

## Optimal computation from fluctuation responses

Jinghao Lyu,<sup>\*</sup> Kyle J. Ray<sup>Ⓜ,†</sup> and James P. Crutchfield<sup>Ⓜ,‡</sup>Complexity Sciences Center and Physics and Astronomy Department, *University of California at Davis*,  
One Shields Avenue, Davis, California 95616, USA

(Received 13 October 2025; accepted 10 May 2026; published 8 June 2026)

The energy cost of computation has emerged as a central challenge at the intersection of physics and computer science. Recent advances in statistical physics—particularly in stochastic thermodynamics—enable precise characterizations of work, heat, and entropy production in information-processing systems driven far from equilibrium by time-dependent control protocols. A key open question is then how to design protocols that minimize thermodynamic cost while ensuring correct outcomes. To this end, we develop a unified framework to identify optimal protocols using fluctuation-response relations (FRRs) and machine learning. Unlike previous approaches that optimize either distributions or protocols separately, our method unifies both using FRR-derived gradients. Moreover, our method is based primarily on iteratively learning from sampled noisy trajectories, which is generally much easier than solving for the optimal protocol directly from a set of governing equations. We apply the framework to canonical examples—bit erasure in a double-well potential and translating harmonic traps—demonstrating how to construct loss functions that trade-off energy cost against task error. The framework extends trivially to underdamped systems, and we show this by optimizing a bit flip in an underdamped system. In all computations we test, the framework achieves the theoretically optimal protocol or achieves work costs comparable to relevant finite-time bounds. In short, the results provide principled strategies for designing thermodynamically efficient protocols in physical information-processing systems. Applications range from quantum gates robust under noise to energy-efficient control of chemical and synthetic biological networks.

DOI: [10.1103/8xjg-rbdl](https://doi.org/10.1103/8xjg-rbdl)

## I. INTRODUCTION

Nowadays, computation—ranging from operating everyday digital devices to training large-scale machine learning models—consumes a significant amount of energy. This raises questions we long to answer: How can computation be performed efficiently, and what is the minimal cost of implementing computation? The most well-known result in this setting is Landauer’s bound that the minimum energetic cost to erase one bit of information at temperature  $T$  is  $k_B T \log_2$ , where  $k_B$  is Boltzmann’s constant [1]. Though Landauer’s original derivation has limitations, it was groundbreaking in establishing a deep connection between computation and thermodynamics.

Since Landauer’s time, we now appreciate that rigorously investigating the energetic costs of computation requires nonequilibrium physics—as all computation takes place in physical systems operating far from equilibrium. A key development was the stochastic description of thermodynamic

resources, which led to *stochastic thermodynamics* [2] and established an arena for exploring the thermodynamics of computation.

Consider a stochastic system  $\mathcal{S}$  with a time-dependent energy landscape control parameter  $a_t$ . Suppose the initial distribution over  $\mathcal{S}$  is  $p_0$ . To implement a computation, we aim to transform the distribution from  $p_0$  to a goal distribution  $p_{\text{target}}$  over a time duration  $\tau$ . We further impose a cyclic boundary condition on the control protocol— $a_0 = a_\tau$ —to enable information storage and to support follow-on computations. The question then becomes: How to design the time-dependent protocol  $a_t$  to transform the distribution from  $p_0 \rightarrow p_{\text{target}}$  at minimal energetic cost?

Two distinct energetic cost optimization approaches have been studied previously [3–22]. The first is distribution oriented: transform an initial distribution into a target distribution, without regard to the final energy landscape. In this, we have the correct distribution but the potential is not correct. As a result, the information after the processing is unstable and cannot be used in the future. The second approach is protocol oriented: drive the system from one protocol configuration to another, without considering the resulting final distribution. However, this is unsuited to computational tasks. Imposing constraints on both final distribution and final protocol values can overconstrain the system, leading to no feasible solution at all.

Unlike previous approaches that only optimize either distributions or protocols, we introduce a framework integrating both optimizations. Following Ref. [23], we use a

<sup>\*</sup>Contact author: [jolyu@ucdavis.edu](mailto:jolyu@ucdavis.edu)<sup>†</sup>Contact author: [kjray@ucdavis.edu](mailto:kjray@ucdavis.edu)<sup>‡</sup>Contact author: [chaos@ucdavis.edu](mailto:chaos@ucdavis.edu)

coarse-grained distribution to store information. In a computational task, we do not ask the final distribution in  $\mathcal{S}$  to be exactly  $p_{\text{target}}$ . Instead, we only need the coarse-grained distribution at final time  $p_{\tau_{\text{cg}}}$  to match the coarse-grained target distribution  $p_{\text{target}_{\text{cg}}}$ . Therefore, we can ensure both a correct coarse-grained distribution and a correct potential landscape. We show how to design loss functions for different computation tasks and use fluctuation-response relations (FRRs) to compute the loss function gradients. We demonstrate that these loss functions successfully drive the distribution toward the target coarse-grained distribution, all the while achieving work costs on the order of Landauer's lower bound.

Our framework has several advantages: (1) Based on universal fluctuation-response relations, it can be used to optimize a large family of stochastic systems including overdamped and underdamped Langevin dynamics, Markovian jump processes, and the like. (2) It allows for highly flexible weighting of various target metrics, like logical error and energetic cost, in the loss function. (3) Because the learning algorithm works by gradient descent and observing samples of simulated trajectories, rather than solving a system of equations, it lends itself to computer implementation that takes advantage of modern machine learning tools that can scale to higher dimensional problems.

The development is organized as follows: Section II reviews FRR basics and turns to discuss loss-function design in Sec. III. Then, we use FRR to optimize loss functions and explore numerical examples. Throughout,  $(\cdot)_t$  denotes the value of a quantity  $(\cdot)$  at time  $t$ .

## II. FLUCTUATION-RESPONSE RELATION

We first review the FRR in the context of overdamped dynamics [24–26]. Consider one-dimensional overdamped Langevin dynamics with an external control parameter  $a$ :

$$dx = \mu\Phi(x, a)dt + \sigma dw_t, \quad (1)$$

where  $\mu$  is the particle mobility,  $dw_t$  is an infinitesimal Wiener process,  $\sigma = \sqrt{2\mu k_B T}$  is the standard deviation of the thermal noise or  $D = \frac{1}{2}\sigma^2$  is the diffusion constant, and  $\Phi(x, a) = -\partial_x U(x, a)$ , with  $U(x, a)$  being the potential of the system as a function of the position  $x$  and the control parameter  $a$ .

An observable (called a cumulant current)  $J$  along a trajectory  $\omega = \{x_t\}_{t=0}^\tau$  is defined as

$$J(\omega, a) = \int_\omega f(x, a) \cdot dx + g(x, a)dt, \quad (2)$$

where  $\cdot$  is the Itô product. For example, the energetic cost (work) of one trajectory  $\omega$  can be written as  $W(\omega, a) = \int_\omega \partial_t U(x, a)dt$ . The expectation value of  $J(\omega, a)$  over an ensemble of trajectories is

$$\mathbb{E}[J(\omega, a)] = \sum_\omega \text{Pr}(\omega, a)J(\omega, a), \quad (3)$$

where  $\text{Pr}(\omega, a)$  is the probability of observing the trajectory  $\omega$  under protocol  $a$ .

The purpose of an FRR is to monitor how the expectation of a current  $J$  changes due to variation in the control

parameters:  $\frac{\partial}{\partial a}\mathbb{E}[J(\omega, a)]$ . From the chain rule, we have

$$\partial_a \mathbb{E}[J(\omega, a)] = \mathbb{E}[\partial_a J(\omega, a)] + \mathbb{E}[J(\omega, a)\partial_a \ln \text{Pr}(\omega, a)]. \quad (4)$$

The first term can be readily calculated given a particular form of  $J$  and a specific protocol. The second term, involving derivatives of the score function, requires special attention. We can compute this term from trajectory data by noting that the probability of observing a trajectory  $\omega$ , using the Onsager-Machlup action formalism, can be written as [27]

$$\text{Pr}(\omega, a) = \mathcal{N} \exp \left\{ -\frac{1}{2\sigma^2} \int_\omega [\dot{x} - \mu\Phi(x, a)]^2 dt \right\}, \quad (5)$$

where  $\mathcal{N}$  is a trajectory-independent constant that plays no role in  $\partial_a \ln \text{Pr}(\omega, a)$ :

$$\partial_a \ln \text{Pr}(\omega, a) = \frac{1}{\sigma^2} \int_\omega \mu \partial_a \Phi(x, a) [\dot{x} - \mu\Phi(x, a)] dt. \quad (6)$$

From the Langevin equation, we recognize that  $[\dot{x} - \mu\Phi(x, a)]dt = \sigma dw_t$ . Thus, plugging the above into Eq. (4) yields

$$\begin{aligned} \partial_a \mathbb{E}[J(\omega, a)] &= \mathbb{E}[\partial_a J(\omega, a)] + \mathbb{E} \left[ J(\omega, a) \frac{1}{\sigma} \int_\omega \mu \partial_a \Phi(x, a) \cdot dw_t \right]. \end{aligned} \quad (7)$$

This is the FRR, which relates the thermal noise in the observed trajectories to the response function. It can also be derived from the Girsanov theorem [28]. The FRR allows us to easily access the gradient of observables with respect to control parameters through trajectory simulation.

## III. LOSS FUNCTIONS IN COMPUTATIONS

Having introduced the essentials behind the FRR, we now associate the observable  $J$  with a loss function designed to be minimal at some desired behavior. From this perspective, the FRR is a flexible and powerful tool to apply machine learning methods to noisy physical systems: We can use it to learn control protocols that push a system of interest to behave in targeted ways. We now discuss how to design a loss function specifically useful for computational tasks. To begin with, we must establish how information is encoded within a physical system.

Consider a classical continuous  $n$ -dimensional system state  $\mathcal{S} \in \mathbb{R}^n$ . Encoding binary information requires coarse-graining  $\mathcal{S}$ —partitioning its continuous space into distinct regions that represent discrete-value states. Mathematically, a coarse-graining  $\sigma$  is a many-to-one map from a distribution  $p$  over system  $\mathcal{S}$  to a distribution  $p_{\mathcal{I}}$  over an informational space  $\mathcal{I}$ . The informational space  $\mathcal{I}$  is a finite-dimensional space spanned by binary values; i.e.,  $b = \{|00000\rangle, |00001\rangle, \dots\}$ . After coarse graining, the degrees of freedom over this finite-dimensional space are called information-bearing degrees of freedom or IBDs [23].

To have a stable information register for  $n$  binary digits, we ask the original potential to have  $2^n$  minima  $\{x_b\}$ , each of which represents  $n$  binary digits. A simple example is a double-well potential  $V(x) = \frac{1}{4}x^4 - \frac{1}{2}x^2$  with two stable minima located at  $x^* = \pm 1$ . Each minimum can be used to store

a binary digit. We use the phase space regions surrounding  $x = -1$  and  $x = +1$  to represent  $|0\rangle$  and  $|1\rangle$ , respectively. The coarse-graining map sends any distribution  $f(x)$  over  $\mathbb{R}$  to the distribution  $p_{\mathcal{I}} = (p_{|0\rangle}, p_{|1\rangle})$ , where  $p_{|0\rangle} = \int_{-\infty}^0 p(x)dx$  and  $p_{|1\rangle} = \int_0^{+\infty} p(x)dx$ .

Despite the eventual goal of digital computation, we will not derive effective dynamics describing the distribution evolution over IBD but rather optimize the dynamics of the continuous underlying physical system  $\mathcal{S}$ . The aim is to design protocols such that the system's final distribution  $f(x_\tau)$  maps to the correct distribution  $p_{\mathcal{I}_{\text{target}}}$  with minimal energetic cost. To do this, we prioritize achieving the target distribution over IBD, minimizing the work cost simultaneously. This motivates the loss function:

$$\mathcal{L}_{\text{loss}}(a) = \mathbb{E}[\alpha_1 \mathcal{L}_{\text{error}}(\omega, a) + \alpha_W W(\omega, a)], \quad (8)$$

where  $\mathcal{L}_{\text{error}}$  represents an appropriate error measurement in  $\mathcal{S}$  and where  $\alpha_1$  and  $\alpha_W$  are hyperparameters.

Suppose we wish to perform a computational task that sends bit value  $b_i$  to  $g(b_i)$ . We ask that, in the physical system  $\mathcal{S}$ , all trajectories starting around the minimum  $x_{b_i}$  representing  $b_i$  end around the minimum  $x_{g(b_i)}$  representing  $g(b_i)$ . To enforce this behavior, we choose

$$\mathcal{L}_{\text{error}} = (x_\tau - x_{g(b_i)})^2, \quad \text{if } x_0 \text{ near } x_{b_i}. \quad (9)$$

We add higher-order moments  $\mathbb{E}[(x_\tau - x_{g(b_i)})^n]$  to the loss function if the L2 norm is insufficient.

Before diving into specific examples, we should discuss the optimization framework underlying this intuitively motivated loss function. The straightforward approach to achieve energetically efficient operations that meet an accuracy threshold is to set an upper bound on error rate  $\epsilon$  and then locate the most energetically efficient protocol with error being no greater than  $\epsilon$ . In optimization, this is the *primal problem* [29]:

$$\underset{a}{\text{minimize}} \quad \mathbb{E}[W(\omega, a)] \quad \text{subject to} \quad \mathbb{E}[\mathcal{L}_{\text{error}}(\omega, a)] - \epsilon \leq 0.$$

The corresponding Lagrangian *dual problem* is

$$\underset{\lambda}{\text{maximize}} \quad \inf_a (\mathbb{E}[W(\omega, a)] + \lambda(\mathbb{E}[\mathcal{L}_{\text{error}}(\omega, a)] - \epsilon)),$$

subject to  $\lambda \geq 0$ .

Suppose the optimal value of the primal problem (the optimal work cost with error being no greater than  $\epsilon$ ) is  $p^*$  and that of the Lagrangian dual problem is  $d^*$ . If strong duality holds, these two problems are equivalent to each other  $p^* = d^*$ ; i.e., for any  $\epsilon$ , we can find a  $\lambda$  such that these two optimizations are equivalent. If we do not have strong duality, then we can simply pick a  $\lambda$  to optimize over and get a lower bound on the primal problem energetic cost  $d^* \leq p^*$ .

*Example: Overdamped Harmonic Trap Translation.* We start with a well-known model for which the optimal protocol can be solved exactly: moving a harmonic trap from 0 to  $a_f$  over the time duration  $\tau$ . The potential is  $V(x, t) = (x - a_t)^2/2$ , where  $a_t$  is the center of the harmonic well and we have full control of  $a_t$  with boundaries fixed at  $a_0 = 0$  and  $a_\tau = a_f$ . The corresponding stochastic differential equation of a unit-mobility particle is

$$dx = -(x - a_t)dt + \sigma dw_t. \quad (10)$$

We assume the system is in equilibrium at  $t = 0$ .




FIG. 1. Optimal protocol for translating a harmonic trap: We fix  $a_f = 1$  and  $\tau = 1$  and use ten break points. For one training step, we simulate 5000 trajectories with protocol value  $\mathbf{a} = \{a_1, \dots, a_{10}\}$  to compute the gradient of the loss functions and use the gradient to update  $\mathbf{a}$ . We train the protocol for 600 iterations and numerically show different types of jumps for different weight combinations.

In addition to minimizing the work, the particle must end up as close as possible to the final target position  $a_f$ . The loss function is taken to be

$$\mathcal{L}_{\text{move}}(a_t) = \alpha_1 \mathbb{E}[(x_\tau - a_f)^2] + \alpha_W \mathbb{E}[W(\omega, a_t)]. \quad (11)$$

Our goal now is to find the control protocol  $a_t^*$  that minimizes the loss function  $\mathcal{L}_{\text{move}}$ . This can be done analytically using the method of Lagrange multipliers (see the Appendix). In short, the optimal protocol is that  $a_t$  is a linear function of time  $t$  except at times  $t = 0$  and  $t = \tau$ .

For numerical optimization, we parametrize the protocol values  $a_t$  by a continuous piecewise-linear function with parameters that are the break points  $\mathbf{a} = \{a_1, \dots, a_n\}$ . We calculate the gradient  $\partial_a \mathcal{L}_{\text{move}}(\mathbf{a})$  with simulated trajectories in each training iteration and update  $\mathbf{a}$  for the next iteration. Figure 1 shows the training results with different weights in the loss function.

*Example: Bit Erasure.* We now move into the domain of optimizing computations by considering single-bit erasure. Bit erasure is a fundamental computational task that maps the uniform distribution over IBD  $p_{\mathcal{I}_0} = (1/2, 1/2)$  to a nearly deterministic state  $p_{\mathcal{I}_{\text{target}}} = (\epsilon, 1 - \epsilon)$ . To implement bit erasure, we use a quartic potential with three parameters  $\mathbf{a}_t = \{a_t, b_t, c_t\}$ :

$$U(x, t) = a_t x^4 - b_t x^2 + c_t x. \quad (12)$$

The parameter boundary conditions are cyclic, with  $a_0 = a_\tau = a$ ,  $b_0 = b_\tau = b$ ,  $c_0 = c_\tau = 0$ , and  $2a = b$ . Parameters  $a_t$  and  $b_t$  control equilibrium positions  $x_{\text{eq}} = \pm\sqrt{(b_t/2a_t)}$  and  $c_t$  tilts the potential. With these boundary conditions imposed,




FIG. 2. Bit erasure optimization: For each training iteration, we simulate 5000 trajectories to compute the gradient and update parameters. We validate with  $10^5$  trajectories to estimate work cost and error. The number of error trajectories almost vanishes after 100 iterations. Three trajectory ensembles are selected from a random training. The plot also includes the upper and lower bounds of the minimal work cost in finite-time bit erasure under full control of the potential [9], denoted with a gray highlight.

the potential  $U$  can encode a binary value as it has two stable equilibrium points  $x = \pm 1$ . The particle is at equilibrium with potential  $U(x, 0)$  at temperature  $T$  initially. At time  $t = \tau$ , all particles are required to be around  $x = +1$ , so we choose the loss function for the bit erasure task to be

$$\mathcal{L}_{\text{erase}}(\mathbf{a}_t) = \mathbb{E}[\alpha_1(x_\tau - 1)^2 + \alpha_W W(\omega, \mathbf{a}_t)]. \quad (13)$$

In numerical experiments, we set  $\alpha_1 = 1$ , work weight  $\alpha_W = 0.1$ ,  $\mu = 1$ ,  $k_B T = 1$ , and  $\tau = 1$ . Figures 2–4 show the numerical optimization results.

In the quasistatic limit, efficient erasure operates as follows. First, lower the energy barrier; next, tilt the potential; finally, raise the energy barrier and remove the tilt. In our optimized finite-time protocol, we see that lowering the barrier and tilting the potential happen all at once at  $t = 0^+$ . Similar jumps in optimal protocols are also found in Refs. [9,10]. The linear potential  $c_t x$ , used to drive the particle toward  $x = +1$ , dominates after  $t = 0.33$  and  $U$  is almost linear in the region of interest ( $x \in [-1, 1]$ ). After  $t = 0.33$ ,  $c_t$  decreases




FIG. 3. Error rate and work dependence on the hyperparameter ratio in erasure: Around  $\alpha_W/\alpha_1 = 0.47$ , we observe a sharp rise in error rate. The inset is error rate and work vs  $(\alpha_W/\alpha_1 - 0.47)$  in log scale, indicating a second-order phase transition.

quickly to around  $-30$ .  $c_1^- \simeq -30$  ensures that the potential  $U(x, t = 1^-)$  has minima around  $x = +1$ , aligned with the minima at the fixed boundary conditions. The work cost after 2000 iterations is  $\tilde{W}_{\min} = (2.942 \pm 0.026)k_B T$ .

To compare our result with other attempts at bit erasure, we compute the ratio  $r = (\tilde{W}_{\min} - \ln 2)/(\text{var}[x_0]/D\tau) = 2.311 \pm 0.027$ . The ratios of other bit erasure frameworks [5,11,30–34] range from 2.89 to 5.67. (See the detailed table in Ref. [9].) While we cannot analytically calculate the optimal protocol for this setting, we see that the training converges to a work cost that is close to the upper bound  $r = 2$  for optimal erasure under full control of the potential [9]. Given that our control is limited by both our choice of piecewise-linear protocol and the form of the potential and the closeness of our distribution at  $t = \tau$  to the local equilibrium distribution, this proximity to the work cost for unrestricted optimal control is evidence that our learning algorithm is finding truly efficient protocols.

We end this example by a brief discussion of the optimal protocol's dependence on the hyperparameter ratio  $\alpha_W/\alpha_1$ . Increasing the ratio  $\alpha_W/\alpha_1$  can lower the work at the cost of higher error rates. To understand this, consider the following two reference protocols. In Protocol 1, the potential  $U(x, t)$  remains unchanged during  $[0, \tau]$  so that  $\mathbb{E}[W] \sim 0$ , but the error is  $\mathbb{E}[\mathcal{L}_{\text{error}}] \sim 0^2 \times \frac{1}{2} + 2^2 \times \frac{1}{2} = 2$ . Protocol 2 is a faithful erasure for which  $\mathbb{E}[W] \sim \ln 2 + 2\text{var}[x_0]/(D\tau) = 2.69$  [9] and  $\mathbb{E}[\mathcal{L}_{\text{error}}] \sim 0$ . Around  $\alpha_W/\alpha_1 \sim 0.7$ , the two protocols are equally favorable, with Protocol 2 being more favorable below this value. Numerically, in Fig. 3, starting from  $\alpha_W/\alpha_1 = 0.1$ , we observe a sharp transition from Protocol 2 to Protocol 1 around 0.47, strongly suggesting a second-order phase transition, which we leave for detailed future study.

*Example: Underdamped-Dynamics Optimization and the Bit Flip.* Let us turn to explore optimization in underdamped dynamics, a regime that is known to be challenging for existing methods. Position trajectories are no longer Markovian in underdamped dynamics and this property can be exploited to accomplish certain computational tasks that are unachievable in Markovian position dynamics (overdamped) without additional hidden states [35]. An illustrative computational task is bit flip in one dimension. In this, trajectories originating from  $-1$  and those from  $+1$  must converge near the middle before reaching their respective final logical values. If the dynamics are purely Markovian, the computation loses the ability to distinguish between these two trajectory classes once they meet, making the bit-flip operation infeasible. In contrast, underdamped dynamics has an additional degree of freedom through momentum, allowing one to distinguish histories and thus implement bit-flip operations effectively.

Consider one-dimensional underdamped dynamics:

$$dx = v dt \quad \text{and} \quad mdv = -\gamma v - \partial_x U(x, a) dt + \sigma dw_t,$$

where  $\gamma$  is the damping coefficient and the thermal noise standard deviation is related to the thermal bath temperature  $T$  by  $\sigma = \sqrt{2\gamma k_B T}$ . The probability of observing a trajectory  $\omega = \{(x_t, v_t)\}_{t=0}^\tau$  with parameter  $a$  is

$$\Pr(\omega, a) = \mathcal{N} \exp \left\{ -\frac{1}{2\sigma^2} \int_0^\tau \left[ m \frac{dv}{dt} - \gamma v - \Phi(x, a) \right]^2 dt \right\}.$$




FIG. 4. Optimal erasure protocol: For each protocol parameter  $\{a_t, b_t, c_t\}$ , we use ten break points for parametrization. (a) They are initialized to 0. (b) Final protocol values after training for 1000 iterations. (c) Potentials  $U$  at different times.

The FRR derived above in Eq. (7) follows directly from the chain rule and so the considerations above work just as well in underdamped regimes. Nothing in the optimization method changes under underdamped dynamics because  $\partial_\alpha \ln \text{Pr}(\omega, a)$  takes exactly the same form here as in Eq. (7).

The theoretical lower bound cost is  $0 k_B T$  as the overall system entropy does not change during the bit-flip operation. Practically speaking, though, energy is dissipated in a finite-time protocol if there is nonzero damping. A more appropriate approximate lower bound is the energy dissipated during constant velocity transport across a distance  $d$  from the bottom of the left well to the bottom of the right well:  $W_{\min} \sim \gamma d^2 / \tau$ .

To compare to our optimization, we use a heuristic bit-flip protocol introduced in Refs. [36,37] and physically implemented in Ref. [12]: By switching to an appropriately chosen harmonic potential at  $t = 0^+$ , the two clusters of particles exchange positions after one half-period of oscillation. We use the same potential, parameters  $\mathbf{a}_t = \{a_t, b_t\}$ , boundary conditions, and training initialization as in the bit erasure task except the linear tilting term  $c_t$  is removed. The error loss is chosen to be

$$\mathcal{L}_{\text{flip}}(\mathbf{a}_t) = \mathbb{E}[\alpha_1(x_\tau - \text{target}(x_0))^2 + \alpha_W W(\omega, \mathbf{a}_t)], \quad (14)$$

where  $\text{target}(x_0) = +1$  if  $x_0 < 0$  and  $\text{target}(x_0) = -1$  if  $x_0 > 0$ . In numerical experiments, we set  $\alpha_1 = 1$ , work weight

$\alpha_W = 0.1$ ,  $\gamma = 0.1$ , and  $\tau = 1$ . Figures 5 and 6 present the training results. The work cost of our converged protocol is  $\approx (0.615 \pm 0.070)k_B T$ . This is close to the lower bound on the work cost based on constant velocity transport:  $W_{\min} \sim 0.4k_B T$ .

#### IV. DISCUSSION AND CONCLUSION

Shifting from distributions to trajectories in contemporary nonequilibrium physics has led to many exciting results [38–76]. Building on this, we introduced a general framework for designing low-cost protocols for computational tasks embedded in physical systems. Given a computational task, we used the linear combination of work cost and statistical moments to construct the loss function and FRRs to compute the gradient and optimize the protocols. We demonstrated that the first- and second-order statistical moments within the loss function effectively guide the system distribution toward the desired configuration in computational tasks.

More broadly, we used the Feynman-Kac formula to transform the optimization in partial differential equations to an optimization over trajectories. This reformulation offers a key advantage: Simulating noisy trajectories is often much easier than solving the corresponding equations for the optimal protocol, either numerically or analytically. Our framework also shares similarity with reinforcement learning: A system (or




FIG. 5. (a) Bit-flip mean work: Initialization of the potential is  $U(x, t) = 0$  for  $t \in (0, 1)$ . The gray dashed line is the work cost of the protocol proposed in Ref. [36], which is  $0.725 k_B T$ . Trajectories are chosen from a single random training. (b) Phase space trajectories: Two trajectory classes in the whole phase space—one cluster starts near  $-1$  and the other starts near  $+1$ .

agent) learns to optimize a protocol (or policy) by minimizing a loss function (or maximizing a reward signal) based on interactions with an environment [77].

The error penalty used in Eq. (9) is not unique. For example, the sign function  $\text{sgn}(x)$  could also be used in the error penalty. Generally, this choice should depend on the specific requirements in the computational task. For example, if the final distribution  $p_\tau$  should be exactly  $p_{\text{target}}$ , then a distribution distance as error penalty is needed; one could use the Kullback-Liebler (KL) divergence  $\mathbb{E}[\ln(p_\tau/p_{\text{target}})]$ . However, computing the gradient of the  $\ln$  distribution  $\partial_a \ln p_\tau$  is typically challenging. A possible solution is to discretize the Fokker-Planck equation in time and iteratively solve for the distribution  $p(x, t = i\Delta t; a)$  at each time step. Then, one can use backpropagation to compute the gradient of KL divergence with respect to  $a$ . Within the context of computational tasks, we observed that a limited number of statistical moments are effective for protocol design.

Notably, FRRs have regained attention in recent years as studies found that they can bound the entropy production and dynamical activity in stochastic systems [25,78–88]. In line with this, our method reveals an application of FRRs—they are a tool to design low-cost protocols for probability-transform tasks within physical systems. The FRR-based gradient yields results consistent with backpropagation and genetic algorithms [13,20,89–93]. Perhaps equally important, the FRR method (1) offers a more physical interpretation, (2)




FIG. 6. (a) For each protocol value  $\{a_t, b_t\}$  in the double-well potential, we use 20 break points for parametrization. They are initialized the same way as in bit erasure. After training on 1000 iterations, we show the final protocol. (b) The potential jumps at  $t = 0^+$  and  $t = 1^-$ .

requires significantly less memory storage, and (3) is relatively easy to implement in code. Additionally, the generality of observables  $J$  for which the FRR can provide a gradient means that the proposed framework can be extended well beyond information-processing tasks simply by swapping out the loss function to encourage any behavior of choice.

The optimization framework presented here offers several promising future research directions and applications. First, it opens low-cost protocol design to different domains, including unconventional computation [94–98], efficiency in information and nanoscale engines [99,100], chemical reactions [101], biophysical systems [102], quantum gates [103], and the like. Second, having proved out piecewise-linear functions as our learning model, one future direction is to replace this with neural networks for enhanced scalability [104]. Third, we fixed the computation-time duration, but one can also treat the total computation time  $\tau$  as a protocol parameter. Adding this to a loss function will enhance discovering protocols that balance duration along with error and work cost [105].

## ACKNOWLEDGMENTS

The authors thank Yibo Jacky Zhang for illuminating discussions on machine learning, as well as the Telluride Science and Innovation Center for its hospitality during visits and the participants of the Information Engines workshop there for their valuable feedback. This material is based on work supported by, or in part by, the U.S. Army Research Laboratory and U.S. Army Research Office under Grant No. W911NF-21-1-0048 and the Art and Science Laboratory.

## DATA AVAILABILITY

The data that support the findings of this article are openly available [92]; embargo periods may apply.

### APPENDIX: OVERDAMPED HARMONIC TRAP TRANSLATION

This Appendix gives the details needed to obtain the optimal protocol in overdamped harmonic trap translation. We denote the position mean as  $\mathbb{E}[x_t] = \mu_t$ . First note that, by averaging the Langevin equation (10), we arrive at a simple constraint between  $\mu_t$  and  $a_t$ :

$$\dot{\mu}_t + \mu_t = a_t.$$

From this, it is straightforward to see how the position mean at  $t = \tau$  can be written in closed form:

$$\mu_\tau \equiv \mathbb{E}[x_\tau] = \int_0^\tau e^{-(\tau-t)} a_t dt,$$

with the mean work being

$$\mathbb{E}[W] = \int_0^\tau (a_t - \mu_t) \dot{a}_t dt.$$

The loss function is

$$\mathcal{L}_{\text{move}}(a_t) = \alpha_1 \mathbb{E}[(x_\tau - a_f)^2] + \alpha_w \mathbb{E}[W(\omega, a_t)].$$

A helpful property of a linear force is that the variance of the distribution does not change with time  $t$  if the distribution starts in equilibrium. So, we can safely replace  $\mathbb{E}[(x_\tau - a_f)^2]$  with  $(\mathbb{E}[x_\tau] - a_f)^2$  up to a constant. Adding a Lagrange multiplier  $\lambda_t$  to enforce the constraint, and using the expressions for the means above, the function we want to minimize is

$$\begin{aligned} \mathcal{L} = & \alpha_1 \left( \int_0^\tau e^{-(\tau-t)} a_t dt - a_f \right)^2 \\ & + \alpha_w \int_0^\tau (a_t - \mu_t) \dot{a}_t dt + \int_0^\tau \lambda_t (\dot{\mu}_t + \mu_t - a_t) dt. \end{aligned}$$

The Euler-Lagrange equations lead to three conditions:

$$\begin{aligned} a_t : & 2\alpha_1 e^{-(\tau-t)} \left( \int_0^\tau e^{-(\tau-s)} a_s ds - a_f \right) + \alpha_w \dot{\mu}_t - \lambda_t = 0, \\ \mu_t : & -\alpha_w \dot{a}_t + \lambda_t - \dot{\lambda}_t = 0, \text{ and} \\ \lambda_t : & \dot{\mu}_t + \mu_t - a_t = 0. \end{aligned}$$

We can eliminate the  $\alpha_1$  term in the  $a_t$  equation by solving for  $\lambda_t$ , and then subtracting the time derivative:

$$\dot{\lambda}_t = 2\alpha_1 e^{-(\tau-t)} \left( \int_0^\tau e^{-(\tau-s)} a_s ds - a_f \right) + \alpha_w \ddot{\mu}_t.$$

This yields the expression

$$\lambda_t - \dot{\lambda}_t = \alpha_w \dot{\mu}_t - \alpha_w \ddot{\mu}_t. \quad (\text{A1})$$

Taking the time derivative of the  $\lambda_t$  equation and plugging it into the  $\mu_t$  equation yields

$$\lambda_t - \dot{\lambda}_t = \alpha_w \dot{\mu}_t + \alpha_w \ddot{\mu}_t. \quad (\text{A2})$$

Thus, we find that  $\ddot{\mu}_t = 0$ . The initial condition  $\mu_0 = 0$ , then, implies  $\mu_t = mt$ , which reduces our constraint condition to  $\lambda_t$  equation  $a_t = mt + m$ . The optimal constant  $m^*$  depends on the hyperparameters, but can be determined exactly by minimizing the loss function  $\mathcal{L}_{\text{move}}$  under the assumption of a linear  $a_t$ . The minimization yields

$$m^* = a_f / (\tau + c), \quad \text{with} \quad c = \frac{2\alpha_w}{2\alpha_1 + \alpha_w}. \quad (\text{A3})$$

Finally, we calculate the means of the final position and work when  $m = m^*$ :

$$\mathbb{E}[x_\tau] = \frac{a_f}{1 + c/\tau}, \quad \mathbb{E}[W] = a_f^2 \frac{\tau + c^2/2}{(\tau + c)^2}. \quad (\text{A4})$$

We can see the effect of hyperparameters: Increasing  $\alpha_1$  (decreasing  $c$ ) drives the system closer to the target value  $a_f$ . We recover a known result that takes only the work cost into account [3] when  $\alpha_1 \rightarrow 0, \alpha_w \rightarrow 1$ . As is often the case, the optimal protocol has jumps at both ends.

- 
- [1] R. Landauer, Irreversibility and heat generation in the computing process, *IBM J. Res. Dev.* **5**, 183 (1961).
  - [2] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012).
  - [3] T. Schmiedl and U. Seifert, Optimal finite-time processes in stochastic thermodynamics, *Phys. Rev. Lett.* **98**, 108301 (2007).
  - [4] A. Gomez-Marín, T. Schmiedl, and U. Seifert, Optimal protocols for minimal work processes in underdamped stochastic thermodynamics, *J. Chem. Phys.* **129**, 024114 (2008).
  - [5] P. R. Zulkowski and M. R. DeWeese, Optimal finite-time erasure of a classical bit, *Phys. Rev. E* **89**, 052140 (2014).
  - [6] T. R. Gingrich, G. M. Rotskoff, G. E. Crooks, and P. L. Geissler, Near-optimal protocols in complex nonequilibrium transformations, *Proc. Natl. Acad. Sci. USA* **113**, 10263 (2016).
  - [7] A. Dechant and Y. Sakurai, Thermodynamic interpretation of Wasserstein distance, [arXiv:1912.08405](https://arxiv.org/abs/1912.08405).
  - [8] A. Zhong and M. R. DeWeese, Limited-control optimal protocols arbitrarily far from equilibrium, *Phys. Rev. E* **106**, 044135 (2022).
  - [9] K. Proesmans, J. Ehrich, and J. Bechhoefer, Finite-time Landauer principle, *Phys. Rev. Lett.* **125**, 100602 (2020).
  - [10] K. Proesmans, J. Ehrich, and J. Bechhoefer, Optimal finite-time bit erasure under full control, *Phys. Rev. E* **102**, 032105 (2020).
  - [11] A. B. Boyd, A. Patra, C. Jarzynski, and J. P. Crutchfield, Shortcuts to thermodynamic computing: The cost of fast and faithful information processing, *J. Stat. Phys.* **187**, 17 (2022).
  - [12] S. Dago and L. Bellon, Logical and thermodynamical reversibility: Optimized experimental implementation of the NOT operation, *Phys. Rev. E* **108**, L022101 (2023).
  - [13] M. C. Engel, J. A. Smith, and M. P. Brenner, Optimal control of nonequilibrium systems through automatic differentiation, *Phys. Rev. X* **13**, 041032 (2023).

- [14] S. Whitelam, Demon in the machine: Learning to extract work and absorb entropy from fluctuating nanosystems, *Phys. Rev. X* **13**, 021005 (2023).
- [15] R. Sabbagh, O. Movilla Miangolarra, and T. T. Georgiou, Wasserstein speed limits for Langevin systems, *Phys. Rev. Res.* **6**, 033308 (2024).
- [16] L. K. Davis, K. Proesmans, and É. Fodor, Active matter under control: Insights from response theory, *Phys. Rev. X* **14**, 011012 (2024).
- [17] J. Sanders, M. Baldovin, and P. Muratore-Ginanneschi, Optimal control of underdamped systems: An analytic approach, *J. Stat. Phys.* **191**, 117 (2024).
- [18] S. A. Loos, S. Monter, F. Ginot, and C. Bechinger, Universal symmetry of optimal control at the microscale, *Phys. Rev. X* **14**, 021032 (2024).
- [19] J. Klinger and G. M. Rotskoff, Minimally dissipative multibit logical operations, *Phys. Rev. E* **113**, 014109 (2026).
- [20] D. Rengifo and G. Téllez, Machine learning approach to fast thermal equilibration, *Phys. Rev. E* **111**, 065311 (2025).
- [21] J. Sanders, M. Baldovin, and P. Muratore-Ginanneschi, Minimal work protocols for inertial particles in nonharmonic traps, *Phys. Rev. E* **111**, 034127 (2025).
- [22] D. Gupta, K. S. Olsen, and S. Krishnamurthy, Thermodynamic cost of recurrent erasure, *Commun. Phys.* **8**, 355 (2025).
- [23] S. Deffner and C. Jarzynski, Information processing and the second law of thermodynamics: An inclusive, Hamiltonian approach, *Phys. Rev. X* **3**, 041003 (2013).
- [24] P. B. Warren and R. J. Allen, Malliavin weight sampling for computing sensitivity coefficients in Brownian dynamics simulations, *Phys. Rev. Lett.* **109**, 250601 (2012).
- [25] J. Zheng and Z. Lu, Universal response inequalities beyond steady states via trajectory information geometry, *Phys. Rev. E* **112**, L012103 (2025).
- [26] J. Klinger and G. M. Rotskoff, Computing nonequilibrium responses with score-shifted stochastic differential equations, *Phys. Rev. Lett.* **134**, 097101 (2025).
- [27] L. Onsager and S. Machlup, Fluctuations and irreversible processes, *Phys. Rev.* **91**, 1505 (1953).
- [28] R. S. Liptser and A. N. Shiryaev, *Statistics of Random Processes II: Applications* (Springer, Berlin, Heidelberg, 2013), Vol. 6.
- [29] S. P. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, 2004).
- [30] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, Experimental verification of Landauer's principle linking information and thermodynamics, *Nature (London)* **483**, 187 (2012).
- [31] A. Bérut, A. Petrosyan, and S. Ciliberto, Detailed Jarzynski equality applied to a logically irreversible procedure, *Europhys. Lett.* **103**, 60002 (2013).
- [32] Y. Jun, M. Gavrilov, and J. Bechhoefer, High-precision test of Landauer's principle in a feedback trap, *Phys. Rev. Lett.* **113**, 190601 (2014).
- [33] A. Bérut, A. Petrosyan, and S. Ciliberto, Information and thermodynamics: Experimental verification of Landauer's erasure principle, *J. Stat. Mech.* (2015) P06015.
- [34] M. Gavrilov, R. Chétrite, and J. Bechhoefer, Direct measurement of weakly nonequilibrium system entropy is consistent with Gibbs–Shannon form, *Proc. Natl. Acad. Sci. USA* **114**, 11097 (2017).
- [35] J. A. Owen, A. Kolchinsky, and D. H. Wolpert, Number of hidden states needed to physically implement a given conditional distribution, *New J. Phys.* **21**, 013022 (2019).
- [36] K. J. Ray, A. B. Boyd, G. W. Wimsatt, and J. P. Crutchfield, Non-Markovian momentum computing: Thermodynamically efficient and computation universal, *Phys. Rev. Res.* **3**, 023164 (2021).
- [37] K. J. Ray and J. P. Crutchfield, Gigahertz sub-Landauer momentum computing, *Phys. Rev. Appl.* **19**, 014049 (2023).
- [38] C. Maes, Response theory: A trajectory-based approach, *Front. Phys.* **8**, 229 (2020).
- [39] S. B. Nicholson, L. P. García-Pintos, A. del Campo, and J. R. Green, Time–information uncertainty relations in thermodynamics, *Nat. Phys.* **16**, 1211 (2020).
- [40] C. Dieball and A. Godec, Direct route to thermodynamic uncertainty relations and their saturation, *Phys. Rev. Lett.* **130**, 087101 (2023).
- [41] C. Dieball and A. Godec, Feynman-Kac theory of time-integrated functionals: Itô versus functional calculus, *J. Phys. A* **56**, 155002 (2023).
- [42] C. Dieball and A. Godec, Thermodynamic bounds on generalized transport: From single-molecule to bulk observables, *Phys. Rev. Lett.* **133**, 067101 (2024).
- [43] P. B. Melo, S. M. Duarte Queirós, and W. A. Morgado, Stochastic thermodynamics of Fisher information, *Phys. Rev. E* **111**, 014101 (2025).
- [44] P. B. Melo, F. Iemini, D. O. Soares-Pinto, S. M. D. Queirós, and W. A. Morgado, Thermodynamic interpretation to stochastic Fisher information and single-trajectory speed limits, *Phys. Rev. E* **112**, 014126 (2025).
- [45] C. Dieball and A. Godec, Perspective: Time irreversibility in systems observed at coarse resolution, *J. Chem. Phys.* **162**, 090901 (2025).
- [46] C. Dieball, Y. M. Satahari, A. B. Zuccolotto-Bernez, S. U. Egelhaaf, M. A. Escobedo-Sánchez, and A. Godec, Precisely controlled colloids: A playground for path-wise nonequilibrium physics, *Soft Matter* **21**, 2170 (2025).
- [47] C. Dieball and A. Godec, Thermodynamic correlation inequalities for finite times and transients, *J. Phys. A* **58**, 255001 (2025).
- [48] G. Fiusa, P. E. Harunari, A. S. Hegde, and G. T. Landi, Counting observables in stochastic excursions, [arXiv:2505.06208](https://arxiv.org/abs/2505.06208).
- [49] G. Fiusa, P. E. Harunari, A. S. Hegde, and G. T. Landi, A framework for fluctuating times and counting observables in stochastic excursions, [arXiv:2506.05160](https://arxiv.org/abs/2506.05160).
- [50] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015).
- [51] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, *Phys. Rev. Lett.* **116**, 120601 (2016).
- [52] J. M. Horowitz and T. R. Gingrich, Proof of the finite-time thermodynamic uncertainty relation for steady-state currents, *Phys. Rev. E* **96**, 020103(R) (2017).
- [53] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nat. Phys.* **16**, 15 (2020).

- [54] K. Liu, Z. Gong, and M. Ueda, Thermodynamic uncertainty relation for arbitrary initial states, *Phys. Rev. Lett.* **125**, 140602 (2020).
- [55] T. Koyuk and U. Seifert, Thermodynamic uncertainty relation for time-dependent driving, *Phys. Rev. Lett.* **125**, 260604 (2020).
- [56] J. S. Lee, J.-M. Park, and H. Park, Universal form of thermodynamic uncertainty relation for Langevin dynamics, *Phys. Rev. E* **104**, L052102 (2021).
- [57] Z. Cao, J. Su, H. Jiang, and Z. Hou, Effective entropy production and thermodynamic uncertainty relation of active Brownian particles, *Phys. Fluids* **34**, 053310 (2022).
- [58] P. P. Potts and P. Samuelsson, Thermodynamic uncertainty relations including measurement and feedback, *Phys. Rev. E* **100**, 052137 (2019).
- [59] P. Pietzonka, A. C. Barato, and U. Seifert, Universal bounds on current fluctuations, *Phys. Rev. E* **93**, 052145 (2016).
- [60] M. Polettini, A. Lazarescu, and M. Esposito, Tightening the uncertainty principle for stochastic currents, *Phys. Rev. E* **94**, 052104 (2016).
- [61] K. Proesmans and C. Van den Broeck, Discrete-time thermodynamic uncertainty relation, *Europhys. Lett.* **119**, 20001 (2017).
- [62] A. Dechant and S.-I. Sasa, Current fluctuations and transport efficiency for general Langevin systems, *J. Stat. Mech.* (2018) 063209.
- [63] A. C. Barato, R. Chetrite, A. Faggionato, and D. Gabrielli, Bounds on current fluctuations in periodically driven systems, *New J. Phys.* **20**, 103023 (2018).
- [64] K. Macieszczak, K. Brandner, and J. P. Garrahan, Unified thermodynamic uncertainty relations in linear response, *Phys. Rev. Lett.* **121**, 130601 (2018).
- [65] K. Brandner, T. Hanazato, and K. Saito, Thermodynamic bounds on precision in ballistic multiterminal transport, *Phys. Rev. Lett.* **120**, 090601 (2018).
- [66] T. Koyuk and U. Seifert, Operationally accessible bounds on fluctuations and entropy production in periodically driven systems, *Phys. Rev. Lett.* **122**, 230601 (2019).
- [67] H.-M. Chun, L. P. Fischer, and U. Seifert, Effect of a magnetic field on the thermodynamic uncertainty relation, *Phys. Rev. E* **99**, 042128 (2019).
- [68] T. Van Vu and Y. Hasegawa, Uncertainty relations for time-delayed Langevin systems, *Phys. Rev. E* **100**, 012134 (2019).
- [69] A. Dechant, Multidimensional thermodynamic uncertainty relations, *J. Phys. A* **52**, 035001 (2019).
- [70] Y. Hasegawa and T. Van Vu, Fluctuation theorem uncertainty relation, *Phys. Rev. Lett.* **123**, 110602 (2019).
- [71] G. Guarnieri, G. T. Landi, S. R. Clark, and J. Goold, Thermodynamics of precision in quantum nonequilibrium steady states, *Phys. Rev. Res.* **1**, 033021 (2019).
- [72] K. Proesmans and J. M. Horowitz, Hysteretic thermodynamic uncertainty relation for systems with broken time-reversal symmetry, *J. Stat. Mech.* (2019) 054005.
- [73] T. Van Vu and Y. Hasegawa, Uncertainty relation under information measurement and feedback control, *J. Phys. A* **53**, 075001 (2020).
- [74] A. M. Timpanaro, G. Guarnieri, J. Goold, and G. T. Landi, Thermodynamic uncertainty relations from exchange fluctuation theorems, *Phys. Rev. Lett.* **123**, 090604 (2019).
- [75] Y. Hasegawa, Quantum thermodynamic uncertainty relation for continuous measurement, *Phys. Rev. Lett.* **125**, 050601 (2020).
- [76] K. J. Ray, A. B. Boyd, G. Guarnieri, and J. P. Crutchfield, Thermodynamic uncertainty theorem, *Phys. Rev. E* **108**, 054126 (2023).
- [77] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, 1998).
- [78] J. A. Owen, T. R. Gingrich, and J. M. Horowitz, Universal thermodynamic bounds on nonequilibrium response with biochemical applications, *Phys. Rev. X* **10**, 011066 (2020).
- [79] H.-M. Chun, Q. Gao, and J. M. Horowitz, Nonequilibrium Green-Kubo relations for hydrodynamic transport from an equilibrium-like fluctuation-response equality, *Phys. Rev. Res.* **3**, 043172 (2021).
- [80] T. Aslyamov and M. Esposito, Nonequilibrium response for Markov jump processes: Exact results and tight bounds, *Phys. Rev. Lett.* **132**, 037101 (2024).
- [81] K. Ptasiński, T. Aslyamov, and M. Esposito, Dissipation bounds precision of current response to kinetic perturbations, *Phys. Rev. Lett.* **133**, 227101 (2024).
- [82] Q. Gao, H.-M. Chun, and J. M. Horowitz, Thermodynamic constraints on kinetic perturbations of homogeneous driven diffusions, *Europhys. Lett.* **146**, 31001 (2024).
- [83] K. Liu and J. Gu, Dynamical activity universally bounds precision of response in Markovian nonequilibrium systems, *Commun. Phys.* **8**, 62 (2025).
- [84] T. Aslyamov, K. Ptasiński, and M. Esposito, Nonequilibrium fluctuation-response relations: From identities to bounds, *Phys. Rev. Lett.* **134**, 157101 (2025).
- [85] R. Bao and S. Liang, Nonequilibrium response theory: From precision limits to strong perturbation, *arXiv:2412.19602*.
- [86] T. Van Vu, Fundamental bounds on precision and response for quantum trajectory observables, *PRX Quantum* **6**, 010343 (2025).
- [87] E. Kwon, H.-M. Chun, H. Park, and J. S. Lee, Fluctuation-response inequalities for kinetic and entropic perturbations, *Phys. Rev. Lett.* **135**, 097101 (2025).
- [88] J. Zheng and Z. Lu, Nonlinear response relations and fluctuation-response inequalities for nonequilibrium stochastic systems, *J. Chem. Phys.* **164**, 184114 (2026).
- [89] S. Whitlam, How to train your demon to do fast information erasure without heat production, *Phys. Rev. E* **108**, 044138 (2023).
- [90] C. Casert and S. Whitlam, Learning protocols for the fast and efficient control of active matter, *Nat. Commun.* **15**, 9128 (2024).
- [91] S. Whitlam, Training thermodynamic computers by gradient descent, *Proc. Natl. Acad. Sci. USA* **123**, e2528413123 (2026).
- [92] S. Whitlam, C. Casert, M. Engel, and I. Tamblyn, Benchmarks for protocol control in nonequilibrium statistical mechanics, *arXiv:2506.15122*.
- [93] N. Barros, S. Whitlam, S. Ciliberto, and L. Bellon, Learning efficient erasure protocols for an underdamped memory, *Phys. Rev. E* **111**, 044114 (2025).
- [94] M. Aifer, K. Donatella, M. H. Gordon, S. Duffield, T. Ahle, D. Simpson, G. Crooks, and P. J. Coles, Thermodynamic linear algebra, *npj Unconv. Comput.* **1**, 13 (2024).

- [95] C. Z. Pratt, K. J. Ray, and J. P. Crutchfield, Controlled erasure as a building block for universal thermodynamically robust superconducting computing, *Chaos* **35**, 043112 (2025).
- [96] K. W. Tang, K. J. Ray, and J. P. Crutchfield, Nonequilibrium thermodynamics of a superconducting Szilard engine, *New J. Phys.* **27**, 043011 (2025).
- [97] C. Z. Pratt, K. J. Ray, and J. P. Crutchfield, Extracting equations of motion from superconducting circuits, *Phys. Rev. Res.* **7**, 013014 (2025).
- [98] Z. Zhang and Z. Lu, Kinetic frustration enables single-molecule computation, *J. Chem. Phys.* **162**, 134102 (2025).
- [99] A. B. Boyd and J. P. Crutchfield, Maxwell demon dynamics: Deterministic chaos, the Szilard map, and the intelligence of thermodynamic systems, *Phys. Rev. Lett.* **116**, 190601 (2016).
- [100] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Correlation-powered information engines and the thermodynamics of self-correction, *Phys. Rev. E* **95**, 012152 (2017).
- [101] S. S. Chittari and Z. Lu, Geometric approach to nonequilibrium hasty shortcuts, *J. Chem. Phys.* **159**, 084106 (2023).
- [102] A. I. Brown and D. A. Sivak, Theory of nonequilibrium free energy transduction by molecular machines, *Chem. Rev.* **120**, 434 (2020).
- [103] T. Van Vu, T. Kuwahara, and K. Saito, Fidelity-dissipation relations in quantum gates, *Phys. Rev. Res.* **6**, 033225 (2024).
- [104] D.-K. Kim, Y. Bae, S. Lee, and H. Jeong, Learning entropy production via neural networks, *Phys. Rev. Lett.* **125**, 140604 (2020).
- [105] T. Van Vu and K. Saito, Time-cost-error trade-off relation in thermodynamics: The third law and beyond, *Phys. Rev. X* **15**, 041029 (2025).