


# Optimal Control of Nonequilibrium Systems through Automatic Differentiation

Megan C. Engel<sup>\*</sup>*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA  
and Department of Biological Sciences, University of Calgary, Calgary, Canada, T2N 1N4*

Jamie A. Smith

*Google Research, Mountain View, California 94043, USA*Michael P. Brenner<sup>†</sup>*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA* (Received 18 May 2023; revised 8 September 2023; accepted 17 October 2023; published 16 November 2023)

Controlling the evolution of nonequilibrium systems to minimize dissipated heat or work is a key goal for designing nanodevices, in both nanotechnology and biology. Progress in computing optimal protocols to extremize thermodynamic variables has, thus far, been limited to either simple systems or near-equilibrium evolution. Here, we present an approach for computing optimal protocols based on automatic differentiation. Our methodology is applicable to complex systems and multidimensional protocols and is valid arbitrarily far from equilibrium. We validate our method by reproducing theoretical optimal protocols for a Brownian particle in a time-varying harmonic trap. We also compute departures from near-equilibrium behavior for magnetization reversal on an Ising lattice and for barrier crossing driven by a harmonic trap, which is used to represent a range of biological processes including biomolecular unfolding reactions. Algorithms based on automatic differentiation outperform the near-equilibrium theory for far-from-equilibrium magnetization reversal and for driven barrier crossing beyond the linear regime. The optimal protocol for far-from-equilibrium driven barrier crossing is found to hasten the approach to, and slow the departure from, the barrier region compared to the near-equilibrium theoretical protocol. We demonstrate the utility of our method in a real-world use case by reducing the work required to unfold a DNA hairpin in the coarse-grained oxDNA model and improving its nonequilibrium free-energy landscape reconstruction compared to a naive linear protocol.

DOI: [10.1103/PhysRevX.13.041032](https://doi.org/10.1103/PhysRevX.13.041032)Subject Areas: Biological Physics,  
Computational Physics,  
Statistical Physics

## I. INTRODUCTION

The control of nonequilibrium phenomena at microscopic scales is central to biology and nanotechnology. Evolution has exquisitely tuned cellular processes to perform out-of-equilibrium tasks, ranging from machines like ATP synthase to metabolic factories converting raw materials and energy into functional macromolecules. Experimental advances allow phenomena on this scale to be probed in unprecedented detail [1–3], but determining precisely how specific processes work and the role of evolutionary optimization remains a major challenge. And

while impressive progress has already been made engineering synthetic DNA [4,5] and protein [6,7] structures, we do not understand how to design *de novo* nanomachines for nonequilibrium environments well enough for nanotechnology to rival the complexity of cellular machines.

For a microscopic system evolving out of equilibrium, thermodynamic quantities like entropy, heat, and work can be meaningfully defined only at the level of individual trajectories [8]. An ensemble of trajectories has distributions of thermodynamic properties with forms that depend on the system's nonequilibrium evolution [8,9]. This immediately suggests an optimization problem whereby a protocol  $\lambda(t)$  drives a system between given initial and final states to produce a desired distribution of thermodynamic properties. A common aim is to minimize average dissipated work. This is important for optimal bit erasure [10–12], as well as experimental measurements of the equilibrium free energy of biomolecules from nonequilibrium force pulling experiments and simulations [9,13,14].

---

<sup>\*</sup>Corresponding author: [megan.engel@ucalgary.ca](mailto:megan.engel@ucalgary.ca)

*Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.*

Other targets include protocols that maximize thermodynamic efficiency for synthetic [15,16] and biological [17] nanoengines and protocols that minimize dissipated heat, for example, in bit-flipping operations [18,19].

Modeling nonequilibrium processes is notoriously difficult, even when the equations of motion are precisely known. A general method to optimize nonequilibrium driving protocols valid for systems of any complexity evolving arbitrarily far from equilibrium has heretofore not been elucidated. In the context of extremizing thermodynamic variables like entropy and work by tuning external parameter protocols, most existing work either has been limited to extremely simple systems, such as a Brownian particle diffusing in a harmonic well or a single quantum dot [20–23], or applies only in the near-equilibrium regime [24–26]. An exception is the recent work of Whitelam [27], which treats two of the examples we consider here through an alternative machine-learning-inspired approach. We discuss the relative benefits and drawbacks of his method in relation to ours in Sec. IV. The assumption of near-equilibrium evolution restricts optimal protocols to free-energy landscapes with low-energy barriers, ruling out many systems of interest, such as RNA molecules with pseudoknots, proteins, and biomolecular motors like ATP synthase.

Other work optimizing nonequilibrium processes has focused on colloidal self-assembly. Goodrich *et al.* present a method for tuning interparticle potentials to manipulate crystal structure and transition rates [28] but do not consider processes where external parameters are continuously manipulated. Tang *et al.* [29,30] and Grover *et al.* [31] present methods for adjusting external parameters to minimize crystalline defects based on feedback-informed Markov state modeling (MSM). While they achieve impressive experimental results increasing colloidal crystal order for hundreds of particles, these rely on computationally costly dynamic programming calculations. Recent work by Trubiano and Hagan circumvents this issue by incorporating gradient information but acknowledges their approach will scale exponentially with large numbers of external control parameters [32]. Additionally, MSMs are constructed for a few discrete values of external control parameter, limiting the complexity of protocols that can be discovered via these methods. Other Monte Carlo-based optimization techniques preserve greater protocol complexity by linearly interpolating between discrete control parameter values [33], but in this case accuracy is limited by how many discrete values are sampled. Whitelam *et al.* recently presented a novel approach to tuning crystal formation based on evolutionary reinforcement learning, using neural networks to parametrize external protocols [34]. Evolutionary reinforcement learning does not use information about loss landscape gradients, however, rendering convergence slower than for methods that incorporate such information.

Inspired by recent computational advances in the machine-learning community, we propose a method for computing optimal nonequilibrium protocols that is valid for complex systems evolving far from equilibrium. In particular, we leverage automatic differentiation (AD) [35–38], a technique for computing gradients that repeatedly applies the chain rule to elementary computational steps. AD optimization has been recently applied in a range of scientific contexts, from quantum devices to self-assembly [28,39]. Using efficient AD algorithms developed in the context of training neural networks [40–42] in conjunction with sophisticated graphical processing unit (GPU) and tensor processing unit (TPU) hardware, we compute gradients by *backpropagating* through entire simulations, allowing us to find optimal protocols via gradient descent for a variety of systems. Using explicit gradient information allows our AD approach to converge faster than methods that do not incorporate it and provides richer information about loss landscapes. Furthermore, our method can be applied to arbitrarily complex external protocols. The computational cost of automatic differentiation is *independent* of the number of protocol parameters in so-called “reverse mode” and scales linearly—not exponentially—with the number of parameters in “forward mode” [43]. Our method, thus, scales well to complex, multidimensional external protocols.

To illustrate the potential of this method, we here consider three canonical examples from the optimal protocol literature and drive evolution much farther out of equilibrium than previously possible. First, we consider Monte Carlo (MC) simulations and use AD to derive optimal protocols for flipping the magnetization of a 2D Ising lattice. The AD protocols perform similarly to existing near-equilibrium theoretical results in the linear regime and outperform the near-equilibrium theory in the far-from-equilibrium regime. Next, we treat molecular dynamics (MD) simulations, reproducing classic analytical results for a single Brownian particle in a time-varying harmonic potential. We then examine barrier crossing on a double-well potential driven by a moving harmonic potential, a problem that maps onto biomolecular unfolding processes [44,45]. After recreating existing results, we probe the far-from-equilibrium regime of barrier crossing, demonstrating the capability of our method to capture departures from the near-equilibrium optimal protocols. Finally, we apply JAX-MD optimized biomolecular unfolding protocols to the complex problem of biomolecular free-energy landscape reconstruction. Reconstructing free-energy landscapes from nonequilibrium data is a key goal of many experimental [46–49] and computational [50,51] studies, and much work has been devoted to improving the accuracy of nonequilibrium free-energy landscape reconstructions [13,52,53]. By improving the reconstructed energy landscape of a DNA hairpin, we demonstrate that

our method constitutes a valuable and practicable addition to these efforts.

## II. RESULTS

### A. Differentiation of MC simulations: Nanomagnetic spin systems

Computers dissipate large amounts of heat when performing logical operations via bit erasure, which reverses nanomagnetic spins [54,55]. This has motivated recent studies investigating minimum-dissipation protocols for magnetization reversal with the 2D Ising model [18,19,33]. The system is described by the Hamiltonian

$$H[B(t)] = -\sum_{\langle i,j \rangle} J_{ij} \sigma_i \sigma_j - B(t) \sum_i \sigma_i, \quad (1)$$

where  $\sigma_i = \pm 1$  are the spins,  $\langle ij \rangle$  indicates a sum over all nearest-neighbor spins,  $J_{ij}$  is the coupling between spins  $i$  and  $j$ , and  $B(t)$  is the (time-dependent) external magnetic field.

In the linear response (near-equilibrium) regime, Crooks and co-workers developed a general formalism for computing optimal protocols based on thermodynamic geometry [24,56]. Rotskoff and Crooks applied this theory to yield the theoretical optimal protocol for varying external magnetic field  $B$  and spin-spin coupling strength  $J$  (equivalent to varying temperature) to reverse magnetization on an Ising lattice [18]. More recently, Gingrich *et al.* [33] explored the same problem using a numeric approach, in which the space of low dissipation protocols is explored with a Monte Carlo scheme. This yields a number of degenerate, near-optimal protocols, but like the work of Rotskoff and Crooks [18] is limited by the assumption of near-equilibrium evolution.

Inspired by previous work, we examine the nonequilibrium magnetization reversal of a 2D periodic lattice of spins driven by a protocol that varies both magnetic field  $B(t)$  and temperature  $T(t)$ :  $\lambda(t) = [B(t), T(t)]$  but push evolution beyond the near-equilibrium regime, benchmarking against the linear response formalism of Rotskoff and Crooks [18].

We seek to minimize the total entropy production  $\Delta S$ , a proxy for the heat dissipated to the environment during the “bit flip” that quantifies the irreversibility of the process [9,33,57–59]:

$$\Delta S = k_B \ln \frac{P^F[\mathbf{x}]}{P^R[\bar{\mathbf{x}}]}. \quad (2)$$

Here,  $P^F[\mathbf{x}]$  is the probability of observing a particular trajectory  $\mathbf{x}$  during the forward evolution of a system, and  $P^R[\bar{\mathbf{x}}]$  is the probability of observing the exact time reversal of that trajectory,  $\bar{\mathbf{x}}$ .

To find the optimal protocol for varying  $B(t)$  and  $T(t)$ , we write a Monte Carlo simulator using JAX [42], a PYTHON library with built-in automatic differentiation and just-in-time compilation. The code carries out standard Glauber dynamics [60] and iteratively updates the grid points with even, then odd lattice index  $i$  using the spin flip probability

$$P(\sigma_i \rightarrow -\sigma_i) = \frac{1}{1 + e^{\beta \Delta E_i}}, \quad (3)$$

where  $\beta = 1/k_B T$  is the usual inverse thermal energy and the change in lattice energy resulting from the flip of spin  $\sigma_i$  can be computed using the sum of its nearest-neighbor spins  $\sum_j \sigma_j$ :

$$\Delta E_i = 2\sigma_i \left( J_{ij} \sum_j \sigma_j - B \right). \quad (4)$$

Our code compiles to run rapidly on GPUs or TPUs and is differentiable: Given a Monte Carlo trajectory of spins under some protocol  $\lambda(t)$ , we can compute the gradient of any function of the trajectory with respect to the protocol. This gradient can be computed in either forward mode or reverse mode. In forward mode, derivatives are computed and stored at each step in the computation, obviating the need to store intermediate states. In reverse mode—commonly known as “backpropagation” in the deep learning community—the computation proceeds in two passes: one forward evaluation and then a backward traversal of the computational graph to compute derivatives. It is, thus, necessary to hold all intermediate states in memory for reverse-mode AD, which can be a significant memory constraint for large systems. Generally, reverse mode is favored for systems with large numbers of inputs and few outputs (as is the case in typical neural network training), as it is more computationally efficient in these cases [43], while forward mode is a viable option for large, memory-intensive computations.

Our Ising optimizations on larger lattices exhaust memory in reverse-mode differentiation, so we opt to compute gradients with forward-mode differentiation in this case. We find it most convenient to parametrize the protocol using a Chebyshev polynomial basis:

$$\lambda(t) = \theta_0 T_0(t) + \theta_1 T_1(t) + \dots + \theta_k T_k(t), \quad (5)$$

where the  $T_i$  is the  $i$ th Chebyshev polynomial and  $k$  is a hyperparameter (typically, we choose  $k = 32$ ), but many other parametrizations are possible. When selecting a parametrization, a balance must be achieved between sufficient complexity to describe the optimal protocol and overfitting. For the problems we consider in this work, we begin by choosing a simple parametrization (e.g., piecewise linear) and then increase complexity until we

obtain a smoothly decreasing loss function and convergence in optimal protocol shape. For our Ising model optimizations, piecewise linear parametrizations are insufficient; losses do not decrease smoothly and protocol shapes do not converge. We try Chebyshev polynomials next and adjust the number of terms until we obtain convergence.

The loss function (2) is computed for our trajectories as follows. Each time step in the forward evolution can be formulated as a sequence of two substeps: First, the external protocol parameters are updated (this is where external work is performed), and then the system performs a transition to a new microstate (this is where heat is exchanged with the bath) [57–59]. If  $\rho_A(x_0)$  is the probability of drawing microstate  $x_0$  from the equilibrium distribution corresponding to the initial value of the control parameter  $\lambda(t=0) = \lambda_0$ , which we call system state  $A$ , then the probability of observing the forward trajectory is given by

$$P^F[\mathbf{x}] = \rho_A(x_0)P(x_0 \rightarrow x_1; \lambda_1) \dots P(x_{N-1} \rightarrow x_N; \lambda_N), \quad (6)$$

where  $P(x_{i-1} \rightarrow x_i; \lambda_i)$  is the transition probability between lattice states  $x_{i-1}$  and  $x_i$  at protocol parameter values  $\lambda_i$ . Correspondingly, the probability of observing the time-reversed trajectory  $\bar{\mathbf{x}}$  is

$$P^R[\bar{\mathbf{x}}] = \rho_B(x_N)P(x_N \rightarrow x_{N-1}; \lambda_N) \dots P(x_1 \rightarrow x_0; \lambda_1). \quad (7)$$

Formulating the evolution as a succession of accepted and rejected spin flips and noting that for the Glauber transition probability (3),  $1 - P(\Delta E) = P(-\Delta E)$ , with  $\Delta E$  given by Eq. (4), we can combine Eqs. (6) and (7) to obtain

$$\frac{P^F[\mathbf{x}]}{P^R[\bar{\mathbf{x}}]} = \frac{\rho_A(x_0) \prod_i P(\Delta E_i) \prod_j P(-\Delta E_j)}{\rho_B(x_N) \prod_i P(-\Delta E_i) \prod_j P(-\Delta E_j)}. \quad (8)$$

Here, the products containing  $i$  include all spins that flip successfully and those containing  $j$  are failed spin flips. The probability for the  $j$ th spin *not* to flip is the same in the forward and reverse trajectories, since the  $\Delta E_i$  term [Eq. (4)] evaluates identically in each case. The probability of *accepting* the  $i$ th spin flip differs in the sign of  $\Delta E_i$  for the forward and reverse trajectories. Note that this does not violate normalization of probability, since the products  $\prod_i$  and  $\prod_j$  run over *different* spins. The ratio of probabilities of drawing the initial and final states from their respective equilibrium ensembles,  $\rho_A(x_0)/\rho_B(x_N)$ , is given by

$$\frac{\rho_A(x_0)}{\rho_B(x_N)} = \frac{Z_B e^{-\beta_0 H(x_0, \lambda_0)}}{Z_A e^{-\beta_N H(x_N, \lambda_N)}}, \quad (9)$$

where  $Z_B/Z_A = 1$  since, by construction, we set the magnitude of the external field identically (and equal to 1) in the initial and final states of our simulations.

Plugging Eq. (3) into Eq. (8) and rearranging Eq. (2), we find that the entropy production in our simulations is given by

$$\frac{\Delta S}{k_B} \equiv \omega = \beta_N H(x_N, \lambda_N) - \beta_0 H(x_0, \lambda_0) - \sum_i \beta_i \Delta E_i, \quad (10)$$

where we take a sum of the system energy changes  $\Delta E_i$  following successful spin flips  $i$ , multiplied by the inverse temperature  $\beta_i$  at which the flip occurred.

We use JAX's automatic differentiation to compute the gradient  $\nabla_\theta \langle \omega \rangle$  of this entropy production. Because of the discrete choices inherent in Monte Carlo simulations, care must be exercised in computing gradients. In particular, the entropy production depends on external parameters through discrete spin flip operations—dictated by a Glauber acceptance criterion—which are not themselves differentiable; a similar issue arises in the context of training stochastic neural networks [61]. Instead of directly backpropagating through the loss function, we proceed as follows. The average entropy production over all possible trajectories  $\mathbf{x}(t)$  for external protocol parameters  $\theta$  is given by (see also Ref. [62])

$$\langle \omega(\theta) \rangle = \int \mathcal{D}\mathbf{x}(t) P[\mathbf{x}(t); \theta] \omega[\mathbf{x}(t), \theta], \quad (11)$$

where  $\int \mathcal{D}\mathbf{x}(t)$  is an integration over all possible trajectories,  $P[\mathbf{x}(t); \theta]$  is the probability weight associated with each trajectory, and  $\omega[\mathbf{x}(t), \theta]$  is the total entropy production (a proxy for heat dissipation) for each trajectory.

Applying the product rule and noting that  $\nabla P = P \nabla \ln P$ , the gradient of  $\langle \omega(\theta) \rangle$  is

$$\nabla_\theta \langle \omega(\theta) \rangle = \langle (\nabla_\theta \ln P) \omega \rangle + \langle \nabla_\theta \omega \rangle. \quad (12)$$

Note that merely averaging the gradient  $\nabla_\theta \omega$  over a batch of simulated trajectories does *not* yield the correct average, since the probability of observing a trajectory is itself dependent on protocol parameters  $\theta$ . A similar approach to finding stochastic gradients is used for taking gradients in the REINFORCE algorithm [63], widely used in reinforcement learning.

We carry out  $N$  Monte Carlo simulations of the 2D Ising lattice evolving under assumed protocol  $\lambda(\theta)$ . For each trajectory, we compute the gradients  $\nabla_\theta \ln P$  and  $\nabla_\theta \omega$ . Plugging these into Eq. (12) and averaging gives us the correct, REINFORCE estimate of the gradient,  $\nabla_\theta \langle \omega(\theta) \rangle$ . We note that we are unable to train our system successfully without using this REINFORCE approach. We then use the Adam optimizer [64] to minimize the loss [Eq. (10)].

Figure 1(a) shows our lowest achieved average entropy production  $\langle \omega \rangle$  from AD protocols alongside  $\langle \omega \rangle$  for the near-equilibrium theoretical protocol of Rotskoff and Crooks [18] for seven different simulation lengths on a

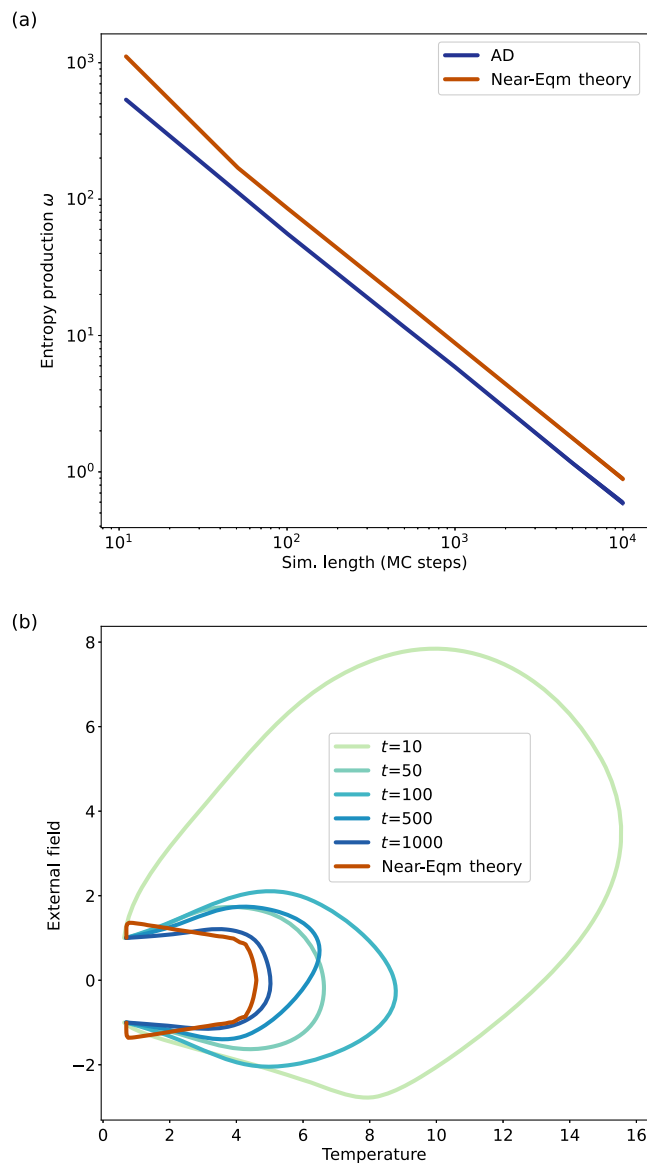


FIG. 1. (a) Minimum entropy production for different simulation lengths on a  $32 \times 32$  lattice, averaged over 2560 trajectories (standard errors of the mean are smaller than the linewidths). The AD protocols in each case are derived using the Adam optimizer [64] with gradients averaged over batches of  $N = 256$ . AD protocols are able to outperform the optimal entropy production that results from using the near-equilibrium theoretical protocol of Rotskoff and Crooks [18] and approach the minimum entropy production value of Rotskoff and Crooks [18] as the near-equilibrium regime is approached. (b) Optimal protocols for reversing magnetization for five simulation lengths:  $t = 10, 50, 100, 500,$  and  $1000$ . Forward evolution in time occurs along the counterclockwise direction. Curves for  $t = 5000$  and  $t = 10000$  are nearly identical to the  $t = 1000$  curve and are not shown. That the AD curves do not converge to the shape of the near-equilibrium theoretical curve, yet produce lower entropy in (a), suggests a flat loss landscape in the optimal region in the near-equilibrium regime, in accordance with previous findings [18,33].

$32 \times 32$  lattice ( $t = 10, 50, 100, 500, 1000, 5000,$  and  $10000$  MC time steps). The longer the simulation, the closer the magnetization reversal is to quasistatic.  $\langle \omega \rangle$  values are averages over 2560 trajectories. In all cases, AD outperforms the near-equilibrium theory. We repeat the  $t = 10, t = 50, t = 100, t = 500,$  and  $t = 1000$  simulations on multiple lattice sizes up to  $512 \times 512$  and find that the normalized entropy production on the  $32 \times 32$  lattice is within 1% of its converged, infinite-lattice value; see Supplemental Material, Fig. S2 [65]. The AD-derived protocols also outperform the near-equilibrium theory regardless of lattice size. This suggests AD optimization can be effectively performed on smaller lattices to find optimal protocols for larger lattices.

Figure 1(b) shows the optimal protocols corresponding to the simulation lengths in Fig. 1(a) along with the near-equilibrium theoretical curve of Rotskoff and Crooks [18]. Curves for  $t = 5000$  and  $t = 10000$  are omitted as they are near identical to the  $t = 1000$  curve. While the near-equilibrium theoretical protocol is necessarily time symmetric [18], our protocols appear to break this symmetry; see Supplemental Material, Fig. S4 [65]. Like the near-equilibrium theoretical result, the AD protocols avoid the critical phase transition region [18], but they do not appear to be converging to the exact shape of the near-equilibrium theoretical curve as equilibrium is approached. We also observe a flat loss function for  $t = 5000$  and  $t = 10000$  over 1000 optimization iterations, as shown in Supplemental Material, Fig. S1 [65]. The fact that our curves perform comparably to the near-equilibrium theoretical curve in the near-equilibrium limit, but are differently shaped, suggests that the entropy production landscape is relatively flat in the region of optimal protocols and that there is a degenerate space of nearly optimal solutions. Indeed, all of the AD protocols in Fig. 1(b) perform comparably well at longer simulation lengths; see Supplemental Material, Fig. S5 [65]. This is in agreement with the findings of Gingrich *et al.* [33] and noted by Rotskoff and Crooks [18], who predict “weakly constrained” protocols in the noncritical region of  $[B(t), T(t)]$  space.

## B. Optimal protocols for Brownian dynamics

We now consider the MD of isothermal evolution of Brownian particles, where the total entropy production is equal to the dissipated external work ( $W_D$ ) [9]. We use the closely related *total* external work supplied as our loss function in the following case studies. Details about our MD implementation can be found in Supplemental Material [65].

### 1. Brownian particle in a harmonic potential

Some of the earliest work identifying optimal non-equilibrium protocols focused on the paradigmatic colloidal Brownian particle in a harmonic trap [8,20,22,62,66].

Exact optimal protocols for varying the stiffness of the center and stiffness of the harmonic potential were found by Schmiedl and Seifert [22] using variational calculus. Strikingly, the solutions have discrete “jumps” in the parameters [62,67]. We note that approaches based on the linear response approximation are incapable of discovering these jumps, since they assume protocols are differentiable [56,66].

We reproduce these early results by using JAX-MD [68] to automatically differentiate molecular dynamics simulations of a colloid subjected to the moving harmonic potential

$$V(x, t) = \frac{1}{2}[x - \lambda(t)]^2, \quad (13)$$

where here  $\lambda(t)$  is the time-dependent position of the trap. We seek a protocol that minimizes the total work done in moving the trap from  $\lambda_i$  to  $\lambda_f$ . As before, and following Crooks [59], evolution can be formulated as proceeding in two stages: (i) The external protocol is updated, and then (ii) the particle makes a random transition to a new state. External work is done in the first step, implying

$$W = \sum_{n=0}^{N-1} [V(x_n; \lambda_{n+1}) - V(x_n; \lambda_n)]. \quad (14)$$

Since the free energy is the same in the initial and final ensembles, this external work is equivalent to the “dissipated work”  $W_D$ .

We perform Brownian dynamics simulations using one of the sets of parameters considered by Schmiedl and Seifert [22]:  $\Delta\lambda = \lambda_f - \lambda_i = 5$  ( $\lambda_i = 5$  and  $\lambda_f = 10$ );  $k_B T = \mu = 1$ , with  $k_B T$  the usual thermal energy and  $\mu$  the mobility of the colloid, and total simulation time of  $t_f = 2.69$  units, the time that theoretically yields the highest ratio between the work dissipated by a “naive,” linear trap protocol and the dissipated work corresponding to the optimal protocol [22].

To parametrize our protocols, we consider piecewise linear  $\lambda(t)$ , specified by the values at eight distinct time points; see Eq. (S1) [65]. Starting from an initial guess of a linear protocol, we perform optimization with Adam [64] on batches of  $N = 5000$  MD simulations with learning rate  $\alpha = 0.1$  (this governs the size of step taken on the loss landscape for each iteration), integration time step  $dt = 0.001$ , and an initial equilibration period of 0.1 simulation time units prior to trap motion. We are able to reproduce the exact theoretical optimal curve derived by Schmiedl and Seifert [22] within 100 optimization iterations, taking a few minutes on a GPU; see Fig. 2. Our calculation reproduces the theoretical ratio between the work dissipated by the optimal and linear protocols,  $W_{\text{lin}}/W_{\text{opt}} \approx 1.14$  for  $t_f = 2.69$  [22]. We also run optimization simulations for two additional protocol durations,

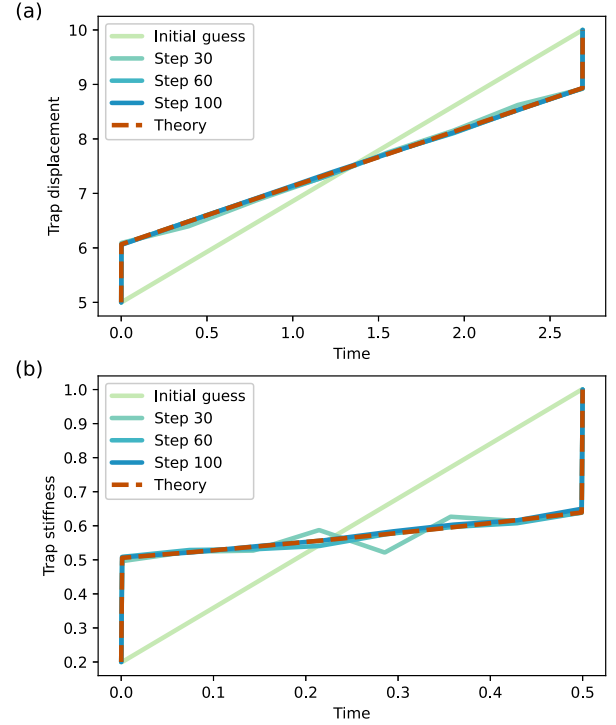


FIG. 2. Automatic differentiation-based optimization rapidly converges to the exact theoretical optimal protocol for the cases of a Brownian particle in (a) a moving harmonic trap and (b) a trap with time-varying stiffness. The discrete jumps at the beginning and the end of the protocol found originally by Schmiedl and Seifert [22] are recaptured by our method.

$t_f = 0.269$  and  $t_f = 26.9$ , to investigate whether our approach is robust even when the signal on which to optimize  $W_{\text{lin}}/W_{\text{opt}}$  is weaker. In both cases, we reproduce the theoretical optimal protocol within 100 optimization steps; see Fig. S7 [65].

Our method also successfully reproduces the exact theoretical optimal protocol for varying the stiffness of a harmonic potential  $V(x, t) = \lambda(t)x^2/2$ . We run for a total simulation time of  $t_f = 0.5$  and use two of the  $\lambda_f/\lambda_i$  ratios considered by Schmiedl and Seifert [22]: one that is expected to maximize the deviation from a linear guess protocol ( $\lambda_i = 0.2$ ,  $\lambda_f = 1.0$ ,  $(\lambda_f/\lambda_i) = 5$ , shown in Fig. 2) and one that optimizes on a weaker signal, expected to produce much smaller jumps in the optimal protocol ( $\lambda_i = 0.2$ ,  $\lambda_f = 0.4$ ,  $(\lambda_f/\lambda_i) = 2$ , shown in Fig. S8 [65]). In the former case, we are able to recover the theoretical optimal protocol within 100 optimization iterations, though we need to increase our batch size to  $N = 10^5$  trajectories as the signal on which to optimize is weaker than in the case of the moving trap [compare Figs. S6(c) and S6(d) [65]]. When the signal on which to optimize is yet weaker still, for  $(\lambda_f/\lambda_i) = 2$ , we find a batch size of  $N = 2 \times 10^5$  and 500 optimization iterations are necessary. Additionally, we need to thermalize the initial system for longer ( $t_{\text{eq}} = 100$ ,

whereas previous examples use  $t_{\text{eq}} = 10$ ), for the same reason a larger batch size is needed: More accurate ensemble averages are required where the signal on which to optimize is weaker.

We note that, while the gradient of the average dissipated work for these optimizations is given by Eq. (12), we find in practice that we are able to achieve convergence with the exact theoretical results for most of the parameter sets considered using only the second term in Eq. (12):  $\langle \nabla_{\theta} \omega \rangle$ . For the case of the weakest optimization signal (Fig. S8 [65]), we use the full gradient expression (12) to improve convergence.

## 2. Driven barrier crossing

We now turn to the more complex situation of driven Brownian motion on a bistable potential, a model used widely in soft matter to represent biomolecular unfolding via AFM or optical tweezers [69,70], as well as optimal protocols for bit erasure [12,71]. Sivak and Crooks [44] consider a Brownian particle driven across a bistable potential energy landscape (see Fig. 3 insets and Supplemental Material [65] for details of the potential) by a harmonic trap with a time-dependent minimum  $\lambda(t)$ . The trap drives barrier crossing from one minimum to the other.

Following Refs. [44,45], we perform molecular dynamics simulations of barrier crossing using parameters approximating DNA hairpin unfolding experiments with optical tweezers; see Supplemental Material [65] for details. For quantitative comparison with previous work, we make the same simplifying assumptions as Sivak and Crooks [44]: (i) The free energies of the initial and final equilibrium states are equal, and (ii) the two landscape wells have equal curvature. Note that our method does not require these assumptions. We proceed by calculating dissipated work (again, equivalent here to the total external work supplied) with Eq. (14). We perform 10 ms simulations and consider two free-energy landscapes, with barrier heights  $2.5k_B T$  and  $10k_B T$ . Since the rate of escape from the potential well varies exponentially with landscape barrier height [73], for a given protocol length, passage over the  $10k_B T$  barrier is much farther from equilibrium than a  $2.5k_B T$  barrier. Further details of how we quantify distance from equilibrium can be found in Supplemental Material, Fig. S9 [65]. A barrier height of  $2.5k_B T$  ( $10k_B T$ ) maps roughly onto the unfolding of a six (20) base pair DNA hairpin [74–76].

Figure 3 presents the results of using automatic differentiation-based optimization to find optimal trap protocols  $\lambda(t)$  for driven barrier crossing; landscape profiles are shown as insets. In the near-equilibrium regime (barrier height  $2.5k_B T$ ), optimizing over a batch of  $N = 2504$  trajectories, our method converges to the near-equilibrium theoretical result of Sivak and Crooks [44] after 1000 optimization iterations, with most of the convergence

achieved after a couple hundred optimization steps (taking a few hours on a TPU). The shape of the optimal protocol—faster trap motion at the beginning and ends of the motion and a slowing down in the central barrier region—reflects the fact that the minimal work is dissipated if the trap slows down in the vicinity of the barrier to “wait” for the system to harness thermal energy kicks to surmount it [44,45].

We compare the limiting probability distributions (across approximately  $10^5$  MD simulations) of dissipated work for the near-equilibrium theoretical protocol, our result, and a naive linear protocol. The AD-optimized and near-equilibrium theoretical distributions agree within error, each having a mean work of  $\langle W_D \rangle = 0.104 \pm 0.001 k_B T$ . Both protocols outperform the naive linear protocol, which gives  $\langle W_D \rangle = 0.136 \pm 0.002 k_B T$ . Errors are standard errors of the mean.

The AD-based optimization allows us to probe far beyond the near-equilibrium regime [Figs. 3(c) and 3(d)]. With a  $10k_B T$  barrier landscape and a protocol length of 10 ms, the automatic differentiation-based optimal protocol outperforms the near-equilibrium theory. Here, our algorithm finds that a nonsymmetric protocol is optimal, whereas linear theory necessarily predicts that it is symmetric [56]. Intuitively, the trap needs to spend more time in the vicinity of the barrier to successfully “drag” the particle along: The bimodal distribution of the near-equilibrium theoretical  $p(W)$  in Fig. 3(d) reveals that not every particle successfully “unfolds” under the near-equilibrium optimal protocol; some are left behind in the “folded” state after the trap completes its motion. These trajectories maximize the external dissipated work: After the trap stops moving, work can no longer accrue according to Eq. (14), even if the particle eventually hops to the unfolded well. In contrast, AD-based optimization finds a protocol that unfolds nearly all molecules in simulation time, leading to a significantly lower average dissipated work of  $\langle W_D \rangle = 3.260 \pm 0.007 k_B T$  compared to the near-equilibrium theoretical mean work of  $\langle W_D \rangle = 5.37 \pm 0.02 k_B T$ . Here, a naive linear protocol ( $\langle W_D \rangle = 5.172 \pm 0.009 k_B T$ ) also outperforms the near-equilibrium theory.

The AD-based  $\lambda(t)$  features discrete jumps at the beginning and end of the protocol that are absent in the linear response optimum [see upper insets in Fig. 3(c)]. Discrete jumps have been observed in multiple other studies of minimum-dissipation protocols [20,67,72] and, indeed, were posited by Schmiedl and Seifert [22] to be a “generic feature of the optimal protocol for arbitrary potentials.” Recent work corroborates the universality of jump features in optimal protocols [77]. Since the near-equilibrium theory assumes protocols to be differentiable, it necessarily misses these features [10,24,56].

We note that, although here we focus on the form of landscape studied in previous literature, our method allows the user to perform a similar barrier-crossing optimization for virtually any free-energy landscape, such as bespoke

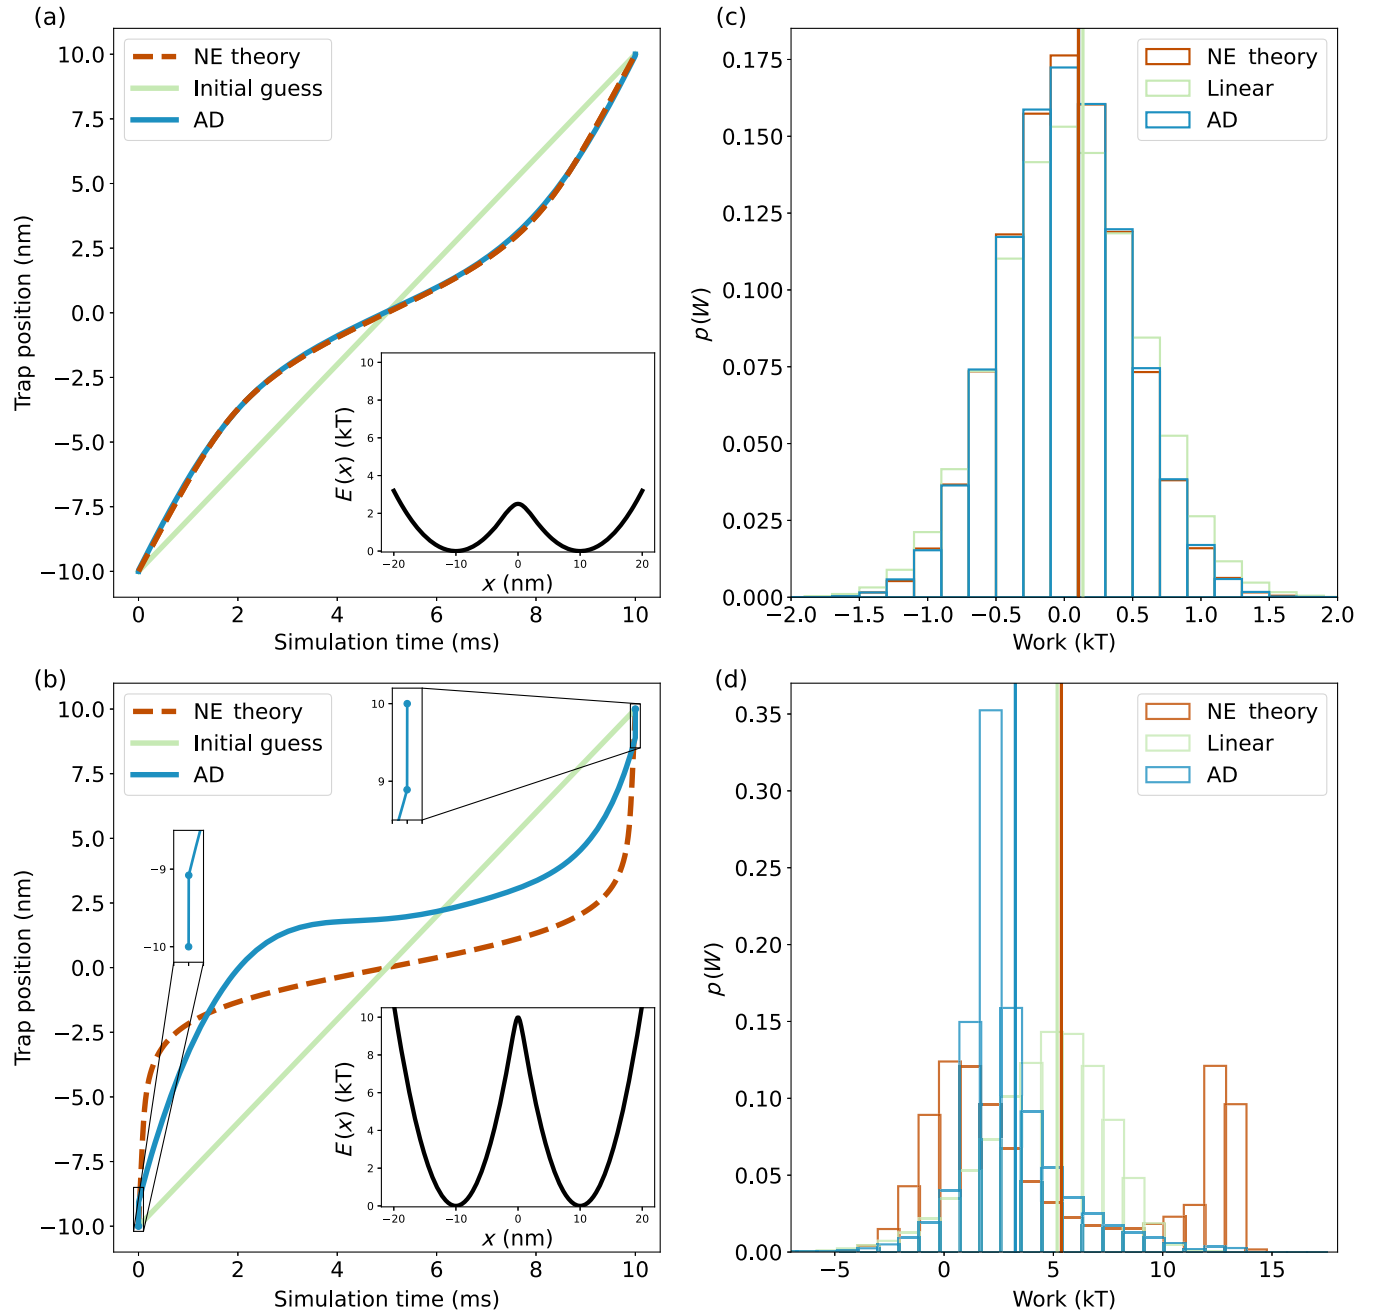


FIG. 3. (a) The automatic differentiation-based optimal protocol converges to the near-equilibrium (NE) theoretical result of Sivak and Crooks [44] in the near-equilibrium regime, corresponding to a  $2.5k_B T$  barrier landscape for a protocol length of 10 ms. Inset: a  $2.5k_B T$  barrier energy landscape. (b) The AD protocol is asymmetric and differs from the NE theoretical result of Sivak and Crooks [44] in the farther-from-equilibrium regime represented by a  $10k_B T$  barrier landscape for a protocol length of 10 ms. Insets near the top in (b), showing enlarged views of the beginning and end of the AD protocol, reveal discrete jumps at the beginning and end of the optimal protocol, in agreement with results for similar systems [20,22,67,72]. Lower right inset: a  $10k_B T$  barrier energy landscape. (c) Probability work distributions associated with the protocols in (a), with the mean dissipated work values under each protocol indicated with vertical lines. AD ( $\langle W_D \rangle = 0.104 \pm 0.001 k_B T$ ) performs as well as the NE theory ( $\langle W_D \rangle = 0.104 \pm 0.001 k_B T$ ) within error, and both outperform a naive linear protocol ( $\langle W_D \rangle = 0.136 \pm 0.002 k_B T$ ). (d) The same as in (c), but for the protocols shown in (b). Here, the AD protocol ( $\langle W_D \rangle = 3.260 \pm 0.007 k_B T$ ) significantly outperforms the NE theory ( $\langle W_D \rangle = 5.37 \pm 0.02 k_B T$ ), which fails to unfold all particles and is, thus, also beaten by a naive linear protocol ( $\langle W_D \rangle = 5.172 \pm 0.009 k_B T$ ).



free-energy landscapes containing nontrivial features—like intermediate states—that map to complex biomolecules.

### 3. *oxDNA* unfolding simulations

To demonstrate the utility of our method, we apply the approach of the previous subsection to the mechanical unfolding of a DNA hairpin. Near-equilibrium optimal protocols were recently applied in this context to reduce the external work required to unfold DNA hairpins [45]. Here, we minimize the external work required to unfold a DNA hairpin in the *far-from-equilibrium* regime by performing coarse-grained simulations with the *oxDNA* model, which has enjoyed widespread use and validation against experiments [78–81]. To do this, we first compute an optimal unfolding protocol using Brownian JAX-MD simulations, and then we apply this protocol in *oxDNA* simulations. Ideally, one would perform automatic differentiation *directly* on the *oxDNA* unfolding simulations, but doing so would necessitate rewriting the *oxDNA* coarse-grained model and simulation software in a differentiable programming language like JAX-MD. This is an avenue we are actively pursuing, but it requires enough additional technical effort to merit a separate, forthcoming publication.

We treat mechanical DNA hairpin unfolding as a driven barrier-crossing problem on a one-dimensional landscape parametrized by the end-to-end extension of the DNA hairpin, as has been done previously [70]. In typical experiments, full underlying free-energy landscapes are not always available for a particular molecule, so, in performing our optimizations, we do not use exact landscape parameters; rather, we estimate barrier height based on unfolding experiments on similar hairpins [82] and take the barrier to be located at an end-to-end extension equal to one-half the contour length of the unfolded ssDNA. We use the bistable potential of Sivak and Crooks [44] as the underlying free-energy landscape, with  $\Delta x \approx 14$  nm and  $\Delta G \approx 6.5k_B T$ . The protocol obtained from minimizing the external work required in JAX-MD Brownian dynamics barrier-crossing simulations (see Fig. S11 in Supplemental Material [65]) is applied to unfold a 20 base pair (bp) hairpin with a 4 bp loop similar to hairpins studied previously [46,82,83] in *oxDNA*. We use an average pulling rate of  $4 \times 10^7$  nm s<sup>-1</sup>. The work distributions resulting from applying our protocol and a naive linear protocol are shown in Fig. 4. Using the optimal protocol derived from automatic differentiation offers a  $13k_B T$  reduction in external energy required to unfold the hairpin.

As the work  $W$  done unfolding a molecule increases, so too does the energy dissipation  $W_{\text{diss}}$  according to

$$W_{\text{diss}} = W - \Delta F, \quad (15)$$

where  $\Delta F$  is the free-energy difference between the folded and unfolded states, which remains constant for different unfolding protocols. The amount of dissipated work

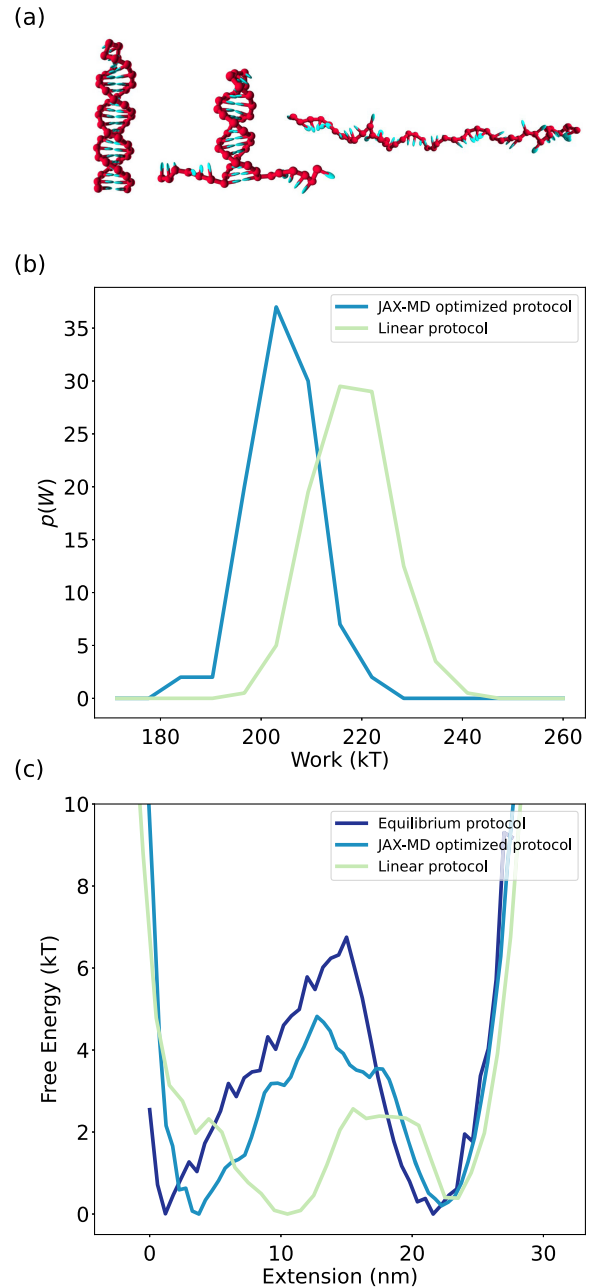


FIG. 4. (a) Snapshots of a 20 bp stem, 4 bp loop DNA hairpin unfolding in *oxDNA* in the far-from-equilibrium regime, simulation length 1  $\mu$ s. (b) Comparison of the work probability distributions for unfolding this DNA hairpin in *oxDNA* with harmonic “traps,” according to both a naive linear protocol (green line,  $\langle W \rangle = 214k_B T$ ) and a protocol derived by automatic-differentiation-based optimization over simplified Brownian dynamics simulations in JAX-MD (blue line,  $\langle W \rangle = 201k_B T$ ); see Supplemental Material [65] for protocol details. (c) Free-energy landscapes for the DNA hairpin reconstructed using nonequilibrium pulling simulations corresponding to both linear (green line) and JAX-MD-optimized (light blue line) protocols. Protocols are compared to a landscape derived using equilibrium umbrella sampling simulations [84] (dark blue line). Nonequilibrium reconstructions are based on Hummer and Szabo’s adaptation of Jarzynski’s equation to biomolecular unfolding [85] and use 100 pulling trajectories each.

governs the accuracy of free-energy estimates derived using nonequilibrium fluctuation theorems like the Jarzynski equality [14,86]. Minimizing the external work done unfolding a molecule should, therefore, improve the accuracy of Jarzynski-based landscape reconstructions. We explore this by reconstructing full free-energy landscapes for our 20 bp hairpin by adapting the Jarzynski equality according to the method proposed by Hummer and Szabo [85], for both linear and optimized unfolding protocols. Figure 4 demonstrates that our work-optimized protocol clearly improves the accuracy of free-energy landscape reconstruction compared to a naive linear protocol. Here, we are comparing both nonequilibrium reconstructions to a free-energy landscape for the hairpin reconstructed using *equilibrium* simulations, which do not suffer from the same convergence errors as Jarzynski-based landscapes [87]. Since barrier crossing is a rare event, advanced sampling techniques are typically required to construct such equilibrium free-energy landscapes, and the dark blue curve in Fig. 4 is calculated using umbrella sampling [84,88]. It has been shown that the proper quantity to minimize when attempting to maximize the accuracy of free-energy reconstructions is not the total work but rather an exponential of the work performed in the reverse protocol [62]. It will be interesting to explore whether JAX-MD optimal protocols using this alternative objective function will be able to improve free-energy landscape reconstructions yet further and is a target of our future work. One can presumably also use the improved free-energy landscape reconstruction to better inform JAX-MD work-optimized protocols in an iterative fashion (for example, the barrier is clearly not located midway between the wells, as assumed when optimizing in JAX-MD). We are actively exploring this possibility.

### III. DISCUSSION

#### A. Form of protocols

In the equilibrium limit, protocols are necessarily symmetric because every microscopic step must be time reversible. Substantial thermodynamic driving forces, however, break time-reversal symmetry. Recent work has explicitly demonstrated a link between the magnitude of thermodynamic driving forces and the degree to which time-reversal symmetry is broken [89]. In light of this, the increasing asymmetry we observe in protocols that exert significant thermodynamic driving forces (i.e., in the farther-from-equilibrium regime) is not unexpected.

For the farthest-from-equilibrium examples treated ( $t = 10$  duration in the Ising example and a barrier height of  $10k_B T$  in the barrier-crossing example), protocols are markedly asymmetric in time. Increasing protocol asymmetry with increasing distance from the equilibrium regime has been observed previously; Blaber, Louwse, and Sivak

find that, in the rapidly driven limit, optimal magnetic field switching protocols for an Ising lattice feature asymmetric jumps to field values beyond the midpoint of their range [77]. Esposito *et al.* find similar departures from symmetry for increasing short (farther-from-equilibrium) protocols for exciting a quantum dot [72]. Recent work on optimal driving protocols for an ATP-synthase-like rotary motor found a similar relationship between protocol symmetry and effective distance from equilibrium [90]. In the case of the Ising spins, one could intuitively imagine that, in order to flip as many spins as possible in the allotted time, the shorter the protocol, the more dramatically the field must change earlier on to incentivize flipping. Any unflipped spins remaining at the end of the procedure incur a large penalty; see Sec. III C.

Recent work by Zhong and DeWeese [91] replicates our findings of symmetry breaking in driven barrier crossing, though in their case protocols become symmetric again in the limit of very fast switching, because they consider a symmetric energy landscape. In the general barrier-crossing case, the time at the position of the barrier is maximized in the fast-switching limit, leading to asymmetric protocols for barriers not equidistant between the wells [77,91]. This is to give particles “as much time as possible” in the vicinity of the barrier in order to maximize their chances of crossing, since any particles that fail to transition to the second well in the protocol duration incur a large penalty in the loss function; again, see Sec. III C.

#### B. Choice of basis functions and approach

Our approach across the case studies detailed above is to begin with (i) a simple choice of basis functions and (ii) reverse-mode automatic differentiation in the first instance. Only when results fail to converge or gradients are poorly behaved do we refine the approach. For example, we cannot derive a convergent optimal protocol for the Ising model using basis functions that linearly interpolated between discrete points and reverse-mode AD. We begin to employ forward-mode AD for the longest simulations to mitigate memory issues. Next, we attempt basis functions that capture the deviation from initial guess functional forms for temperature and magnetic field using Chebyshev polynomials. Our gradients remain poorly behaved due to the fact that Monte Carlo simulations are inherently nondifferentiable. We then attempt to use so-called “straight-through” gradient estimation, where gradients of undifferentiable steps are set to 1 [61], but still do not achieve convergence. Finally, we employ the REINFORCE calculation of Eq. (12) to compute gradients [61], and this leads to the successful optimizations detailed in the text. In general, we expect simulations containing undifferentiable “coin flips” to necessitate the use of Eq. (12). Longer, memory-intensive simulations will likely require forward-mode AD or checkpointing [92]. The optimal choice of basis functions is likely context dependent, but beginning

with simple ones and gradually adding complexity is a sensible approach.

### C. Ensuring successful transitions

Nonequilibrium protocol optimization problems seek protocols that force *successful* transitions from initial state  $A$  to final state  $B$ . For the Ising magnetization reversal, the loss function [Eq. (1)] contains a term proportional to the Hamiltonian of the final state (which is larger for larger numbers of “unflipped” spins):  $\beta_N H(x_N, \lambda_N)$ . For our purposes, this term is large enough to ensure that none of the AD-optimal protocols feature incomplete transitions by penalizing protocols that end without flipping *all* spins.

In the barrier-crossing examples, the form of the loss [Eq. (14)] likewise ensures that protocols that fail to drag the Brownian particle to the second minimum on the energy landscape are *always* less favorable than protocols that succeed in doing so, since “failed” transitions incur a term like  $V = \frac{1}{2}(x_{\text{final}} - \lambda_f)^2$  in the loss, incentivizing protocols that minimize this distance.

For the general use case, successful transitions can be ensured by incorporating a penalty into the loss function that renders protocols that do *not* accomplish the transition to *always* be less favorable than those that do. For example, including a term like

$$\text{Loss} = \eta[A(x_{\text{final}}) - B(x_{\text{final}})] + \dots, \quad (16)$$

where  $\eta$  is a weight that can be increased until no failed transitions are observed,  $A$  is some function capturing how “ $A$ -like” the final state  $x_{\text{final}}$  is, and  $B$  is a function capturing how “ $B$ -like” the final state  $x_{\text{final}}$  is. In the Ising example,  $A$  and  $B$  could simply be the number of unflipped and flipped spins, respectively, and in the barrier-crossing examples,  $A$  and  $B$  could be particle distance from the second well on the energy landscape. Again, while we do not find this to be necessary in practice for the examples we consider, such a constrained optimization should not present significant additional computational difficulty.

## IV. CONCLUSION

We have demonstrated the viability of AD to identify optimal nonequilibrium protocols for both Monte Carlo and molecular dynamics simulations. The method performed as well as existing near-equilibrium theoretical results in the near-equilibrium regime for both magnetization reversal on a 2D Ising lattice and driven barrier crossing. Critically, the AD algorithm easily extends to far-from-equilibrium conditions, where it significantly outperforms existing near-equilibrium theoretical protocols. We have evidenced the applicability of our method to real-world biomolecular unfolding experiments by showing AD-optimized protocols can reduce the energy required to unfold a DNA hairpin in a realistic coarse-grained DNA simulation and,

thereby, improve the accuracy of nonequilibrium free-energy landscape reconstructions.

Our work considers fixed simulation times with one- and two-dimensional protocols, though the framework is much more general than this, and essentially arbitrary constraints and protocol parameters are possible, e.g., multidimensional external protocols and arbitrary loss functions. For example, one could attempt to optimize the speed or efficiency of nanoengines; minimize the time taken to unfold a molecule; or maximize the accuracy of Jarzynski-based free-energy landscape reconstructions for a given amount of experimental or computational time. Furthermore, in the JAX-MD code suite, non-Brownian dynamics can easily be simulated, opening up the possibilities of optimizing protocols for systems like the recently proposed active-matter-based thermodynamic engine [15]. All that our strategy requires are known dynamical equations that govern simulations, so there is no necessary assumption of Markovian evolution.

Our approach is complementary to recent work of Whitelam [27], which presents a method for computing optimal external parameter protocols parametrized by deep neural networks whose parameters are iteratively adjusted and refined either via adaptive Monte Carlo (aMC) [93] or genetic algorithms (GA) [94]. Whitelam benchmarks his approach against our results, detailed in Ref. [95], and finds it a viable alternative to the gradient-based strategies we present here. It has particular promise for applications in which gradient information is unavailable or unreliable: experimental settings and cases where gradients are poorly behaved [96]. Where gradients *are* accessible, however, AD algorithms converge significantly more quickly than aMC and GA approaches, particularly in complex settings with tens of millions of parameters to train [27,93], and also have the advantage of providing information about the loss landscape. Further, aMC and GA require many ( $10^2$ – $10^5$ ) iterations for convergence, which may be unfeasible in some experimental contexts. In particular, in cases of inverse system design (for example, optimizing the sequence of a DNA duplex to tune unbinding kinetics), regenerating a new system thousands of times in response to aMC or GA perturbations may be costly or impractically time consuming. In such cases, an AD approach would be preferable, and we have demonstrated here that protocols derived from simulations can nonetheless be useful in experimentally realistic settings. Whitelam’s use of neural networks to parametrize external protocols could be easily combined with the AD approach we describe here. He also demonstrates the use of feedback to achieve negative entropy production and dissipated work, effectively employing Maxwell’s demon to convert information into other forms of energy. Incorporating feedback information into the pipelines we have described above is an exciting avenue for future research and one we are pursuing.

We are actively working to extend our framework to truly complex examples. Forthcoming publications treat AD within the full oxDNA model (which requires transposing the model into the differentiable JAX-MD framework) and for active matter systems containing tens of thousands of particles. The main challenge of AD-based protocol optimization in the context of increasingly complex systems is the high computational cost of back-propagating gradients through entire simulations. We have already demonstrated a possible route to addressing this in forward-mode automatic differentiation. In addition to using forward-mode automatic differentiation, one can use gradient checkpointing with reverse-mode autodiff, which mitigates the cost of storing unrolled long, complex simulations in memory by saving and rematerializing snapshots of the gradient [92]. Other authors have had success unrolling only the last  $N$  steps of simulation trajectories for automatic differentiation [28]. There are also alternative gradient estimation methods which still employ AD but do not require full simulation trajectories to be unrolled: implicit differentiation [97] and differentiable trajectory reweighting [98]. The application of AD in the present context of nonequilibrium thermodynamics is in its infancy, and the field will continue to benefit from ongoing improvements in gradient estimation and AD performance from other domains as well as continual hardware improvements. We have, nevertheless, already demonstrated that protocols derived in simplified 1D JAX-MD simulations can be successful when applied to more complex contexts, namely, oxDNA simulations, suggesting they can already be profitably applied to experiments. The longest optimizations in the current work (barrier crossing for  $t = 0.01$  s) took 5 GPU hours for 1000 iterations, and the longest Ising lattice optimizations for  $N = 1024$  spins took less than an hour on an A100 GPU. These are well below typical simulation times for molecular dynamics studies, suggesting a broad scope already exists for treating more complex problems before computational time concerns rival those of some MD simulations.

AD provides a valuable complement to existing near-equilibrium approaches to find optimal protocols, as it makes more complex systems and the far-from-equilibrium regime accessible. We are eager to see its manifold applications unfold in nonequilibrium protocol optimization and beyond.

The code used to perform this research is freely available [99].

### ACKNOWLEDGMENTS

We thank Carl Goodrich, Ella King, and Daniel Fisher for important conversations and Grant Rotskoff for generously sharing data with us. This research was supported by the Office of Naval Research through ONR N00014-17-

1-3029 and the Simons Foundation. M. C. E. also thanks Schmidt Futures in partnership with The Rhodes Trust for funding this work.

- 
- [1] Y. Rondelez, G. Tresset, T. Nakashima, Y. Kato-Yamada, H. Fujita, S. Takeuchi, and H. Noji, *Highly coupled ATP synthesis by F1-ATPase single molecules*, *Nature (London)* **433**, 773 (2005).
  - [2] L. M. Alexander, D. H. Goldman, L. M. Wee, and C. Bustamante, *Non-equilibrium dynamics of a nascent polypeptide during translation suppress its misfolding*, *Nat. Commun.* **10**, 2709 (2019).
  - [3] F. Wruck, A. Katranidis, K. H. Nierhaus, G. Büldt, and M. Hegner, *Translation and folding of single proteins in real time*, *Proc. Natl. Acad. Sci. U.S.A.* **114**, E4399 (2017).
  - [4] P. W. K. Rothmund, *Folding DNA to create nanoscale shapes and patterns*, *Nature (London)* **440**, 297 (2006).
  - [5] N. C. Seeman and H. F. Sleiman, *DNA nanotechnology*, *Nat. Rev. Mater.* **3**, 17068 (2018).
  - [6] P.-S. Huang, S. E. Boyken, and D. Baker, *The coming of age of de novo protein design*, *Nature (London)* **537**, 320 (2016).
  - [7] I. V. Korendovych and W. F. DeGrado, *De novo protein design, a retrospective*, *Q. Rev. Biophys.* **53**, e3 (2020).
  - [8] U. Seifert, *Stochastic thermodynamics, fluctuation theorems and molecular machines*, *Rep. Prog. Phys.* **75**, 126001 (2012).
  - [9] R. E. Spinney and I. J. Ford, *Fluctuation relations: A pedagogical overview*, [arXiv:1201.6381](https://arxiv.org/abs/1201.6381).
  - [10] P. R. Zulkowski and M. R. DeWeese, *Optimal control of overdamped systems*, *Phys. Rev. E* **92**, 032117 (2015).
  - [11] K. Proesmans, J. Ehrich, and J. Bechhoefer, *Finite-time Landauer principle*, *Phys. Rev. Lett.* **125**, 100602 (2020).
  - [12] K. Proesmans, J. Ehrich, and J. Bechhoefer, *Optimal finite-time bit erasure under full control*, *Phys. Rev. E* **102**, 032105 (2020).
  - [13] C. Dellago and G. Hummer, *Computing equilibrium free energies using non-equilibrium molecular dynamics*, *Entropy* **16**, 41 (2013).
  - [14] N. Yunger Halpern and C. Jarzynski, *Number of trials required to estimate a free-energy difference, using fluctuation relations*, *Phys. Rev. E* **93**, 052144 (2016).
  - [15] P. Pietzonka, É. Fodor, C. Lohrmann, M. E. Cates, and U. Seifert, *Autonomous engines driven by active matter: Energetics and design principles*, *Phys. Rev. X* **9**, 041032 (2019).
  - [16] T. K. Saha, J. N. E. Lucero, J. Ehrich, D. A. Sivak, and J. Bechhoefer, *Maximizing power and velocity of an information engine*, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2023356118 (2021).
  - [17] E. Lathouwers, J. N. E. Lucero, and D. A. Sivak, *Non-equilibrium energy transduction in stochastic strongly coupled rotary motors*, *J. Phys. Chem. Lett.* **11**, 5273 (2020).
  - [18] G. M. Rotskoff and G. E. Crooks, *Optimal control in nonequilibrium systems: Dynamic Riemannian geometry of the Ising model*, *Phys. Rev. E* **92**, 060102(R) (2015).

- [19] G. M. Rotskoff, G. E. Crooks, and E. Vanden-Eijnden, *Geometric approach to optimal nonequilibrium control: Minimizing dissipation in nanomagnetic spin systems*, *Phys. Rev. E* **95**, 012148 (2017).
- [20] A. Gomez-Marín, T. Schmiedl, and U. Seifert, *Optimal protocols for minimal work processes in underdamped stochastic thermodynamics*, *J. Chem. Phys.* **129**, 024114 (2008).
- [21] E. Aurell, C. Mejía-Monasterio, and P. Muratore-Ginanneschi, *Optimal protocols and optimal transport in stochastic thermodynamics*, *Phys. Rev. Lett.* **106**, 250601 (2011).
- [22] T. Schmiedl and U. Seifert, *Optimal finite-time processes in stochastic thermodynamics*, *Phys. Rev. Lett.* **98**, 108301 (2007).
- [23] A. P. Solon and J. M. Horowitz, *Phase transition in protocols minimizing work fluctuations*, *Phys. Rev. Lett.* **120**, 180605 (2018).
- [24] P. R. Zulkowski, D. A. Sivak, G. E. Crooks, and M. R. DeWeese, *Geometry of thermodynamic control*, *Phys. Rev. E* **86**, 041148 (2012).
- [25] M. V. S. Bonança and S. Deffner, *Optimal driving of isothermal processes close to equilibrium*, *J. Chem. Phys.* **140**, 244119 (2014).
- [26] S. J. Large and D. A. Sivak, *Optimal discrete control: Minimizing dissipation in discretely driven nonequilibrium systems*, *J. Stat. Mech.* (2019) 083212.
- [27] S. Whitelam, *Demon in the machine: Learning to extract work and absorb entropy from fluctuating nanosystems*, *Phys. Rev. X* **13**, 021005 (2023).
- [28] C. P. Goodrich, E. M. King, S. S. Schoenholz, E. D. Cubuk, and M. P. Brenner, *Designing self-assembling kinetics with differentiable statistical physics models*, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2024083118 (2021).
- [29] X. Tang, B. Rupp, Y. Yang, T. D. Edwards, M. A. Grover, and M. A. Bevan, *Optimal feedback controlled assembly of perfect crystals*, *ACS Nano* **10**, 6791 (2016).
- [30] X. Tang, J. Zhang, M. A. Bevan, and M. A. Grover, *A comparison of open-loop and closed-loop strategies in colloidal self-assembly*, *J. Process Control* **60**, 141 (2017).
- [31] M. A. Grover, D. J. Griffin, X. Tang, Y. Kim, and R. W. Rousseau, *Optimal feedback control of batch self-assembly processes using dynamic programming*, *Journal of Process Control* **88**, 32 (2020).
- [32] A. Trubiano and M. F. Hagan, *Optimization of non-equilibrium self-assembly protocols using Markov state models*, *J. Chem. Phys.* **157**, 244901 (2022).
- [33] T. R. Gingrich, G. M. Rotskoff, G. E. Crooks, and P. L. Geissler, *Near-optimal protocols in complex nonequilibrium transformations*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 10263 (2016).
- [34] S. Whitelam and I. Tamblyn, *Learning to grow: Control of material self-assembly using evolutionary reinforcement learning*, *Phys. Rev. E* **101**, 052604 (2020).
- [35] A. E. Bryson and W. F. Denham, *A steepest-ascent method for solving optimum programming problems*, *J. Appl. Mech.* **29**, 247 (1962).
- [36] A. Bryson and Y.-C. Ho, *Applied Optimal Control: Optimization, Estimation, and Control*, 1st ed. (Routledge, Oxfordshire, 1975).
- [37] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning representations by back-propagating errors*, *Nature (London)* **323**, 533 (1986).
- [38] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, *Automatic differentiation in machine learning: A survey*, *J. Mach. Learn. Res.* **18**, 1 (2018).
- [39] R. A. Vargas-Hernández, R. T. Q. Chen, K. A. Jung, and P. Brumer, *Fully differentiable optimization protocols for non-equilibrium steady states*, *New J. Phys.* **23**, 123006 (2021).
- [40] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16)*, 2016, Savannah, GA, USA (2016), p. 265, <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>
- [41] A. Paszke et al., *PyTorch: An imperative style, high-performance deep learning library*, in *Advances in Neural Information Processing Systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., New York, 2019), <https://www.proceedings.com/content/053/053719webtoc.pdf>.
- [42] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne, *JAX: Composable transformations of Python + NumPy programs*, GitHub (2018), <http://github.com/google/jax>.
- [43] C. C. Margossian, *A review of automatic differentiation and its efficient implementation*, *WIREs Data Mining Knowledge Discovery* **9**, e1305 (2019).
- [44] D. A. Sivak and G. E. Crooks, *Thermodynamic geometry of minimum-dissipation driven barrier crossing*, *Phys. Rev. E* **94**, 052106 (2016).
- [45] S. Tafuya, S. J. Large, S. Liu, C. Bustamante, and D. A. Sivak, *Using a system's equilibrium behavior to reduce its energy dissipation in nonequilibrium processes*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 5920 (2019).
- [46] A. N. Gupta, A. Vincent, K. Neupane, H. Yu, F. Wang, and M. T. Woodside, *Experimental validation of free-energy-landscape reconstruction from non-equilibrium single-molecule force spectroscopy measurements*, *Nat. Phys.* **7**, 631 (2011).
- [47] M. T. Woodside and S. M. Block, *Reconstructing folding energy landscapes by single-molecule force spectroscopy*, *Annu. Rev. Biophys.* **43**, 19 (2014).
- [48] P. R. Heenan, H. Yu, M. G. W. Siewny, and T. T. Perkins, *Improved free-energy landscape reconstruction of bacteriorhodopsin highlights local variations in unfolding energy*, *J. Chem. Phys.* **148**, 123313 (2018).
- [49] V. Walhorn, A.-K. Möller, C. Bartz, T. Dierks, and D. Anselmetti, *Exploring the sulfatase I catch bond free energy landscape using Jarzynski's equality*, *Sci. Rep.* **8**, 16849 (2018).
- [50] S. Park, F. Khalili-Araghi, E. Tajkhorshid, and K. Schulten, *Free energy calculation from steered molecular dynamics simulations using Jarzynski's equality*, *J. Chem. Phys.* **119**, 3559 (2003).
- [51] Y. Zhuang, H. R. Bureau, S. Quirk, and R. Hernandez, *Adaptive steered molecular dynamics of biomolecules*, *Mol. Simul.* **47**, 408 (2021).

- [52] J. Gore, F. Ritort, and C. Bustamante, *Bias and error in estimates of equilibrium free-energy differences from non-equilibrium measurements*, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 12564 (2003).
- [53] O. Perišić and H. Lu, *On the improvement of free-energy calculation from steered molecular dynamics simulations using adaptive stochastic perturbation protocols*, *PLoS One* **9**, e101810 (2014).
- [54] B. Lambson, D. Carlton, and J. Bokor, *Exploring the thermodynamic limits of computation in integrated systems: Magnetic memory, nanomagnetic logic, and the Landauer limit*, *Phys. Rev. Lett.* **107**, 010604 (2011).
- [55] J. Hong, B. Lambson, S. Dhuey, and J. Bokor, *Experimental test of Landauer's principle in single-bit operations on nanomagnetic memory bits*, *Sci. Adv.* **2**, e1501492 (2016).
- [56] D. A. Sivak and G. E. Crooks, *Thermodynamic metrics and optimal paths*, *Phys. Rev. Lett.* **108**, 190602 (2012).
- [57] C. Jarzynski, *Nonequilibrium work relations: Foundations and applications*, *Eur. Phys. J. B* **64**, 331 (2008).
- [58] G. E. Crooks, *Excursions in statistical dynamics*, Ph.D. thesis, University of California at Berkeley, 1999.
- [59] G. E. Crooks, *Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems*, *J. Stat. Phys.* **90**, 1481 (1998).
- [60] R. J. Glauber, *Time-dependent statistics of the Ising model*, *J. Math. Phys. (N.Y.)* **4**, 294 (1963).
- [61] Y. Bengio, N. Léonard, and A. Courville, *Estimating or propagating gradients through stochastic neurons for conditional computation*, arXiv:1308.3432.
- [62] P. Geiger and C. Dellago, *Optimum protocol for fast-switching free-energy calculations*, *Phys. Rev. E* **81**, 021127 (2010).
- [63] R. J. Williams, *Simple statistical gradient-following algorithms for connectionist reinforcement learning*, *Mach. Learn.* **8**, 229 (1992).
- [64] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, arXiv:1412.6980.
- [65] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.13.041032> for additional figures and simulation details.
- [66] M. de Koning, *Optimizing the driving function for non-equilibrium free-energy calculations in the linear regime: A variational approach*, *J. Chem. Phys.* **122**, 104106 (2005).
- [67] H. Then and A. Engel, *Computing the optimal protocol for finite-time processes in stochastic thermodynamics*, *Phys. Rev. E* **77**, 041105 (2008).
- [68] S. S. Schoenholz and E. D. Cubuk, *JAX M.D. A framework for differentiable physics*, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., New York, 2020), Vol. 33, p. 11428, <https://www.proceedings.com/content/059/059066webtoc.pdf>.
- [69] O. K. Dudko, G. Hummer, and A. Szabo, *Intrinsic rates and activation free energies from single-molecule pulling experiments*, *Phys. Rev. Lett.* **96**, 108101 (2006).
- [70] M. T. Woodside and S. M. Block, *Reconstructing folding energy landscapes by single-molecule force spectroscopy*, *Annu. Rev. Biophys.* **43**, 19 (2014).
- [71] E. Aurell, K. Gawędzki, C. Mejía-Monasterio, R. Mohayaei, and P. Muratore-Ginanneschi, *Refined second law of thermodynamics for fast random processes*, *J. Stat. Phys.* **147**, 487 (2012).
- [72] M. Esposito, R. Kawai, K. Lindenberg, and C. Van den Broeck, *Finite-time thermodynamics for a single-level quantum dot*, *Europhys. Lett.* **89**, 20003 (2010).
- [73] H. Kramers, *Brownian motion in a field of force and the diffusion model of chemical reactions*, *Physica (Utrecht)* **7**, 284 (1940).
- [74] E. Pfitzner, C. Wachauf, F. Kilchherr, B. Pelz, W. M. Shih, M. Rief, and H. Dietz, *Rigid DNA beams for high-resolution single-molecule mechanics*, *Angew. Chem., Int. Ed. Engl.* **52**, 7766 (2013).
- [75] M. Engel, *DNA Systems under Internal and External Forcing: An Exploration Using Coarse-Grained Modelling*, Springer Theses (Springer International, New York, 2019).
- [76] M. T. Woodside, P. C. Anthony, W. M. Behnke-Parks, K. Larizadeh, D. Herschlag, and S. M. Block, *Direct measurement of the full, sequence-dependent folding landscape of a nucleic acid*, *Science* **314**, 1001 (2006).
- [77] S. Blaber, M. D. Louwerse, and D. A. Sivak, *Steps minimize dissipation in rapidly driven stochastic systems*, *Phys. Rev. E* **104**, L022101 (2021).
- [78] T. E. Ouldridge, A. A. Louis, and J. P. K. Doye, *Structural, mechanical, and thermodynamic properties of a coarse-grained DNA model*, *J. Chem. Phys.* **134**, 085101 (2011).
- [79] P. Šulc, F. Romano, T. E. Ouldridge, L. Rovigatti, J. P. K. Doye, and A. A. Louis, *Sequence-dependent thermodynamics of a coarse-grained DNA model*, *J. Chem. Phys.* **137**, 135101 (2012).
- [80] B. E. K. Snodin, F. Randisi, M. Mosayebi, P. Šulc, J. S. Schreck, F. Romano, T. E. Ouldridge, R. Tsukanov, E. Nir, A. A. Louis, and J. P. K. Doye, *Introducing improved structural properties and salt dependence into a coarse-grained model of DNA*, *J. Chem. Phys.* **142**, 234901 (2015).
- [81] A. Sengar, T. E. Ouldridge, O. Henrich, L. Rovigatti, and P. Šulc, *A primer on the oxDNA Model of DNA: When to use it, how to simulate it and how to interpret the results*, *Front. Mol. Biosci.* **8**, 693710 (2021).
- [82] M. T. Woodside, W. M. Behnke-Parks, K. Larizadeh, K. Travers, D. Herschlag, and S. M. Block, *Nanomechanical measurements of the sequence-dependent folding landscapes of single nucleic acid hairpins*, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 6190 (2006).
- [83] M. C. Engel, D. B. Ritchie, D. A. N. Foster, K. S. D. Beach, and M. T. Woodside, *Reconstructing folding energy landscape profiles from nonequilibrium pulling curves with an inverse Weierstrass integral transform*, *Phys. Rev. Lett.* **113**, 238104 (2014).
- [84] M. Engel, *DNA Systems under Internal and External Forcing: An Exploration Using Coarse-Grained Modelling*, Springer Theses (Springer, New York, 2019).
- [85] G. Hummer and A. Szabo, *Free energy reconstruction from nonequilibrium single-molecule pulling experiments*, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 3658 (2001).
- [86] C. Jarzynski, *Rare events and the convergence of exponentially averaged work values*, *Phys. Rev. E* **73**, 046105 (2006).
- [87] M. T. Woodside and S. M. Block, *Reconstructing folding energy landscapes by single-molecule force spectroscopy*, *Annu. Rev. Biophys.* **43**, 19 (2014).

- [88] G. Torrie and J. Valleau, *Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling*, *J. Comput. Phys.* **23**, 187 (1977).
- [89] N. Ohga, S. Ito, and A. Kolchinsky, *Thermodynamic bound on the asymmetry of cross-correlations*, *Phys. Rev. Lett.* **131**, 077101 (2023).
- [90] D. Gupta, S. J. Large, S. Toyabe, and D. A. Sivak, *Optimal Control of the F<sub>1</sub>-ATPase molecular motor*, *J. Phys. Chem. Lett.* **13**, 11844 (2022).
- [91] A. Zhong and M. R. DeWeese, *Limited-control optimal protocols arbitrarily far from equilibrium*, *Phys. Rev. E* **106**, 044135 (2022).
- [92] A. Griewank and A. Walther, *Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation*, *ACM Trans. Math. Softw.* **26**, 19 (2000).
- [93] S. Whitelam, V. Selin, I. Benlolo, C. Casert, and I. Tamblyn, *Training neural networks using Metropolis Monte Carlo and an adaptive variant*, *Mach. Learn.: Sci. Technol.* **3**, 045026 (2022).
- [94] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, *Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning*, arXiv:1712.06567.
- [95] M. C. Engel, J. A. Smith, and M. P. Brenner, *Optimal control of nonequilibrium systems through automatic differentiation*, arXiv:2201.00098v1.
- [96] L. Metz, C. D. Freeman, S. S. Schoenholz, and T. Kachman, *Gradients are not all you need*, arXiv:2111.05803.
- [97] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, *Efficient and modular implicit differentiation*, arXiv:2105.15183.
- [98] S. Thaler and J. Zavadlav, *Learning neural network potentials from experimental data via differentiable trajectory reweighting*, *Nat. Commun.* **12**, 6884 (2021).
- [99] [https://github.com/mc2engel/noneq\\_opt](https://github.com/mc2engel/noneq_opt).