades to solve this problem. For example, this includes methods like free energy perturbation, Markov chain Monte Carlo (MC) methods, the replica trick and others (2–7).

In this study, we propose a generative AI based approach that characterizes phase transitions by effectively learning the temperature dependence of the partition function, and therefore the free energy. Our method, which we call “thermodynamic maps” (TM) incorporates score-based generative modeling into the framework of free energy perturbation within statistical mechanics (8–11). The central idea underlying TM is to map the temperature dependence of configurational ensembles of a complex system to and from the temperature dependence of a simple, idealized system. We show how such a bidirectional mapping can be obtained. This allows for efficient generation of physically realistic samples of the complex system with the correct Boltzmann weights. Within the framework of free energy perturbation, the mapping also allows for temperature-dependent free energy estimates.

We demonstrate the broad applicability of TM for three complex systems where we compare against benchmarks from theory, extensive computational studies, and experiments. The first system we consider is the Ising model on a two-dimensional square lattice. With observations made at two temperatures, one deep in the paramagnetic regime and the other deep in the ferromagnetic regime, we are able to correctly infer critical behavior. We then study two different RNA systems: the GCAA tetraloop and HIV-TAR RNA (12, 13). For both of these, we infer the temperature dependence of the equilibrium distribution across temperatures using TM trained on data generated by short molecular dynamics (MD) simulations. For both RNA systems, we predict the temperature dependence of conformational ensemble and compute melting temperatures, finding agreement with computational studies and experiment.

Given this demonstration of applicability, we believe that TM will be found useful for the characterization of complex phase transitions in diverse systems, especially those with multiple phases. For example, local minima within the energy landscape of biomolecular systems are often biologically and functionally relevant (14), and while they have been historically challenging to characterize computationally, we are able to study them with TM. The computational efficiency of the learning algorithm and scalability due to not requiring samples from the global equilibrium distribution make TM especially suitable for studying large-scale systems exhibiting complex behavior across long timescales.

1. Materials and Methods

1.1. Targeted Free Energy Perturbation. To motivate TM, we start with relevant background work in the longstanding problem of free energy estimation. Since pioneering work by Zwanzig and Feynman in the 1950s, many frameworks have been developed to do this efficiently (2, 5, 6, 15, 16). The guiding principle behind all such methods is that one is usually interested in differences of free energies, rather than their absolute values. Such differences can be estimated from ratios of partition functions. One relatively recent framework is that of targeted free energy perturbation (TFEP), which is especially relevant for machine learning approaches to free energy estimation (17).

Estimates of free energy differences between states converge according to the overlap between the states. Avoiding a technical definition, the overlap between states A and B measures how likely it is to generate a sample from B by sampling from A . The central idea behind TFEP is to increase the overlap between states with an invertible mapping of the configuration space onto itself:

$$\mathcal{M} : \mathbf{x} \rightarrow \mathbf{x}' \quad [4]$$

A well-chosen \mathcal{M} might map $\mathbf{x} \in A$ so that A' is a subset of B . Under such a mapping, sampling from A and transforming the sample using Eq. 4 always generates a sample from B ; it is therefore more efficient to estimate free energies under \mathcal{M} and reweight back to the original configuration space. Reweighting is possible if \mathcal{M} has a well-defined Jacobian $J_{\mathcal{M}}$ by using the following property: if $p(\mathbf{x})$ is the probability density of $\mathbf{x} \in A$ and $q(\mathbf{x}')$ of $\mathbf{x}' \in A'$, then the densities are related through the identity

$$p(\mathbf{x}) = \frac{q(\mathbf{x}')}{|\det J_{\mathcal{M}}(\mathbf{x}')|} \quad [5]$$

Using Eq. 5 free energy differences computed under \mathcal{M} can be reweighted to free energy differences in the original configuration space.

Clearly using Eqs. 4 and 5 depends on specifying a map that increases the overlap between states and has a tractable Jacobian. Finding a suitable mapping is difficult due to the often complex, high-dimensional nature of the configurations; as a result, the most successful approaches to TFEP represent \mathcal{M} with a neural network (18, 19). Likelihood-based models have been especially appealing since they meet both conditions (20–22), but have drawbacks of reduced expressivity and robustness (23). Recent lines of work have improved both aspects—for example, by pursuing more expressive black-box operations or incorporating stochasticity (24–27). In the following sections, we bypass these issues of finding a suitable mapping by exploring nonequilibrium diffusion processes as general-purpose maps.

1.2. Leveraging Nonequilibrium Thermodynamics. We begin by recounting the main results of a line of work that explores using properties of stochastic processes for general inference (8, 10, 28–31). Suppose that the diffusion process is modeled as a Fokker–Planck equation cast in the form of a Liouville equation

$$\frac{\partial}{\partial t} p(\mathbf{x}, t) = -\lambda(t) \nabla \cdot [b(\mathbf{x}, t) p(\mathbf{x}, t)] \quad \text{with} \quad [6]$$

$$b(\mathbf{x}, t) = \mathbf{x} - \nabla \log p(\mathbf{x}, t), \quad [7]$$

where gradients are assumed to be with respect to spatial coordinates \mathbf{x} (31, 32). Eq. 6 is a partial differential equation (PDE) describing the evolution of a probability distribution $p(\mathbf{x}, t)$ toward a Gaussian distribution under the influence of the vector field $b(\mathbf{x}, t)$ (Eq. 7). Time is dilated by the noise schedule $\lambda(t)$ —a monotonically increasing scalar-valued function that ensures the process converges sufficiently close to a Gaussian distribution at $t = 1$. Equivalently, the process can be expressed in terms of the evolution of a particle using a stochastic differential equation (SDE) of the form

$$d\mathbf{x} = -\lambda(t) \mathbf{x} dt + \sqrt{2\lambda(t)} dB_t, \quad [8]$$

where B_t is a standard Brownian motion and \mathbf{x} is implicitly a function of t .

In Fig. 1, we solve Eqs. 6 and 8 numerically in one spatial dimension. Fig. 1A shows the solution of Eq. 6 when $p(\mathbf{x}, 0)$ is a mixture of Gaussians. In this case, it is possible to evaluate the PDE exactly since $\log p(\mathbf{x}, t)$ is known for all t . The dynamics of Eq. 6 are presented in the *Center* panel, where an initial distribution is transported to a single Gaussian; during the process, the barriers that originally separated the states are destroyed. Similarly, the trajectories generated by the SDE begin at points on $\log p(\mathbf{x}, t)$ (marked as stars) and end in a Gaussian distribution at $t = 1$. The diffusion process demonstrably increases overlap by mapping all states onto a single reference state and does so for arbitrary initial conditions. However, if the diffusion process is to be a valid TFEP map it must meet the invertibility requirement in Section 1.1.

Indeed, Eq. 6 is invertible under the time-reversal $\tau \equiv 1 - t$. Remarkably, it can be shown that the time-reverse of Eq. 8 also exists (33) and is an SDE of the form

$$d\mathbf{x} = -\lambda(\tau) [\mathbf{x} + \nabla_{\mathbf{x}} \log p(\mathbf{x}, \tau)] d\tau + \sqrt{2\lambda(\tau)} dB_{\tau}. \quad [9]$$

Note that the drift term is the same as in Eqs. 6 and 7 but with the opposite sign.

While particles simulated under Eq. 8 approach equilibrium, particles under the influence of Eq. 9 are driven out of equilibrium. When the variance of B_{τ} is

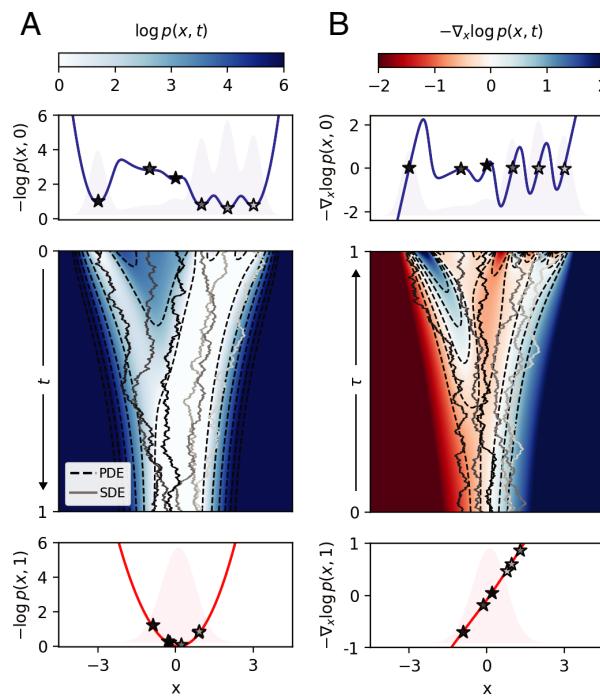


Fig. 1. Transporting a probability distribution with a diffusion process. (A) The Top panel shows the negative log-likelihood for a mixture of Gaussians. The Center panels show the dynamics of a diffusion process defined as a PDE (contours) and an SDE (solid lines). Over a time interval from 0 to 1 the diffusion process relaxes to a Gaussian distribution. (B) The same dynamics are depicted in terms of the score. In both panels, the dynamics of the diffusion process can be viewed forward or backward in time (denoted by t and τ , respectively).

zero, Eq. 9 reduces to an ordinary differential equation (ODE) with a well-defined Jacobian, implying the existence of inverse mappings \mathcal{M}_t and \mathcal{M}_τ . For the SDE the inverse mapping exists in a probabilistic sense since Eq. 9 samples from $p(\mathbf{x}, 0)$. In either case, the existence of such an inverse mapping is critical to the class of models studied here. Here on, t will denote the forward-time process and τ the time-reverse process.

Since the diffusion process meets the TFEP criteria, we are now faced with evaluating the forward and reverse processes using the PDE in Eq. 6 or the pair of SDEs in Eqs. 8 and 9. In either case, one must estimate either the log-likelihood-log $p(\mathbf{x}, t)$, or the score- $-\nabla \log p(\mathbf{x}, t)$. Generative models can be broadly categorized based on which quantity they estimate: likelihood-based models solve the PDE by parameterizing the log-likelihood (20, 22, 34), while score-based models solve the SDEs by parameterizing the score (9, 10). The relationship between the likelihood and the score is depicted in Figs. 1A and B. In this work, we focus primarily on score-based models.

Score-Based Models. Score-based models (SBMs) have emerged as a robust method for generating independent samples from intractable probability distributions (10). Using the reverse diffusion in Eq. 9, SBMs transport a simple *prior* distribution to a potentially complex target probability density that is characterized by a collection of samples \mathcal{D} . The model approximates the score of Eq. 8, which is used to simulate Eq. 9—having the effect of transporting samples from the prior distribution to the target. In an abuse of notation, we define the target distribution as $p(\mathbf{x}) \equiv p(\mathbf{x}, t = 1)$ and the prior distribution as $q(\mathbf{x}) \equiv p(\mathbf{x}, t = 0)$.

The main advantage of defining the diffusion process with the form in Eq. 8 is that realizations of the forward diffusion at t can be generated without simulation; in ref. 10, the distribution of positions originating from $\mathbf{x}_0 \equiv \mathbf{x}(t = 0)$ after time t is derived from Eq. 8 as

$$\mathbf{x}(t) \sim \mathcal{N}\left(\mathbf{x}_0 e^{-\int_0^t \lambda(s) ds}, \mathbf{I}(1 - e^{-\int_0^t \lambda(s) ds})\right), \quad [10]$$

and it can be shown that $\nabla_{\mathbf{x}} \log p(\mathbf{x}, t)$ is proportional to the mean displacement vector from \mathbf{x}_0 to $\mathbf{x}(t)$ (35). The corresponding distribution for the reverse-time diffusion does not have a closed-form solution; to evaluate Eq. 9 one must estimate the score for all \mathbf{x} and τ and simulate from an initial condition $\mathbf{x}(\tau = 0)$.

A SBM uses a neural network to model $s_\theta(\mathbf{x}, t) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x}, t)$ from displacement vectors produced by Eq. 10. The network predicts the score by minimizing the score-matching objective function:

$$\operatorname{argmin}_\theta \mathbb{E}_{\mathbf{x}_0 \in \mathcal{D}} \mathbb{E}_{t \in \mathcal{U}(0,1)} \|s_\theta(\mathbf{x}(t), t) - \nabla_{\mathbf{x}} \log p(\mathbf{x}, t)\|^2. \quad [11]$$

Once trained, estimates of the score predicted by the network can be used to simulate Eq. 9. If the estimates are sufficiently accurate, then the endpoints of the trajectories launched from $q(\mathbf{x})$ are independent and identically distributed samples from $p(\mathbf{x})$. In the next section, we provide an alternative construction wherein $q(\mathbf{x})$ is viewed as the distribution arising from the dynamics of a physical system.

1.3. Thermodynamic Maps. The Einstein-Smoluchowski relation states that temperature is proportional to the diffusivity of a Brownian motion (36). Accordingly, we introduce temperature to the SBM SDEs (Eqs. 8 and 9) by augmenting the positions $\mathbf{x} \in \mathbb{R}^d$ with auxiliary inverse temperature-type variables $\beta \in \mathbb{R}^d$. Together, the coordinates and temperatures form a state vector $(\mathbf{x}, \beta)^\top \in \mathbb{R}^{2d}$, and the thermodynamic map is parameterized in the joint \mathbf{x}, β -space as a SBM of the form

$$\begin{pmatrix} \frac{d\mathbf{x}}{d\beta^{-1}} \end{pmatrix} = -\lambda(t) \begin{pmatrix} \mathbf{x} \\ \beta^{-1} \end{pmatrix} dt + \sqrt{2\lambda(t)} \begin{pmatrix} \sqrt{\beta_0^{-1}} \\ 1 \end{pmatrix} d\beta_t \quad \text{and} \quad [12]$$

$$\begin{pmatrix} \frac{d\mathbf{x}}{d\beta^{-1}} \end{pmatrix} = -\lambda(\tau) \left[\begin{pmatrix} \mathbf{x} \\ \beta^{-1} \end{pmatrix} + s_\theta(\mathbf{x}, \beta^{-1}, \tau) \right] d\tau + \sqrt{2\lambda(\tau)} \begin{pmatrix} \sqrt{\beta_0^{-1}} \\ 1 \end{pmatrix} d\beta_\tau. \quad [13]$$

Both \mathbf{x} and β fluctuate during the dynamics of Eqs. 12 and 13, relaxing to $\mathcal{N}(\mathbf{0}, \beta_0^{-1})$ and $\mathcal{N}(\mathbf{0}, \mathbf{1})$ respectively at $t = 1$. The value of the auxiliary variables β at $t = 0$ is denoted by β_0 and is a constant that sets the temperature of the diffusive dynamics for each component of \mathbf{x} by rescaling the Brownian motion. The construction ensures that the diffusion process for \mathbf{x} relaxes to a normal distribution at temperature β_0^{-1} .

A thermodynamic map (Fig. 2) maps the thermodynamics of a target system $p(\mathbf{x}, \beta)$ onto the thermodynamics of a simple prior system $q(\mathbf{x}', \beta')$, where in both systems each degree of freedom is coupled to separate baths defined by β . Our construction is motivated by treating the prior distribution as the equilibrium distribution arising from the dynamics of a physical system. If the prior is a Gaussian distribution then the corresponding physical system is a simple harmonic oscillator, and generating samples from $p(\mathbf{x})$ at β amounts to sampling from the harmonic oscillator $q(\mathbf{x}')$ at the corresponding β' (Fig. 2). Details on training and inference are provided in *SI Appendix, section A*.

In the generalized construction above we have assumed that each degree of freedom of \mathbf{x} is coupled to a separate bath defined by the entries of β , but in most cases physical systems are coupled to a single bath. If samples of $p(\mathbf{x}, \beta)$ are from the global equilibrium distribution of an ensemble defined by β_{bath} , each entry of β should be equal to β_{bath} . However, global equilibrium distribution is often not attainable in practice (e.g., in molecular dynamics simulations). In this case, allowing each degree of freedom of \mathbf{x} to be coupled to separate baths β allows the global equilibrium distribution to be inferred from samples gathered under a local equilibrium distribution. In *SI Appendix, section G* we discuss how the global equilibrium distribution can be inferred from local equilibrium samples if one assumes that β varies linearly with β_{bath} . In Section 1.5, we demonstrate the validity of the assumption by inferring the global equilibrium distribution from short, independent molecular dynamics simulations of two RNAs.

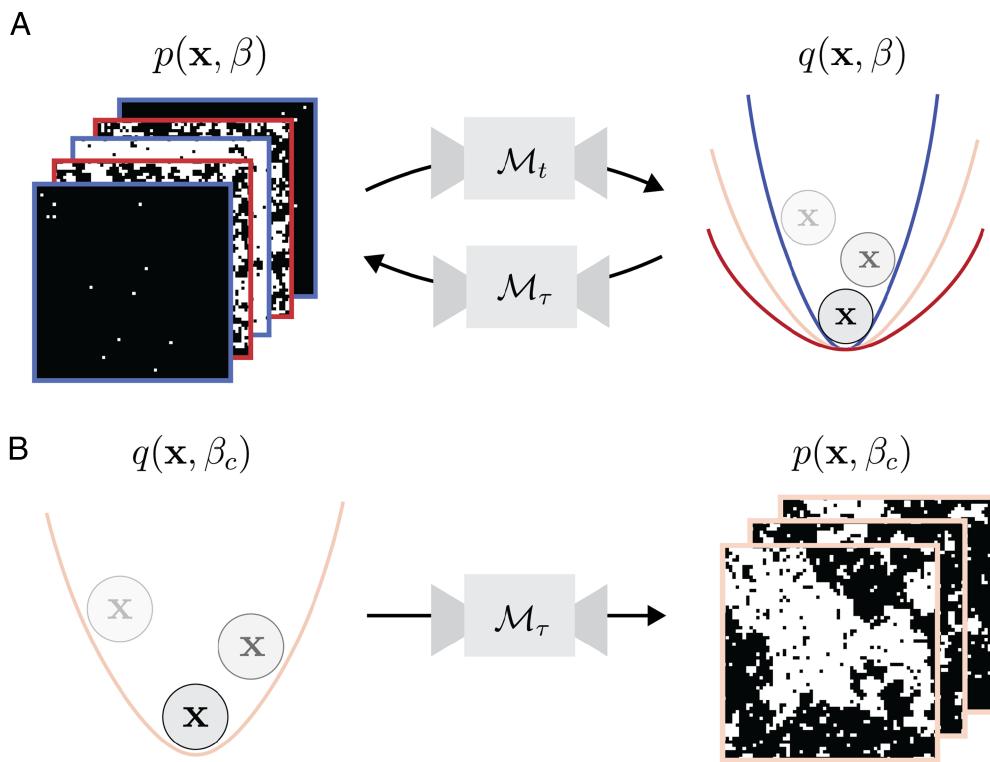


Fig. 2. Illustration of a thermodynamic map between systems. (A) The thermodynamic map is parameterized as a score-based model transporting densities from a target distribution $p(\mathbf{x}, \beta)$ to a prior distribution $q(\mathbf{x}, \beta)$. In the schematic, data from an Ising model at two temperatures are mapped onto a simple harmonic oscillator. The generative process \mathcal{M}_τ is parameterized from realizations of the forward process \mathcal{M}_t . (B) Once learned, the thermodynamic map allows samples of the complex system to be generated from the simple prior system at any temperature—even those showing nontrivial behavior—by sampling the prior at the correct temperature.

1.4. Critical Behavior of the Ising Model. We apply TM to the most widely studied formulation of the classical Ising model as a system of interacting spins arranged on a two-dimensional square lattice, without the presence of an external magnetic field. Arguably, this model serves as a prototypical example of a simple system with a complex phase transition. Each spin can have one of two states with value $\sigma = +1$ or $\sigma = -1$, and the spins interact with their nearest neighbors through an interaction term J and Hamiltonian given by

$$\mathcal{H} = J \sum_{\langle i,j \rangle} \sigma_i \sigma_j \quad [14]$$

where $\langle i,j \rangle$ denotes nearest-neighbor pairs. When $J < 0$ the model exhibits ferromagnetic behavior in its ground state. Going forward, we set $J = -1$ without loss of generality.

As the temperature increases, an Ising model in two or higher dimensions transitions from an ordered magnetic phase to a disordered paramagnetic phase. For our set-up of a two-dimensional Ising model on a square lattice with $J = -1$, this critical temperature is known to be $T_c \approx 2.27$ (37). The two phases can be distinguished from each other with the magnetization order parameter, M , defined as the absolute value of the average of the spins. Well above T_c , as $T \rightarrow \infty$, the spins are equally likely to be $+1$ or -1 , regardless of their neighbors, resulting in a net magnetization of zero. Well below T_c , all spins in the lattice align, leading to a magnetization of 1. Power law scaling is a signature of critical behavior, and near the critical temperature the magnetization M and heat capacity C of the Ising model scale as

$$M \sim |\tau|^\beta \quad \text{and} \quad C \sim |\tau|^{-\alpha} \quad \text{where} \quad \tau = \frac{T - T_c}{T_c}, \quad [15]$$

and with critical exponents $\alpha = 0$ and $\beta = 0.125$ (not to be confused with the inverse temperature variables β), which can be derived analytically in the thermodynamic limit (38). For systems that are not solvable, the critical

exponents must be measured numerically, which is often done through MC sampling, wherein proposals for spin flips are generated and their acceptance or rejection is determined based on the detailed balance condition. Near the critical temperature, the presence of long-range correlations causes MC dynamics to slow down exponentially, leading to difficulty in sampling the phase transition (39).

We investigate the ability of TM to infer such critical behaviors by generating configurations of an Ising model through MC sampling across temperatures and using the data from two temperatures asymmetrically spaced about T_c to train a TM (details provided in *SI Appendix* C). In the formalism of Eqs. 12 and 13 the lattice configurations are $\mathbf{x}_0 \in \mathbb{R}^{32 \times 32}$ and each entry of $\beta_0^{-1} \in \mathbb{R}^{32 \times 32}$ is the bath temperature; a straightforward choice since MC sampling guarantees the global equilibrium distribution. Once the TM is trained, we generate configurations at each temperature by evaluating Eq. 13 with each entry of β_0^{-1} equal to T_{bath} . The generated continuous-valued configurations \mathbf{x} are then transformed into discrete spins by $\sigma = \text{sgn}(\mathbf{x})$, and the average magnetization and heat capacity at each temperature is estimated directly from the spin configurations. Fig. 3 shows the behavior of M and C for the MC and TM-generated samples. The TM infers the correct value of $T_c(\tau = 0)$ even based on limited, deliberately misleading training data and generates samples with divergences in M and C .

Fig. 3A shows the behavior of the magnetization for MC and TM-generated samples across temperatures, with critical exponents measured as $\beta_{\text{MC}} \approx 0.149 \pm 0.004$ and $\beta_{\text{TM}} \approx 0.178 \pm 0.012$. Similarly, Fig. 3B shows a divergence in the heat capacity at $T_{\text{MC}} \approx 2.25$ and $T_{\text{TM}} \approx 2.30$, with critical exponents $\alpha_{\text{MC}} \approx 0.338 \pm 0.050$ and $\alpha_{\text{TM}} \approx 0.236 \pm 0.061$. For both exponents, the MC and TM exponents agree with each other, although far from the ideal value due to finite size effects.

These results highlight the ability of our model to infer the nontrivial thermodynamic behavior of phase transitions without being shown samples from the transition region; the physical meaning of temperature in the prior

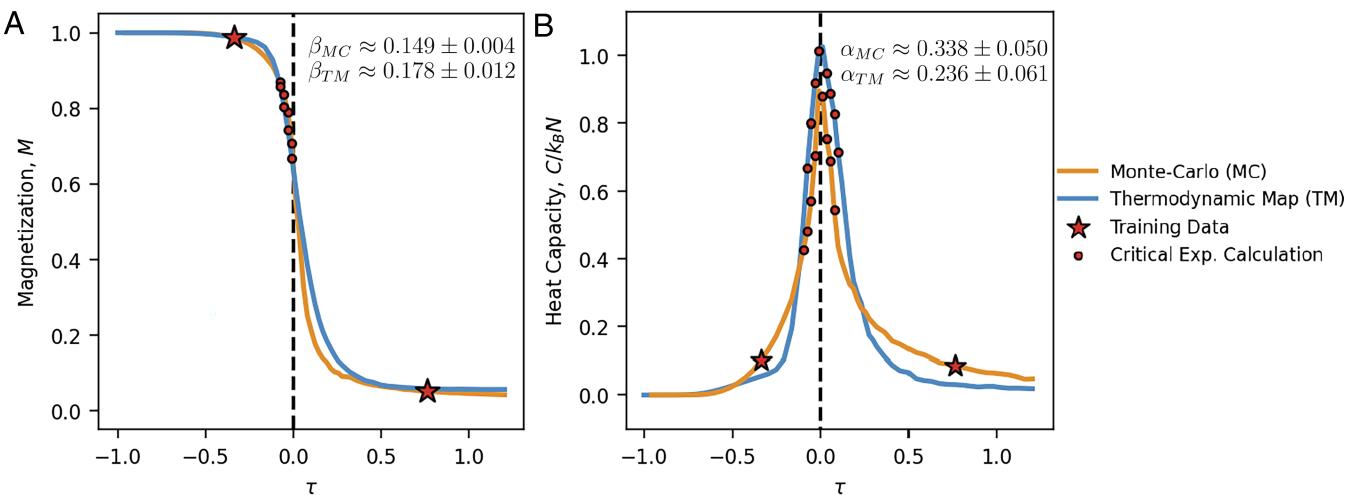


Fig. 3. Inferring the phase transition of the 2D Ising model from limited sampling. (A) The magnetization is plotted for samples of a 32×32 square Ising model generated through MC sampling (orange) and the thermodynamic map (blue). The thermodynamic map predicts change in magnetization at T_c when trained on samples generated at $T = 1.5$ and $T = 4$ (red stars). Similar results were obtained (not reported) for a symmetric choice of training data temperatures. (B) The heat capacity of samples generated from MC sampling (orange) and the thermodynamic map (blue) is plotted. The thermodynamic map infers the divergence in the heat capacity, numerically computed for the red dots, when trained on the same samples as panel A (red stars).

system has been transferred to the complex system, complete with the statistics properties associated with critical behavior. The correct prediction of T_m , even with asymmetrically spaced training temperatures, indicates our model has learned the physics of the Ising model, not merely the distribution of structures at each temperature.

1.5. Exploring RNA Conformational Landscapes with Thermodynamic Map-accelerated Molecular Dynamics (TM-aMD). To show the broad applicability of TM, we now study conformational transitions and melting in two different RNA (RNA) systems. Studying atomic-resolution conformational ensembles of RNAs through molecular dynamics simulations has proven crucial for understanding RNA structural dynamics, yet remains challenging due to the disordered, glassy nature of RNA energy landscapes (40, 41).

Increasing evidence points toward some RNAs having glassy energy landscapes where multiple minima are separated by high barriers (42–44). The ruggedness of the landscape results in many competing degrees of freedom and no clear-cut separation of timescales within the dynamics (43–45). The most striking feature of glassy energy landscapes as they relate to polymers is the presence of conformational heterogeneity at equilibrium (43, 46); while the conformational landscape of proteins is often dominated by a single well-defined fold (energy minimum), the RNA conformational landscape may not be dominated by a single structure (47, 48). This difference is analogous to the single magnetized phase of an Ising model and the multiple phases of long-ranged spin glasses (7, 49–51). Since multiple members of the ensemble contribute substantially to the free energy and other thermodynamic observables, RNAs are best described as a weighted ensemble of conformers (48, 52).

Exploring energy landscapes by biasing dynamics along a small number of slow degrees of freedom has proven successful for exploring the conformational landscape of proteins, but the lack of timescale separation and ensemble nature of the RNA conformational landscape violates fundamental assumptions of dimensionality made in biasing methods (53). On the other hand, TM have the advantage of being able to learn the conformational landscape directly in a much higher-dimensional configuration space.

We train thermodynamics maps on limited information generated by bioinformatic approaches and multiensemble molecular dynamics simulations to efficiently characterize RNA conformational ensembles. Our starting point is the physics and knowledge-based potential that is central to Rosetta to generate a putative conformational ensemble (54). These structures serve as the starting point to explore a more realistic energy landscape through all-atom, explicit solvent molecular dynamics performed over a range of temperatures. Between rounds of molecular dynamics simulation, the global equilibrium distribution is inferred using TM alongside assumptions of local equilibrium

(SI Appendix, section G). Initial conditions are resampled from the inferred equilibrium distribution using a geometric RNA order parameter. Here on, we refer to this protocol as thermodynamic map-accelerated Molecular Dynamics (TM-aMD).

Although agreement between the input equilibrium distribution and the output from the thermodynamic map is a necessary, though not sufficient condition that the true equilibrium distribution has been attained, we leave rigorously addressing convergence to equilibrium to future work. Here, we take solace in the numerical results for the two challenging test systems, described next.

1.5.1. GCAA tetraloop. With a shift in perspective toward viewing RNAs as dynamic entities, there has been interest in studying the variation in dynamics between different tetraloops. Studies combining computation and experiment have demonstrated that even so-called simple tetraloops can exhibit rich dynamics (55). We study the GCAA tetraloop, a well-studied model system, enabling us to compare the equilibrium distribution generated by our model with extensive molecular dynamics simulations and experimental data (40, 56, 57). The GCAA Tetraloop is a small, highly stable, 12-nucleotide RNA sequence that adopts a hairpin structure consisting of an eight-nucleotide helix and a four-nucleotide loop (PDB: 1ZIH) (12). Consistent with an ensemble perspective, the variable arrangement of nucleotides in the loop gives rise to alternative conformations, which we investigate with TM-aMD.

The thermodynamic map learns to generate RNA structures represented as G -vectors—an internal coordinate system for RNAs that effectively clusters distinct folded states (58). The principal components of G -vectors have been shown to be a convenient visualization of RNA structural diversity, which we use to guide adaptive sampling. Further information on G -vectors, along with details of TM-aMD can be found in SI Appendix, section D.

Fig. 4 summarizes the result of nine iterations of our enhanced sampling procedure, with a total of $50 \mu\text{s}$ of simulation. Although we performed extensive MD simulations reaching long timescales, we still used two orders of magnitude less compute compared to our reference millisecond replica exchange simulation, even with suboptimal scheduling of TM learning with MD simulation (40). SI Appendix, Fig. S2A suggests that TM-aMD would benefit from more frequent reseeding of simulations.

In Fig. 4A, we present the projection of the learned equilibrium distribution onto the first two principal components of the G -vectors. The contours represent the free energy landscape inferred by the thermodynamic map, while the shaded regions represent the distribution of structures observed in the simulation, shaded by the number of base pairs. The 10 conformers reported in the Protein Data Bank (PDB), represented as green stars, lie within the most dominant TM-predicted cluster. Fig. 4B shows the learned, temperature-dependent free energy

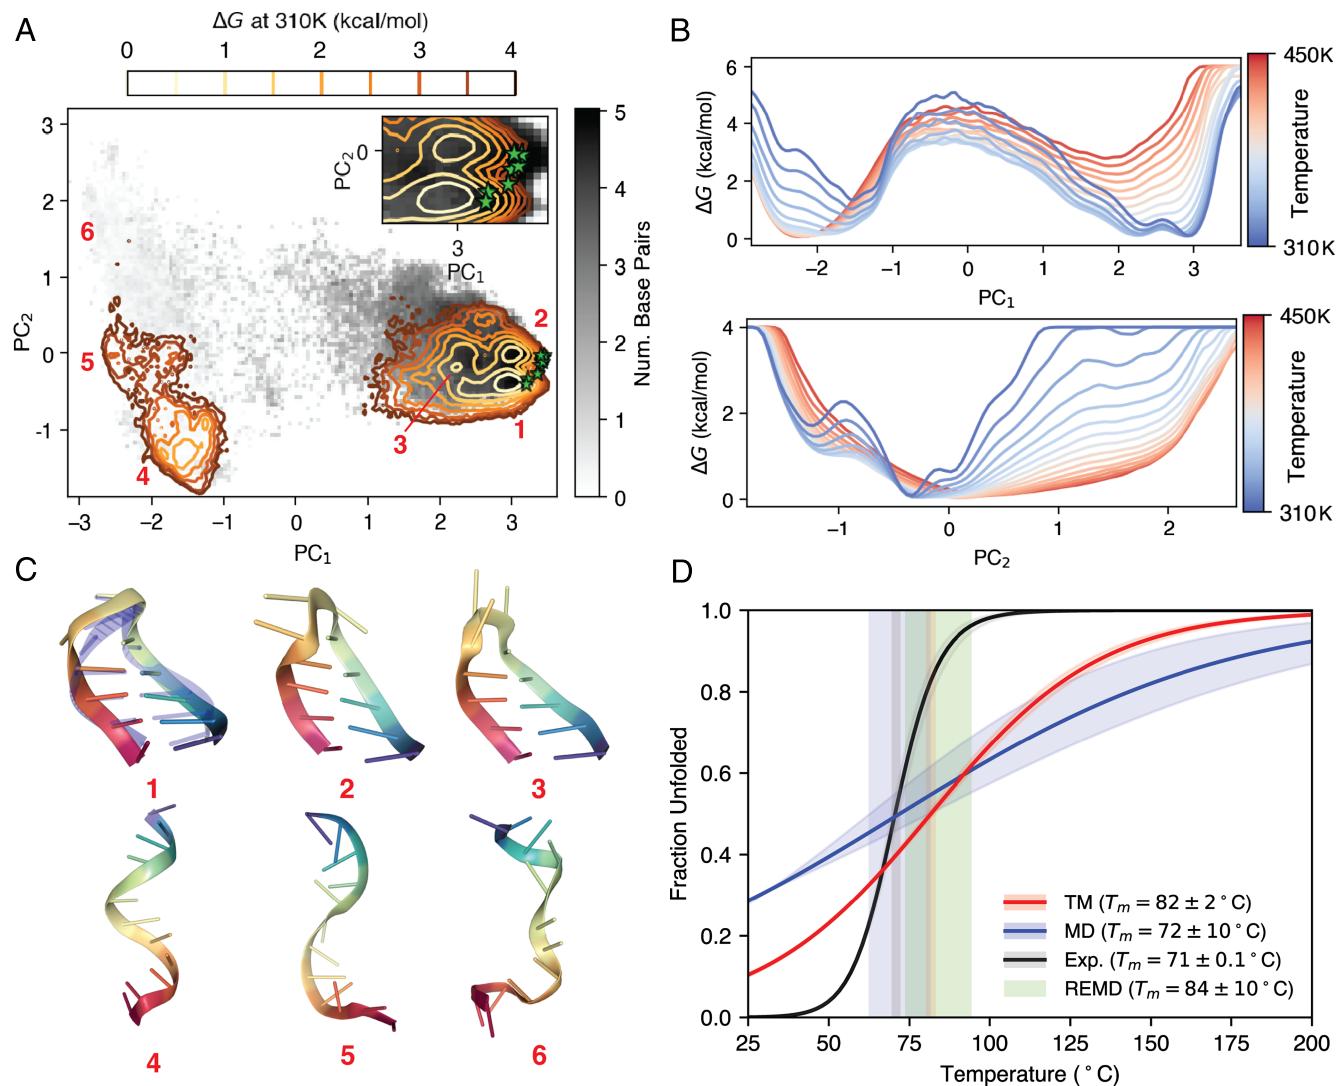


Fig. 4. GCAA Tetraloop Conformational Landscape. (A) Joint distribution of the first two principal components of G-vectors for GCAA. Contours representing TM-generated samples at 310 K are overlaid on MD samples shaded by the number of base pairs, with selected regions of the conformational landscape annotated as 1-6. The 10 NMR conformers reported in the Protein Data Bank (PDB: 1ZIH) are depicted as green stars. (B) Temperature-dependent free energy profiles along the first two principal components shown in panel A. Colors indicate temperatures ranging from 310 K to 450 K at 10 K intervals. (C) Representative structures sampled from labeled minima in panel A, with a representative NMR conformer shown in translucent blue. The structures are colored from blue to red from the 5' to 3' end. (D) A fraction unfolded curve is obtained by reweighting the MD data in panel A with the TM conformational ensemble for each temperature in panel B, and fitting a two-state model. Uncertainties are computed from the last three iterations of TM-aMD. The cutoff for folded and unfolded states is three base pairs to match a reference REMD study (40).

profiles along each of the principal components in Fig. 4A. The first principal component clearly shows the reweighting of the folded and unfolded states with temperature. Fig. 4C shows the structures associated with the minima in Fig. 4A. The first three conformations are abundant at 310 K and are consistent with MD studies (40, 56) and experiment (an example NMR conformer is depicted in blue). The last three represent unfolded states that are stabilized by base stacking and are weakly present in the 310 K ensemble. The first principal component corresponds to the folding-unfolding transition, while the second captures the conformational heterogeneity of the loop region.

Fig. 4D displays the melting curve of the GCAA tetraloop computed from the last three rounds of MD simulation, with the melting curve of the learned equilibrium distribution. Both curves exhibit agreement, matching the range of melting temperatures predicted from a millisecond replica exchange molecular dynamics (REMD) simulation conducted with the same force-field (40). However, both the TM-generated melting curve and the range of melting temperatures from the reference are higher than the experimental melting temperature, indicating that current force-fields do not yet accurately capture the temperature dependence of the tetraloop ensemble.

1.5.2. HIV-TAR RNA. The HIV-TAR RNA is an extensively studied, 29 nucleotide RNA hairpin from the HIV genome that displays rich conformational diversity. The secondary structure consists of lower and upper helices separated by a three-nucleotide bulge with an apical loop closing the upper helix (PDB: 1ANR) (13). The disordered loop and bulge regions mediate interactions with proteins and small molecules (59). We investigate the conformational landscape of the HIV-TAR RNA using six iterations of TM-aMD, requiring a total of 70 μ s of simulation time. We infer relative free energies between the dominant conformer observed through NMR spectroscopy and alternative conformers, which are unattainable through MD simulation alone.

The conformational landscape predicted by the learned TMs, projected along the principal components of the G-vectors is shown in Fig. 5A. Structures corresponding to the NMR ensemble (green stars) are well separated from other misfolded structures within the two-dimensional projection. The temperature dependence of the free energy along each principal component is shown in Fig. 5B, where we find that the free energy barrier separating the NMR conformers from unfolded states reaches a height of 5 kcal/mol. Fig. 5C depicts representative samples from each cluster. The first cluster shows agreement with the NMR

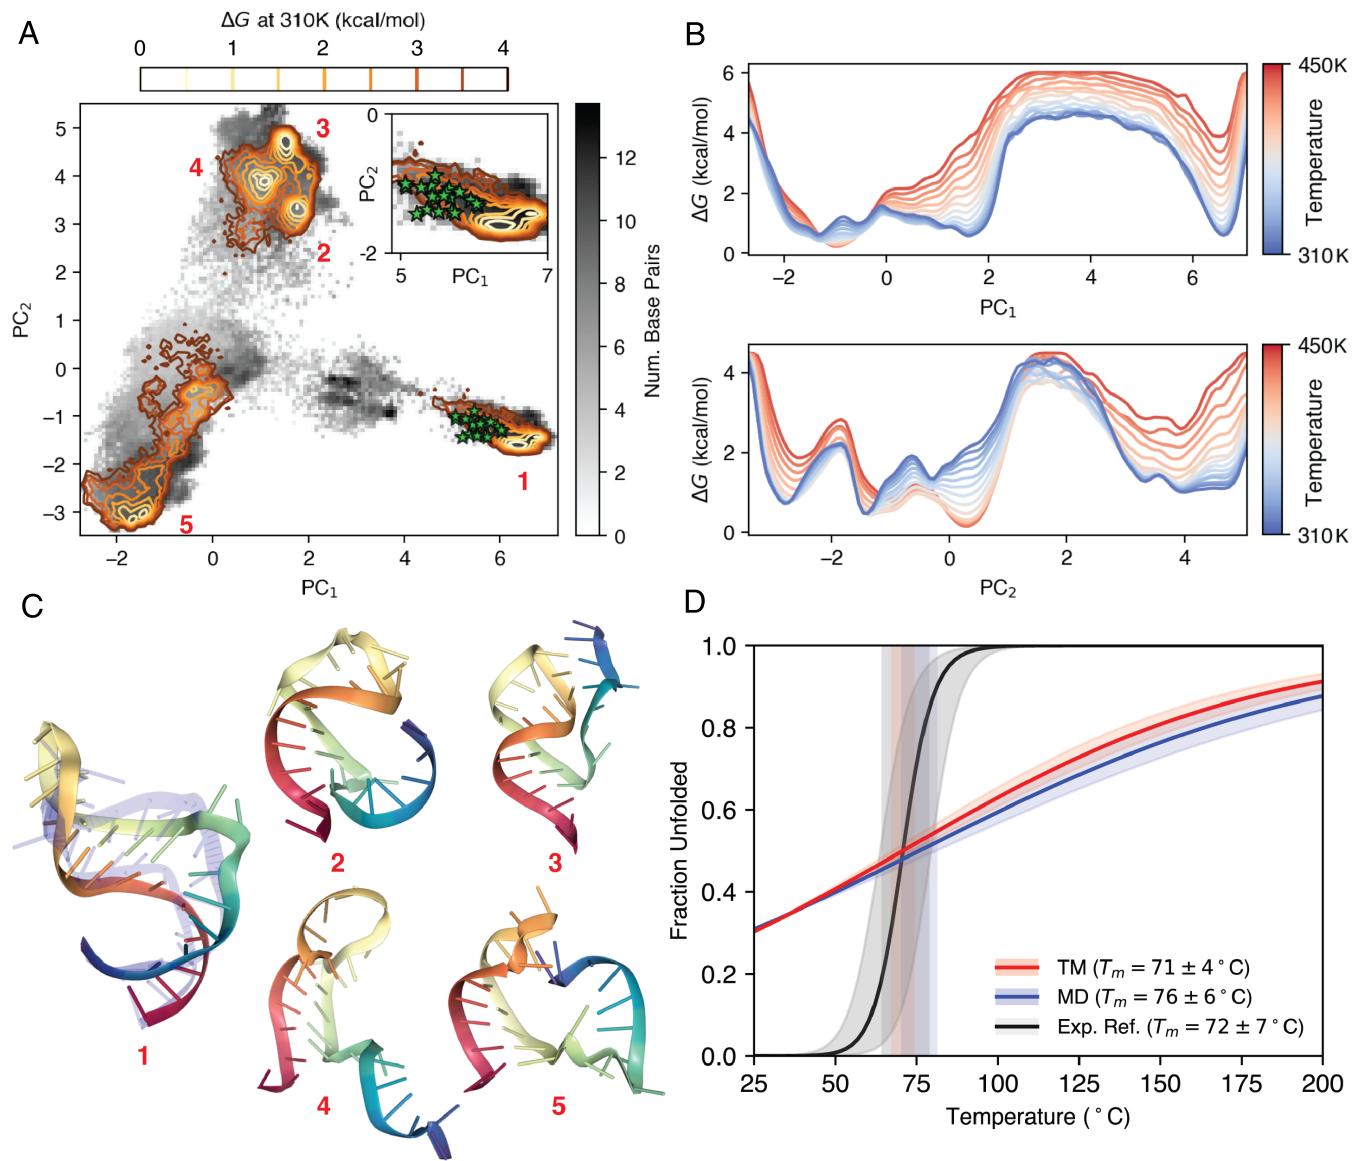


Fig. 5. HIV-TAR RNA conformational landscape. (A) Joint distribution of the first two principal components of G-vectors for HIV-TAR. Contours representing TM-generated samples at 310 K are overlaid on MD samples shaded by the number of base pairs, with basins of the TM landscape labeled 1–5. The 20 reported NMR conformers in the Protein Data Bank (PDB: 1ANR) are depicted as green stars. (B) Temperature-dependent free energy profiles along the first two principal components are shown in panel A. Colors indicate temperatures ranging from 310 K to 450 K at 10 K intervals. (C) Representative structures sampled from labeled clusters in panel A, with a representative NMR conformer shown in translucent blue. The structures are colored from blue at the 5' end to red at the 3' end. (D) A fraction unfolded curve is obtained by reweighting the MD data in panel A with the TM conformational ensemble for each temperature in panel B, and fitting a two-state model. Uncertainties are computed from successive iterations of the TM-aMD algorithm showing agreement between MD and TM-predicted melting temperatures (SI Appendix, Fig. S2B). The cutoff for folded and unfolded states is nine base pairs, as determined by the NMR conformers.

conformers, with the first reported conformer shown in translucent blue. The other clusters correspond to varied secondary structure motifs. Generally, across the conformational landscape, we find that the contribution of folded states to the conformational ensemble diminishes as temperature increases. This can be clearly shown by computing a fraction folded curve (Fig. 5D), which shows agreement between melting temperatures of the MD and TM-derived ensembles, and the experimental melting temperature (60).

Our results support the idea that RNAs evolve over a rugged free energy landscape punctuated by many long-lived states (61). We find that the native state is indeed the most stable at physiological temperatures (310 K), but misfolded states consisting of secondary structure elements still have substantial contributions to the ensemble. Minima of the energy landscape corresponding to states 2 to 5 in Fig. 5 differ from the NMR conformers (state 1) by less than 1 kcal/mol. Overall, our findings are in agreement with theoretical studies of RNA folding pointing towards a rugged energy landscape (44, 61).

Clearly, our findings are dependent on the accuracy of the simulation force field, which for RNAs have not yet reached experimental accuracy. Moreover, TM-aMD is incapable of reaching experimental folding timescales. Despite this, the fraction folded curve for HIV-TAR agrees with experimental results (60). Improved forcefields are out of reach due to the difficulty of attaining equilibrium through simulation; a potential application for TM-aMD may be accelerating convergence to equilibrium, allowing for more rapid evaluation of force-field parameters.

2. Discussion

We have demonstrated that TM are capable of inferring the nontrivial behavior of the free energy with temperature from limited data for the two-dimensional square Ising model. Additionally, we also show how TM may be integrated with

molecular dynamics simulation to enhance sampling. In this section, we outline prospective, potential applications for TM. Our experiments for inferring critical behavior in the Ising model suggest a number of potential research directions that might shed light on some general principles underpinning TM. We outline three directions below.

In this work, we have not explored the mechanism by which TM learn phase transitions. In Section 1.4, as the Ising lattice grows larger the ability to infer the phase transition from only two observations diminishes. Inversely, as the lattice size decreases, the phase transition is learned more consistently. These observations suggest a connection between fluctuations due to finite-size effects and robustness of the algorithm. Spin-glasses with multiple phases present a setting for exploring these connections without the additional complexity molecular dynamics brings.

A second avenue worth exploring is the extent to which TM infer dramatic changes in Boltzmann weights across temperature. Here too, spin-glasses are an appealing system; some spin-glasses exhibit temperature chaos, where, in the thermodynamic limit, the weights change discontinuously with temperature (62). In the framework provided in the introduction, this behavior manifests as a phase transition (i.e., a discontinuity in $Z(\beta)$). Recently, descriptions of temperature chaos have been generalized to finite systems, where it is prevalent on a characteristic length scale that diverges as the change in temperature becomes infinitesimal (63).

Another direction closely related to those outlined above is determining whether TM can capture critical behavior arising from correlation functions. The order parameters explored in this work (magnetization and heat capacity) are measured by averaging over the entire system, but critical exponents are derived from the asymptotic behavior of correlation functions that capture the microscopic details of fluctuations. To measure some critical exponents, computing correlation functions is unavoidable.

In terms of applications, our results in Section 1.5 obtained by TM-aMD indicate that TM are well suited for inferring equilibrium distributions from swarms of short, independent simulations. We show that TM-aMD is capable of exploring the conformational landscape orders of magnitude faster than REMD, and inferring global equilibrium from samples gathered under local equilibrium. Still, TM-accelerated MD stands to benefit greatly from optimization. And since TMs allow for truly parallel simulations across arbitrary temperatures, TMs are poised to make efficient use of distributed computing resources (64).

As pointed out in Section 1.5, we believe that our implementation of TM-aMD is hindered by the scheme used to adaptively resample initial states. The one we employ aims for agreement between the TM and MD-generated distributions. However, in practice, it is more productive to balance thermodynamic

accuracy with sampling of diverse conformations. Methods such as FAST and REAP quantify this trade-off as a reward function which is to be optimized by the choice of adaptive sampling strategy (65, 66).

Finally, as presented, TM are not informative of kinetics since the samples they generate are independent. In principle, one may use the height of the barrier separating states and recrossing related dynamical corrections to estimate rates (67). Still, we do not expect this approach to be useful without some form of further enhanced sampling (68).

Overall, we have presented TM as an exciting way of integrating score-based generative modeling with the theoretical framework of statistical mechanics. We have demonstrated the capability of TM to learn physics by predicting the critical behavior associated with the Ising model despite limited, deliberately misleading sampling. We also use TM to accelerate Molecular Dynamics (TM-aMD) and show that TM-aMD efficiently infers the equilibrium distribution of two model RNA systems using a fraction of the computational resources required by replica exchange simulations. The TM-aMD method, combined with experimentally measured melting curves, might also prove useful for improving force fields for RNA and similar systems. Based on these results, we believe that TM are suitable for widespread use and have great potential for further theoretical development and computational optimization.

Data, Materials, and Software Availability. A Python implementation of thermodynamic maps trained to predict the phase transition of the Ising model has been deposited in GitHub (<https://github.com/lherron2/thermomaps-isng>) (69). All other data are included in the manuscript and/or *SI Appendix*.

ACKNOWLEDGMENTS. This research was supported by the Intramural Research Program of the NIH, National Cancer Institute (NCI), Center for Cancer Research, Project BC011585 09 (J.S.S.). L.H. was supported by the NSF, Grant No. CHE-2044165. J.S.S. and P.T. thank the NCI-UMD Partnership for Integrative Cancer Research. P.T. was an Alfred P. Sloan Foundation fellow during preparation of this manuscript. We thank Zaratan and XSEDE (Project CHE180027P) for computational resources. L.H. would like to thank Yihang Wang and Disha Sanwal for insightful discussions, and Eric Beyerle, Suemin Lee, Giovanni Bussi, Alessandro Laio, and Christopher Jarzynski for critical reading of our manuscript.

Author affiliations: ^aBiophysics Program and Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742; ^bUniversity of Maryland Institute for Health Computing, Bethesda, MD 20852; ^cChemical Biology Laboratory, National Cancer Institute, Frederick, MD 21702; and ^dDepartment of Chemistry and Biochemistry and Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742

Author contributions: L.H., J.S.S., and P.T. designed research; L.H., K.M., and P.T. performed research; L.H., K.M., and P.T. analyzed data; and L.H., J.S.S., and P.T. wrote the paper.

1. D. J. Wales, Exploring energy landscapes. *Annu. Rev. Phys. Chem.* **69**, 401–425 (2018).
2. R. W. Zwanzig, High-temperature equation of state by a perturbation method. I. Nonpolar gases. *J. Chem. Phys.* **22**, 1420–1426 (1954).
3. F. James, Monte Carlo theory and practice. *Rep. Prog. Phys.* **43**, 1145 (1980).
4. S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, P. A. Kollman, The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *J. Comput. Chem.* **13**, 1011–1021 (1992).
5. C. H. Bennett, Efficient estimation of free energy differences from Monte Carlo data. *J. Comput. Phys.* **22**, 245–268 (1976).
6. M. R. Shirts, J. D. Chodera, Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **129** (2008).
7. M. Mézard, G. Parisi, M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific Publishing Company, 1987), vol. 9.
8. J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics" in *International Conference on Machine Learning* (PMLR, 2015), pp. 2256–2265.
9. J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* **33**, 6840–6851 (2020).
10. Y. Song *et al.*, Score-based generative modeling through stochastic differential equations. arXiv [Preprint] (2020). <https://arxiv.org/abs/2011.13456> (Accessed 26 November 2024).
11. Y. Wang, L. Herron, P. Tiwary, From data to noise to data for mixing physics across temperatures with generative artificial intelligence. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2203656119 (2022).
12. F. M. Jucker, H. A. Heus, P. F. Yip, E. H. Moors, A. Pardi, A network of heterogeneous hydrogen bonds in GNRA tetraloops. *J. Mol. Biol.* **264**, 968–980 (1996).
13. F. Aboul-ela, J. Karn, G. Varani, Structure of HIV-1 TAR RNA in the absence of ligands reveals a novel conformation of the trinucleotide bulge. *Nucleic Acids Res.* **24**, 3974–3981 (1996).
14. L. R. Ganser *et al.*, Probing RNA conformational equilibria within the functional cellular context. *Cell Rep.* **30**, 2472–2480 (2020).
15. R. P. Feynman, Slow electrons in a polar crystal. *Phys. Rev.* **97**, 660 (1955).
16. A. Gelman, X. L. Meng, Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Stat. Sci.* **13**, 163–185 (1998).
17. C. Jarzynski, Targeted free energy perturbation. *Phys. Rev. E* **65**, 046122 (2002).

18. P. Wirnsberger *et al.*, Targeted free energy estimation via learned mappings. *J. Chem. Phys.* **153**, 144112 (2020).
19. P. Wirnsberger *et al.*, Normalizing flows for atomic solids. *Mach. Lear. Sci. Tech.* **3**, 025009 (2022).
20. I. Kobyzev, S. J. Prince, M. A. Brubaker, Normalizing flows: An introduction and review of current methods. *IEEE Trans.* **43**, 3964–3979 (2020).
21. L. Dinh, J. Sohl-Dickstein, S. Bengio, Density estimation using real NVP. arXiv [Preprint] (2016). <https://arxiv.org/abs/1605.08803> (Accessed 26 November 2024).
22. R. T. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations. *Adv. Neural Inf. Process. Syst.* **31**, 6571–6583 (2018).
23. A. K. Dhaka *et al.*, Challenges and opportunities in high dimensional variational inference. *Adv. Neural Inf. Process. Syst.* **34**, 7787–7798 (2021).
24. J. Köhler, A. Krämer, F. Noé, Smooth normalizing flows. *Adv. Neural Inf. Process. Syst.* **34**, 2796–2809 (2021).
25. C. Durkan, A. Bekasov, I. Murray, G. Papamakarios, Neural spline flows. *Adv. Neural Inf. Process. Syst.* **32**, 7511–7522 (2019).
26. L. I. Midgley, V. Stimpf, G. N. Simm, B. Schölkopf, J. M. Hernández-Lobato, Flow annealed importance sampling bootstrap. arXiv [Preprint] (2022). <https://arxiv.org/abs/2208.01893> (Accessed 26 November 2024).
27. H. Wu, J. Köhler, F. Noé, Stochastic normalizing flows. *Adv. Neural Inf. Process. Syst.* **33**, 5933–5944 (2020).
28. A. Hyvärinen, P. Dayan, Estimation of non-normalized statistical models by score matching. *J. Mach. Learn. Res.* **6**, 695–709 (2005).
29. Y. Song, S. Ermon, Generative modeling by estimating gradients of the data distribution. *Adv. Neural Inf. Process. Syst.* **32**, 11895–11907 (2019).
30. D. Maoutsa, S. Reich, M. Opper, Interacting particle solutions of Fokker-Planck equations through gradient-log-density estimation. *Entropy* **22**, 802 (2020).
31. N. M. Boffi, E. Vanden-Eijnden, Probability flow solution of the Fokker-Planck equation. *Mach. Learn. Sci. Technol.* **4**, 035012 (2023).
32. H. Risken, H. Risken, *Fokker-Planck equation* (Springer, 1996).
33. B. D. Anderson, Reverse-time diffusion equation models. *Stoc. Proc. Appl.* **12**, 313–326 (1982).
34. F. Noé, S. Olsson, J. Köhler, H. Wu, Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science* **365**, eaaw1147 (2019).
35. C. Luo, Understanding diffusion models: A unified perspective. arXiv [preprint] (2022). <https://arxiv.org/abs/2208.11970> (Accessed 1 December 2023).
36. K. Dill, S. Bromberg, *Molecular Driving Forces: Statistical Thermodynamics* (Garland Science, 2010).
37. S. M. Bhattacharjee, A. Khare, Fifty years of the exact solution of the two-dimensional Ising model by Onsager. *Curr. Sci.* **69**, 816–821 (1995).
38. J. Cardy, *Scaling and Renormalization in Statistical Physics* (Cambridge University Press, 1996), vol. 5.
39. U. Wolff, Critical slowing down. *Nucl. Phys. B-Proc. Suppl.* **17**, 93–102 (1990).
40. D. Tan, S. Piana, R. M. Dirks, D. E. Shaw, RNA force field with accuracy comparable to state-of-the-art protein force fields. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E1346–E1355 (2018).
41. P. Kuhrova *et al.*, Computer folding of RNA tetraloops: Identification of key force field deficiencies. *J. Chem. Theor. Comput.* **12**, 4534–4548 (2016).
42. H. P. Hsu, K. Kremer, Glass transition temperature of (ultra-)thin polymer films. *J. Chem. Phys.* **159**, 071104 (2023).
43. A. Pagnani, G. Parisi, F. Ricci-Tersenghi, Glassy transition in a disordered model for the RNA secondary structure. *Phys. Rev. Lett.* **84**, 2026 (2000).
44. D. Chakraborty, R. Collepardo-Guevara, D. J. Wales, Energy landscapes, folding mechanisms, and kinetics of RNA tetraloop hairpins. *J. Am. Chem. Soc.* **136**, 18052–18061 (2014).
45. A. T. Frank, Q. Zhang, H. M. Al-Hashimi, I. Andricioaei, Slowdown of interhelical motions induces a glass transition in RNA. *Biophys. J.* **108**, 2876–2885 (2015).
46. A. Banerjee, A. Iscen, K. Kremer, O. Kukharenko, Determining glass transition in all-atom acrylic polymeric melt simulations using machine learning. *J. Chem. Phys.* **159**, 074108 (2023).
47. J. Lee, E. A. Dethoff, H. M. Al-Hashimi, Invisible RNA state dynamically couples distant motifs. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 9485–9490 (2014).
48. A. C. Stelzer *et al.*, Discovery of selective bioactive small molecules by targeting an RNA dynamic ensemble. *Nat. Chem. Biol.* **7**, 553–559 (2011).
49. M. E. Fisher, W. Selke, Infinitely many commensurate phases in a simple Ising model. *Phys. Rev. Lett.* **44**, 1502 (1980).
50. J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554–2558 (1982).
51. D. J. Amit, H. Gutfreund, H. Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* **55**, 1530 (1985).
52. Q. Vicens, J. S. Kieft, Thoughts on how to think (and talk) about RNA structure. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2112677119 (2022).
53. O. Valsson, P. Tiwary, M. Parrinello, Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint. *Annu. Rev. Phys. Chem.* **67**, 159–184 (2016).
54. R. Das, D. Baker, Macromolecular modeling with Rosetta. *Annu. Rev. Biochem.* **77**, 363–382 (2008).
55. A. Oxenfarth *et al.*, Integrated NMR/molecular dynamics determination of the ensemble conformation of a thermodynamically stable CUUG RNA tetraloop. *J. Am. Chem. Soc.* **145**, 16557–16572 (2023).
56. G. H. Zerze, P. M. Piaggi, P. G. Debenedetti, A computational study of RNA tetraloop thermodynamics, including misfolded states. *J. Phys. Chem. B* **125**, 13685–13695 (2021).
57. J. P. Sheehy, A. R. Davis, B. M. Znosko, Thermodynamic characterization of naturally occurring RNA tetraloops. *RNA* **16**, 417–429 (2010).
58. S. Bottaro, F. Di Palma, G. Bussi, The role of nucleobase interactions in RNA structure and dynamics. *Nucleic Acids Res.* **42**, 13306–13314 (2014).
59. S. S. Chavali, S. M. Mali, J. L. Jenkins, R. Fasan, J. E. Wedekind, Co-crystal structures of HIV TAR RNA bound to lab-evolved proteins show key roles for arginine relevant to the design of cyclic peptide TAR inhibitors. *J. Biol. Chem.* **295**, 16470–16486 (2020).
60. A. L. Smith, J. Kassman, K. J. Sour, A. M. Soto, Effect of salt concentration on the conformation of TAR RNA and its association with aminoglycoside antibiotics. *Biochemistry* **50**, 9434–9445 (2011).
61. S. J. Chen, K. A. Dill, RNA folding energy landscapes. *Proc. Natl. Acad. Sci. U.S.A.* **97**, 646–651 (2000).
62. F. Ritort, Static chaos and scaling behavior in the spin-glass phase. *Phys. Rev. B* **50**, 6844 (1994).
63. M. Baity-Jesi *et al.*, Temperature chaos is present in off-equilibrium spin-glass dynamics. *Commun. Phys.* **4**, 74 (2021).
64. A. L. Beberg, D. L. Ensign, G. Jayachandran, S. Khalil, V. S. Pande, “Folding@ home: Lessons from eight years of volunteer distributed computing” in *2009 IEEE International Symposium on Parallel & Distributed Processing* (IEEE, 2009), pp. 1–8.
65. M. I. Zimmerman, G. R. Bowman, Fast conformational searches by balancing exploration/exploitation trade-offs. *J. Chem. Theory Comput.* **11**, 5747–5757 (2015).
66. Z. Shamsi, K. J. Cheng, D. Shukla, Reinforcement learning based adaptive sampling: Reaping rewards by exploring protein conformational landscapes. *J. Phys. Chem. B* **122**, 8386–8395 (2018).
67. B. J. Berne, M. Borkovec, J. E. Straub, Classical and modern methods in reaction rate theory. *J. Phys. Chem.* **92**, 3711–3725 (1988).
68. P. Tiwary, M. Parrinello, From metadynamics to dynamics. *Phys. Rev. Lett.* **111**, 230602 (2013).
69. L. Herron, K. Mondal, J. S. Schneekloth Jr., P. Tiwary, thermomaps-isng. GitHub. <https://github.com/lherron2/thermomaps-isng>. Deposited 1 February 2024.