Optimal Computation from Fluctuation Responses

Jinghao Lyu,* Kyle J. Ray,† and James P. Crutchfield‡
Complexity Sciences Center and Physics and Astronomy Department, University of California at Davis, One Shields Avenue, Davis, CA 95616
(Dated: October 13, 2025)

We show how to design computation protocols that minimize thermodynamic cost while ensuring correct outcomes using fluctuation response relations (FRR) and machine learning. Unlike current approaches, our method simultaneously optimizes distributions and protocols, using FRR-derived gradients and learning iteratively from sampled trajectories. We design protocols for bit erasure in a double-well potential and translating harmonic traps. The framework extends to both overdamped and underdamped systems while achieving the theoretically optimal protocol and work costs comparable to relevant finite-time bounds.

Keywords: stochastic thermodynamics, computation, control protocol, entropy production

I. INTRODUCTION

Nowadays, computation—ranging from operating personal digital devices to training very large-scale machine learning models—consumes a significant amount of energy. This raises questions we long to answer: *how to perform computation 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 $k_{\mathrm{B}}T\log 2$, 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. As a key innovation in this, researchers successfully developed a stochastic dynamics for thermodynamic resources. This, in turn, 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_{target} over a time duration τ. 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 $p_0 \to p_{\mathrm{target}}$ 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 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 over-constrain 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 coarse-grained distribution to store information. In a computational task, we do not ask the final distribution in $\mathcal{S}$ to be exactly p_{target}. Instead, we only require the coarse-grained distribution at final time $p_{\tau_{\mathrm{cg}}}$ to match the coarse-grained target distribution $p_{\mathrm{target}_{\mathrm{cg}}}$, ensuring 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 (FRR) 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 close to Landauer's bound.

Our framework introduces several advantages: (1) Leveraging universal fluctuation-response relations, it can 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 target metrics, including logical error and energetic cost, in the loss function. (3) Since the learning algorithm optimizes via gradient descent and tracks samples of simulated trajectories, rather than solving a system of equations, it lends itself to algorithmic implementations which take advantage of modern machine learning tools that scale to high-dimension problems.

* jolyu@ucdavis.edu
† kjray@ucdavis.edu
‡ chaos@ucdavis.edu

II. FLUCTUATION RESPONSE RELATION

We first review the fluctuation-response relation (FRR) in the setting of overdamped dynamics [24–26]. Consider one-dimensional Langevin dynamics with an external control parameter a:

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

where μ is the particle mobility, dw_t is an infinitesimal Wiener process, $\sigma = \sqrt{2\mu k_{\rm B}T}$ 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:

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

where $\cdot$ is the Itô product. For example, the energetic cost (work) of one trajectory ω 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 then $\mathbb{E}\left[J(\omega,a)\right] = \sum_\omega \Pr(\omega,a)J(\omega,a)$, where $\Pr(\omega,a)$ is the probability of observing trajectory ω under protocol a.

The purpose of an FRR is to monitor how the expectation of current J changes due to variation in the control parameters: $\frac{\partial}{\partial a}\mathbb{E}\left[J(\omega,a)\right]$. From the chain rule:

$$\partial_a\mathbb{E}\left[J(\omega,a)\right] = \mathbb{E}\left[\partial_a J(\omega,a) + J(\omega,a)\partial_a \log\Pr(\omega,a)\right] . \tag{1}$$

The first term is readily calculated given the 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 ω, using the the Onsager-Machlup action formalism, can be written as [27]:

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

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

$$\partial_a\log\Pr(\omega,a) = \frac{1}{\sigma^2}\int_\omega \mu\partial_a\Phi(x,a)[\dot{x}-\mu\Phi(x,a)]dt \ . \tag{3}$$

From the Langevin equation, we recognize that $[\dot{x}-\mu\Phi(x,a)]dt = \sigma dw_t$. Thus, plugging the above into Eq. (1) yields:

$$\partial_a\mathbb{E}\left[J(\omega,a)\right] = \mathbb{E}\left[\partial_a J(\omega,a)\right] + \mathbb{E}\left[J(\omega,a)\frac{1}{\sigma}\int_\omega \mu\partial_a\Phi(x,a)\cdot dw_t\right] . \tag{4}$$

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 ready access to the gradient of observables with respect to control parameters through trajectory simulation.

III. LOSS FUNCTIONS IN COMPUTATIONS

Having introduced the FRR essentials, we now associate the observable J with a loss function designed to be minimal for a desired behavior. From this perspective, the FRR is a flexible and powerful tool for applying machine learning methods to noisy physical systems: we use it to learn control protocols that drive a system to behave in targeted ways. We now specialize to design loss functions for computational tasks. To begin, we establish how information is encoded within a physical system.

Consider a classical continuous n-dimensional system space $\mathcal{S}$. Encoding binary information requires coarse-graining $\mathcal{S}$—partitioning its continuous space into distinct regions that represent discrete-value states. Mathematically, a coarse-graining σ 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-dimension space spanned by binary values; i.e., $\mathrm{b} = \{|00000\rangle, |00001\rangle, \dots\}$. After coarse-graining, the degrees of freedom over this space are called *information-bearing degrees of freedom* (IBDs) [1, 23]. To support a stable information register for n binary digits, we reqire the original potential in $\mathcal{S}$ to have 2^n minima $\{x_{\mathrm{b}}\}$, each of which represents n binary digits.

The aim is to design protocols such that the system's final distribution $f(x_\tau)$ maps to the correct distribution $p_{\mathcal{I}_{\rm target}}$ 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:

$$\mathcal{L}_{\rm loss}(a) = \mathbb{E}\left[\alpha_1\mathcal{L}_{\rm error}(\omega,a) + \alpha_W W(\omega,a)\right] \ , \tag{5}$$

where $\mathcal{L}_{\rm error}$ represents an appropriate error measurement in $\mathcal{S}$ and α_1 and α_W are hyperparameters.

Suppose we wish to perform a computation task sending bit value b_i to $g(\mathrm{b}_i)$, where $g(\cdot)$ is the desired computational function. We ask that in the physical system S, all trajectories starting around the minimum x_{b_i} representing b_i to end around the minimum $x_{g(\mathrm{b}_i)}$ representing $g(\mathrm{b})$. To enforce this, we set:

$$\mathcal{L}_{\rm error} = (x_\tau - x_{g(\mathrm{b}_i)})^2 \ , \text{ if } x_0 \text{ near } x_{\mathrm{b}_i} \ . \tag{6}$$

We add higher-order moments $\mathbb{E}\left[(x_\tau - x_{g(\mathrm{b}_i)})^n\right]$ to the loss function if the L2 norm is insufficient. The error penalty used in Eq. (6) is not unique. For example, the sign function $\mathrm{sgn}(x)$ could also be used.

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 ϵ and then locate the most energetically-efficient protocol with error being no greater than ϵ. In optimization this is the *primal problem* [29]:

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

The corresponding Lagrangian *dual problem* is:

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

subject to $\lambda\geq 0$.

Suppose the optimal value of the primal problem—the optimal work cost with error being no greater than ϵ—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 ϵ we can find a λ such that these two optimizations are equivalent. If we do not have strong duality, then we simply select a λ 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 τ. The potential is $U(x,t)=(x-a_t)^2/2$, 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 $a_\tau=a_f$. We assume the system is in equilibrium at $t=0$ and $\mu=1$.

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

$$\mathcal{L}_{\text{move}}(a_t)=\alpha_1\mathbb{E}\left[(x_\tau-a_f)^2\right]+\alpha_W\mathbb{E}\left[W(\omega,a_t)\right]\ . \quad (7)$$

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 IV.) 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 parameterize the protocol values a_t by a continuous piecewise-linear function with 10 parameters that are the break points $\boldsymbol{a}=\{a_1,\ldots,a_{10}\}$. 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 $\boldsymbol{a}_t=\{a_t,b_t,c_t\}$:

$$U(x,t)=a_tx^4-b_tx^2+c_tx\ . \quad (8)$$

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 $x_{\text{eq}}=\pm b_t/2a_t$ and c_t tilts the potential. With these boundary conditions imposed, the potential U can encode a binary value as it has two stable equilibrium points $x=\pm 1$.

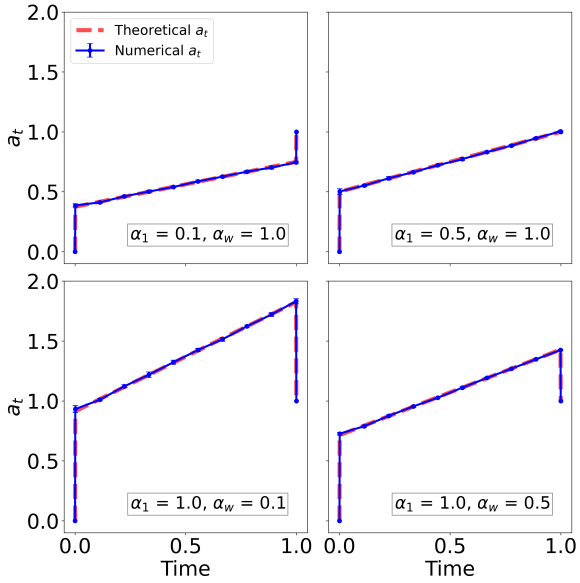

FIG. 1. Optimal protocol for translating a harmonic trap: We fix $a_f=1$ and $\tau=1$ and show different optimal protocols with different weight combinations after 600 iterations.

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:

$$\mathcal{L}_{\text{erase}}(\boldsymbol{a}_t)=\mathbb{E}\left[\alpha_1\left(x_\tau-1\right)^2+\alpha_WW(\omega,\boldsymbol{a}_t)\right]\ . \quad (9)$$

In numerical experiments, we set $\alpha_1=1$, work weight $\alpha_W=0.1$, $\mu=1$, $k_{\text{B}}T=1$, and $\tau=1$. Figures 2 and 3 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 $t=0^+$. The linear potential c_tx, used to drive the particle towards $x=+1$, dominates after $t=0.33$ and U is almost linear in the region of interests ($x\in[-1,1]$). After $t=0.33$, c_t decreases 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 $\tilde{W}_{\text{min}}=(2.942\pm0.026)k_{\text{B}}T$.

To compare our result with other attempts at bit erasure, we calculate the ratio $r=(\tilde{W}_{\text{min}}-\ln 2)/(\text{var}[x_0]/D\tau)=2.311\pm0.027$. 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 one cannot analytically calculate the optimal protocol for this

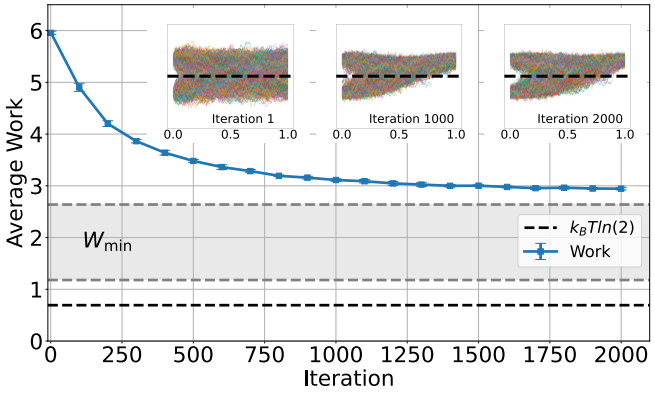

FIG. 2. Bit erasure optimization: For each training iteration, we use 5000 trajectories to compute the gradient and update parameters. We validate with 10^5 trajectories to estimate work cost. Three trajectory ensembles are selected from a random training. The plot also includes the upper and lower bounds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

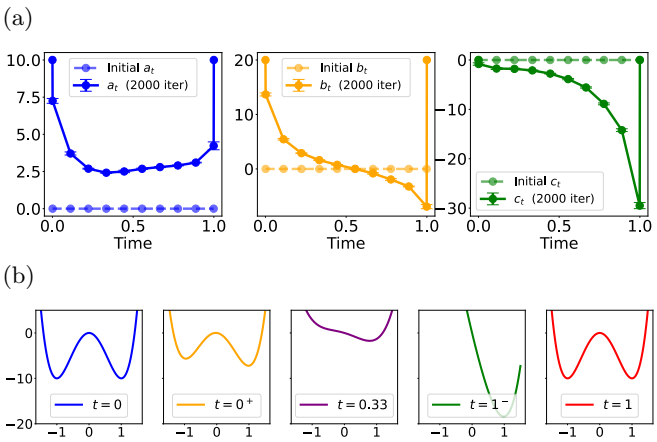

FIG. 3. Optimal erasure protocol: For each protocol parameter $\{a_t, b_t, c_t\}$, we use 10 break points for parameterization. (a) They are initialized to 0. Final protocol values after training for 2000 iterations. (b) Potentials U at different times.

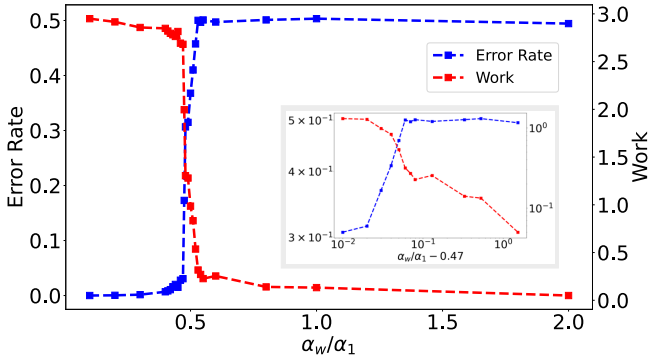

FIG. 4. Error rate and work dependence on the hyperparameter ratio in erasure: Around $\alpha_w/\alpha_1 = 0.47$, we observe a sharp rise in error rate. The inset is error rate and work versus $(\alpha_w/\alpha_1 - 0.47)$ in log scale, indicating a second order phase transition.

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 our choice of piecewise linear protocol, the form of the potential, and the closeness of our distribution at $t = \tau$ to the local equilibrium distribution, the proximity to the work cost for unrestricted optimal control demonstrates that our learning algorithm finds truly efficient protocols.

We briefly discuss the optimal protocol's dependence on the hyperparameter ratio α_W/α_1. Increasing the ratio α_w/α_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 $\mathbb{E}[\mathcal{L}_{\text{error}}] \sim 0^2 \cdot \frac{1}{2} + 2^2 \cdot \frac{1}{2} = 2$. Protocol 2 is a faithful erasure for which $\mathbb{E}[W] \sim \log 2 + 2\text{var}[x_0]/(D\tau) = 2.69$ [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. In Figure 4, 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.

Extension: Underdamped-dynamics optimization Let's turn to explore optimizing 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], such as the bit flip [36].

Consider one-dimensional underdamped dynamics:

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

where γ 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 dynamics in (x,v) space is still Markovian. The FRR derived above in Eq. (4) follows directly from the chain rule and so the considerations above work just as well in underdamped regimes. To illustrate, we successfully optimize a bit flip in one dimension. (See Fig. 5 and details in App. IV.)

IV. DISCUSSION AND CONCLUSION

Broadly speaking, we used the Feynman–Kac formula to transform an optimization using partial differential

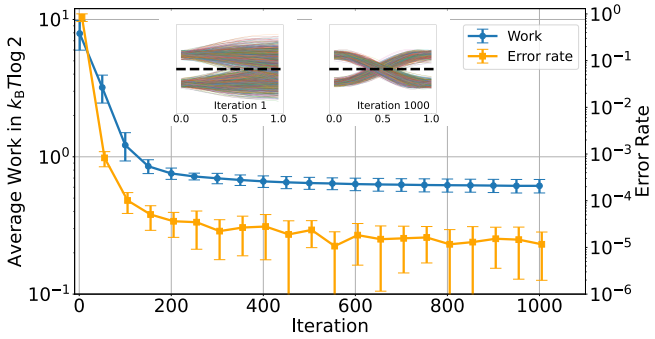

FIG. 5. Bit-flip mean work and error rate. The potential used is $U(x,t) = a_t x^4 - b_t x^2$. The work cost of our converged protocol is $\approx (0.615 \pm 0.070) k_B T$.

equations to an optimization over trajectories [37]. This reformulation offered a key advantage: simulating noisy trajectories is much easier than solving the corresponding equations for the optimal protocol, either numerically or analytically. This shift, from distributions to trajectories, has led to many exciting results in contemporary nonequilibrium physics [38–75]. 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 employed a linear combination of work cost and statistical moments to construct the loss function and FRRs to compute its gradient, allowing optimization of the protocol.

Generally, the choice of loss function should depend on the specific requirements of the task. For example, if the final distribution p_τ must be exactly p_{target}, then a distribution distance as error penalty is needed; one could use the KL-divergence $\mathbb{E}\left[\log\left(p_\tau/p_{\text{target}}\right)\right]$. However, computing the gradient of the log distribution $\partial_a \log p_\tau$ is typically challenging. Within the context of computational tasks, however, we observed that including a limited number of statistical moments (first and second-order only) in the loss function is effective for protocol design.

Of late, FRRs regained attention as studies found that they can bound the entropy production and dynamical activity in stochastic systems [25, 76–87]. 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 back-propagation and genetic algorithms [13, 20, 88–92]. Perhaps equally important, the FRR method (i) offers a more-physical interpretation, (ii) requires significantly less memory storage, and (iii) is relatively easy to implement in code; see [93]. Additionally, the generality of observables J for which the FRR can provide a gradient means that our proposed framework extends 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, as well. 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], biophyisical systems [102], quantum gates [103], and the like. Second, having proved-out linear piecewise 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 τ 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 and Jinghan Jia for helpful discussions, 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.

[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-Marin, T. Schmiedl, and U. Seifert, Optimal protocols for minimal work processes in underdamped stochastic thermodynamics, J. Chem. Phys. **129** (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 (2019).
[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 multi-bit logical operations, arXiv:2506.24021 (2025).

[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 coefficientsin Brownian dynamics simulations, Phys. Rev. Lett. **109**, 250601 (2012).

[25] J. Zheng and Z. Lu, Universal non-equilibrium response theory beyond steady states, arXiv:2403.10952 (2024).

[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*, Vol. 6 (Springer Science & Business Media, 2013).

[29] S. P. Boyd and L. Vandenberghe, *Convex optimization* (Cambridge university press, 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 **483**, 187 (2012).

[31] A. Bérut, A. Petrosyan, and S. Ciliberto, Detailed Jarzynski equality applied to a logically irreversible procedure, EPL **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.: Theory Exp. **2015** (6), 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] C. Dieball and A. Godec, Feynman-Kac theory of time-integrated functionals: Itô versus functional calculus, J. Phys. A **56**, 155002 (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, Thermodynamic bounds on generalized transport: From single-molecule to bulk observables, Phys. Rev. Lett. **133**, 067101 (2024).

[42] P. B. Melo, S. M. Duarte Queirós, and W. A. Morgado, Stochastic thermodynamics of Fisher information, Phys. Rev. E **111**, 014101 (2025).

[43] 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).

[44] C. Dieball and A. Godec, Perspective: Time irreversibility in systems observed at coarse resolution, J. Chem. Phys. **162** (2025).

[45] C. Dieball, Y. M. Satalsari, A. B. Zuccolotto-Bernez, S. U. Egelhaaf, M. A. Escobedo-Sánchez, and A. Godec, Precisely controlled colloids: a playground for path-wise non-equilibrium physics, Soft Matter **21**, 2170 (2025).

[46] C. Dieball and A. Godec, Thermodynamic correlation inequalities for finite times and transients, J. Phys. A (2025).

[47] G. Fiusa, P. E. Harunari, A. S. Hegde, and G. T. Landi, Counting observables in stochastic excursions, arXiv:2505.06208 (2025).

[48] 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 (2025).

[49] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, Phys. Rev.

Lett. **114**, 158101 (2015).
[50] 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).
[51] J. M. Horowitz and T. R. Gingrich, Proof of the finite-time thermodynamic uncertainty relation for steady-state currents, Phys. Rev. E **96**, 020103 (2017).
[52] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, Nat. Phys **16**, 15 (2020).
[53] K. Liu, Z. Gong, and M. Ueda, Thermodynamic uncertainty relation for arbitrary initial states, Phys. Rev. Lett. **125**, 140602 (2020).
[54] T. Koyuk and U. Seifert, Thermodynamic uncertainty relation for time-dependent driving, Phys. Rev. Lett. **125**, 260604 (2020).
[55] J. S. Lee, J.-M. Park, and H. Park, Universal form of thermodynamic uncertainty relation for Langevin dynamics, Phys. Rev. E **104**, L052102 (2021).
[56] Z. Cao, J. Su, H. Jiang, and Z. Hou, Effective entropy production and thermodynamic uncertainty relation of active Brownian particles, Phys. Fluids **34** (2022).
[57] P. P. Potts and P. Samuelsson, Thermodynamic uncertainty relations including measurement and feedback, Phys. Rev. E **100**, 052137 (2019).
[58] P. Pietzonka, A. C. Barato, and U. Seifert, Universal bounds on current fluctuations, Phys. Rev. E **93**, 052145 (2016).
[59] M. Polettini, A. Lazarescu, and M. Esposito, Tightening the uncertainty principle for stochastic currents, Phys. Rev. E **94**, 052104 (2016).
[60] K. Proesmans and C. Van den Broeck, Discrete-time thermodynamic uncertainty relation, EPL **119**, 20001 (2017).
[61] A. Dechant and S.-i. Sasa, Current fluctuations and transport efficiency for general langevin systems, J. Stat. Mech.: Theory Exp. **2018** (6), 063209.
[62] A. C. Barato, R. Chetrite, A. Faggionato, and D. Gabrielli, Bounds on current fluctuations in periodically driven systems, New J. Phys. **20**, 103023 (2018).
[63] K. Macieszczak, K. Brandner, and J. P. Garrahan, Unified thermodynamic uncertainty relations in linear response, Phys. Rev. Lett. **121**, 130601 (2018).
[64] K. Brandner, T. Hanazato, and K. Saito, Thermodynamic bounds on precision in ballistic multiterminal transport, Phys. Rev. Lett. **120**, 090601 (2018).
[65] T. Koyuk and U. Seifert, Operationally accessible bounds on fluctuations and entropy production in periodically driven systems, Phys. Rev. Lett. **122**, 230601 (2019).
[66] 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).
[67] T. Van Vu and Y. Hasegawa, Uncertainty relations for time-delayed Langevin systems, Phys. Rev. E **100**, 012134 (2019).
[68] A. Dechant, Multidimensional thermodynamic uncertainty relations, J. Phys. A **52**, 035001 (2018).
[69] Y. Hasegawa and T. Van Vu, Fluctuation theorem uncertainty relation, Phys. Rev. Lett. **123**, 110602 (2019).
[70] 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).
[71] K. Proesmans and J. M. Horowitz, Hysteretic thermodynamic uncertainty relation for systems with broken time-reversal symmetry, J. Stat. Mech.: Theory Exp. **2019** (5), 054005.
[72] T. Van Vu and Y. Hasegawa, Uncertainty relation under information measurement and feedback control, J. Phys. A **53**, 075001 (2020).
[73] A. M. Timpanaro, G. Guarnieri, J. Goold, and G. T. Landi, Thermodynamic uncertainty relations from exchange fluctuation theorems, Phys. Rev. Lett. **123**, 090604 (2019).
[74] Y. Hasegawa, Quantum thermodynamic uncertainty relation for continuous measurement, Phys. Rev. Lett. **125**, 050601 (2020).
[75] K. J. Ray, A. B. Boyd, G. Guarnieri, and J. P. Crutchfield, Thermodynamic uncertainty theorem, Phys. Rev. E **108**, 054126 (2023).
[76] 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).
[77] 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).
[78] T. Aslyamov and M. Esposito, Nonequilibrium response for Markov jump processes: Exact results and tight bounds, Phys. Rev. Lett. **132**, 037101 (2024).
[79] K. Ptaszyński, T. Aslyamov, and M. Esposito, Dissipation bounds precision of current response to kinetic perturbations, Phys. Rev. Lett. **133**, 227101 (2024).
[80] J. Zheng and Z. Lu, Universal response inequalities beyond steady states via trajectory information geometry, Phys. Rev. E **112**, L012103 (2025).
[81] Q. Gao, H.-M. Chun, and J. M. Horowitz, Thermodynamic constraints on kinetic perturbations of homogeneous driven diffusions, EPL **146**, 31001 (2024).
[82] K. Liu and J. Gu, Dynamical activity universally bounds precision of response in Markovian nonequilibrium systems, Commun. Phys. **8**, 62 (2025).
[83] T. Aslyamov, K. Ptaszyński, and M. Esposito, Nonequilibrium fluctuation-response relations: From identities to bounds, Phys. Rev. Lett. **134**, 157101 (2025).
[84] R. Bao and S. Liang, Nonequilibrium response theory: From precision limits to strong perturbation, arXiv:2412.19602 (2024).
[85] T. Van Vu, Fundamental bounds on precision and response for quantum trajectory observables, PRX Quantum **6**, 010343 (2025).
[86] 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).
[87] J. Zheng and Z. Lu, Nonlinear response relations and fluctuation-response inequalities for nonequilibrium stochastic systems, arXiv:2509.19606 (2025).
[88] S. Whitelam, How to train your demon to do fast information erasure without heat production, Phys. Rev. E **108**, 044138 (2023).
[89] C. Casert and S. Whitelam, Learning protocols for the fast and efficient control of active matter, Nat. Commun. **15**, 9128 (2024).
[90] S. Whitelam, Training thermodynamic computers by gradient descent, arXiv:2509.15324 (2025).

[91] S. Whitelam, C. Casert, M. Engel, and I. Tamblyn, Benchmarks for protocol control in nonequilibrium statistical mechanics, arXiv:2506.15122 (2025).

[92] N. Barros, S. Whitelam, S. Ciliberto, and L. Bellon, Learning efficient erasure protocols for an underdamped memory, Phys. Rev. E **111**, 044114 (2025).

[93] J. Lyu, K. J. Ray, and J. P. Crutchfield, Fluctuation response relation protocol optimization: Open source (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** (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** (2023).

[102] A. I. Brown and D. A. Sivak, Theory of nonequilibrium free energy transduction by molecular machines, Chem. Rev. **120**, 434 (2019).

[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, arXiv:2408.04576 (2024).

[106] K. J. Ray and J. P. Crutchfield, Gigahertz sub-Landauer momentum computing, Phys. Rev. Appl. **19**, 014049 (2023).