Designing Optimal Computation Protocols from Fluctuation Response Relations

Jinghao Lyu
Complexity Sciences Center and Physics and Astronomy Department
University of California at Davis
Davis, CA 95616
jolyu@ucdavis.edu

Abstract

The energy cost of computation has emerged as a central topic 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 computational systems driven far from equilibrium by time-dependent protocols. A key open question is then how to design protocols that minimize thermodynamic cost while still ensuring correct outcomes. In this work, we propose a unified framework to design optimal protocols using Fluctuation-Response Relation (FRR) and machine learning. Unlike previous approaches that either optimize distributions or protocols separately, our method unifies both using FRR-derived gradients. We construct loss functions for various computational objectives and then apply FRR, combined with gradient-based optimization, to efficiently locate the corresponding optimal protocols. We apply this framework to bit erasure in a double well potential—demonstrating how to construct loss functions that trade off energy cost against task error. Extending to underdamped systems, we show how momentum memory enables more complex operations like bit-flips, which are infeasible in purely one-dimensional Markovian settings. In both examples, we achieve work costs on the same order of Landauer's bound. Beyond these examples, our approach provides a general foundation for designing low-cost protocols in physical computation systems, with potential applications ranging from robust quantum gates under noise to energy-efficient control of chemical and synthetic biological networks.

1 Introduction

Nowadays, computation—from operating everyday digital devices to training large-scale machine learning models—consumes a significant amount of energy. This raises a question we long to answer: *how can we perform computation efficiently, or what is the minimal cost of computation?* The most well-known result in this context is Landauer's bound which states the minimum work to erase one bit of information at temperature T is $k_{\mathrm{B}}T\log 2$ where k_{B} is the Boltzmann constant [1]. Though Landauer's original derivation has limitations, it was groundbreaking in establishing a deep connection between computation and thermodynamics. To study the energetic cost in computation rigorously, we must turn to nonequilibrium physics, since all computation is performed in physical systems operating far away from equilibrium.

Recently, researchers have successfully assigned thermodynamic quantities to stochastic dynamics, leading to the emergence of the field known as *stochastic thermodynamics* [2]. With this foundation, the playground for exploring the thermodynamics of computation has been established. We consider a stochastic system $\mathcal{S}$ with a time-dependent energy landscape controlled by parameter a_t. Here,

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

the energy landscape have multiple local minima to stably store information. Suppose the initial distribution over system $\mathcal{S}$ is p_0. We aim to transform the distribution from p_0 to p_{target} over a time duration τ. We further ask the cyclic boundary condition in control parameter $a_0 = a_\tau$ to enable information storage and support future computation. The question is: how can we design the time-dependent protocol a_t to achieve distribution transform $p_0 \to p_{\text{target}}$ at minimal energetic cost.

Previously, two distinct energetic cost optimization approaches have been studied [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13]. The first is distribution-oriented: one seeks to transform the initial distribution p_0 into a target distribution p_{target}, without regard to the final energy landscape. In this approach, we do have the correct distribution but the potential is not correct. Then the information after the processing is unstable and can not be used in the future. The second is protocol-oriented: one aims to drive the system from one protocol configuration to another, without considering the resulting final distribution. This is not suitable for computational tasks. Imposing constraints on the final distribution and the final protocol values can over-constrain the system, leading to no feasible solution at all.

Unlike previous approaches that only optimize distributions or protocols, we propose a framework integrates both. We follow [14] to use 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 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 ensure both correct coarse-grained distribution and consistent potential landscape. We design loss functions for different computation tasks and use FRR to compute the gradient of the loss function. We show that our loss functions successfully drive the distribution toward the target coarse-grained distribution and achieve work on the order of Landauer's bound.

2 Fluctuation Response Relation

We first review fluctuation-response relations in overdamped dynamics. The question is that how expectation values of observables change with control parameters. Consider a one-dimension overdamped Langevin dynamics with a control parameter a:

$$dx = -\partial_x U(x, a)dt + \sigma dw \ , \tag{1}$$

where U is the potential of the system as a function of the position x and the control parameter a, dw is an infinitesimal Wiener process, and $\frac{1}{2}\sigma^2 = k_{\text{B}}T$. An observable J along a trajectory $\omega = \{x_t\}_{t=0}^{\tau}$ can be written as

$$J(\omega, a) = \int_\omega \left[f(x, a) \cdot dx + g(x, a)dt\right] \ , \tag{2}$$

where $\cdot$ is the Ito production. For example, the energetic cost (work) of one trajectory ω can be written in this form: $W = \int_\omega \partial_t U dt$. The expectation value of J with protocol a is then given by $\mathbb{E}_a[J] = \sum_\omega \Pr_a(\omega) J(\omega, a)$ where $\Pr_a(\omega)$ is the probability of observing the trajectory ω with protocol a. We want to study how the expectation of this current J changes with the change in control parameters $\frac{\partial}{\partial a}\mathbb{E}_a[J]$. This can be derived from taking derivative $\partial_a \log \Pr_a(\omega)$ directly. To see this, we write down the probability of observing a trajectory ω:

$$\Pr_a(\omega) = \mathcal{N} \exp\left\{-\frac{1}{2\sigma^2}\int_\omega [\dot{x} - \Phi(x, a)]^2 dt\right\} \ , \tag{3}$$

where $\mathcal{N}$ is a trajectory independent constant and $\Phi = -\partial_x U$ is the force. From the chain rule:

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

This can also be derived from Girsanov theorem. Fluctuation-response relation Eq. (4) allows us to access the gradient of observables with respect to control parameters in one simulation.

3 Loss function and optimization in computation tasks

In this section, we discuss the loss function design in computational tasks. To begin with, we need to establish how information is encoded within a physical system. Consider a classical continuous $n-$dimensional system $\mathcal{S} \in \mathbb{R}^n$. Encoding binary information requires coarse-graining the

system—partitioning its continuous state space into distinct regions that can represent discrete states. Mathematically, a coarse-graining is a many-to-one map from a distribution p over system $\mathcal{S}$ to a distribution $p_{\mathcal{I}}$ over a informational space $\mathcal{I}$. The informational space $\mathcal{I}$ is a finite dimension space spanned by bit values, i.e., $\mathrm{b} = \{|00000\rangle, |00001\rangle, \dots\}$. After coarse-grainig, the degree of freedom over this finite dimension space is called *information-bearing degrees of freedom*, or IBD [14]. To have a stable information register for $n-$binary digits, we ask the original potential to have 2^n minima $\{x_{\mathrm{b}}\}$ each of which represents an $n-$binary digit. A simple example is a double well potential $V(x) = \frac{1}{2}x^4 - x^2$ in which we have two minima $x^* = \pm 1$. This can be used to store 1 binary digit. We use $x = -1$ and $x = +1$ to represent $|0\rangle$ and $|1\rangle$, respectively. The physical system is $\mathcal{S} = \mathbb{R}$ and 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 $p_{|1\rangle} = \int_{0}^{+\infty} p(x)dx$.

We aim to design the dynamics such that the final distribution of the system corresponds to the correct distribution over IBD $p_{\mathcal{I}}$. We prioritize achieving the desired distribution over IBD rather than focusing solely on minimizing the work cost. This motivates a loss function:

$$\mathcal{L}_a = \mathbb{E}_a \left[\mathcal{L}_{\text{error}} + \alpha_W W \right] , \tag{5}$$

where $\mathcal{L}_{\text{error}}$ is a proper error penalty, W is the energetic cost and α_W is a hyperparameter.

Let us go back to the dynamics Eq.(1) and consider a computational task sending bit value b_i to $\sigma(\mathrm{b}_i)$. In the physical system S, it means that trajectories starting around the minimum x_{b_i} will end around the minimum x_{b_f} at $t = \tau$. The error loss we choose is $\mathcal{L}_{\text{error}} = (x_\tau - x_{\sigma(\mathrm{b}_i)})^2$. After we design the loss function $\mathcal{L}_a$, then we can use Eq. (4) to find the gradient of the loss function. Through this work, we will use Adam to optimize [15].

4 Examples

4.1 Bit erasure

The first computational task example we consider is bit erasure. Bit erasure is a fundamental computational task where the system starts with the distribution over IBD $p_{\mathcal{I}0} = (1/2, 1/2)$ and end with $p_{\mathcal{I}_{\text{target}}} = (0, 1)$. To implement this, we embed a particle in a time-dependent quartic potential $U(x,t) = a_t x^4 - b_t x^2 + c_t x$. The parameter boundary conditions are $a_0 = a_\tau = a,\ b_0 = b_\tau = b,\ c_0 = c_\tau = 0,\ 2a = b$. Parameters a_t, b_t control equilibrium locations 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 $x = \pm 1$. Initially, the system is at equilibrium with potential $U(x, 0)$ at temperature T. At time $t = \tau$, all particles should be around $x = +1$ so we choose the loss function to be

$$\mathcal{L}_{\text{erasure}} = \mathbb{E}_{a_t, b_t, c_t} \left[\alpha_1 (x_\tau - 1)^2 + \alpha_W W \right] . \tag{6}$$

In numerical experiments, we set $\alpha_1 = 1$, $\alpha_W = 0.1$, and $\tau = 1$. We use linear-piecewise functions to approximate these three time-dependent protocol values a_t, b_t, and c_t. We use 10 breaking points as parameters in each protocol values, resulting in a total of 30 parameters. Figs. 1 and 2 show our numerical results. In quasi-static limit, the efficient erasure is as follows: lowering the energy barrier first, tilting the potential, and raising the energy barrier in the end. In the finite time erasure, we see that lowering and tilting happen all at once at $t = 0^+$. Fig. 3 shows the trade-off between work and error rate for different values of the hyperparameter ratio α_w/α_1. The discontinuities in a_t, b_t, and c_t are ubiquitous in these protocol optimization and play a crucial role in enabling the distribution to evolve smoothly.

4.2 Underdamped Dynamics and Bit Flip

Next, we explore the optimization scheme in the underdamped regime. The position trajectories are longer Markovian in underdamped dynamics and this property can be exploited to accomplish certain computational tasks which are unachievable in purely Markovian (overdamped) dynamics. An illustrative computation task is bit-flip in one dimension system. In one dimension bit-flip, trajectories originating around -1 and those around $+1$ must converge near the middle before reaching their respective final logical values. If the dynamics are purely Markovian, we lose the ability to distinguish between these two trajectory classes once they meet, making the bit-flip operation

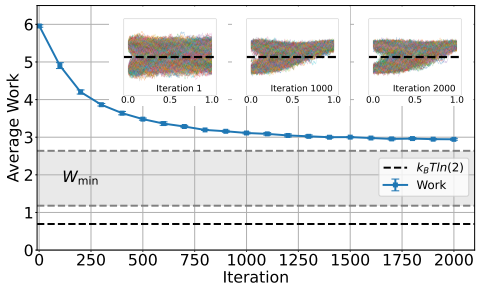

Figure 1: Bit erasure optimization: The ideal work cost of bit erasure is $1k_{\mathrm{B}}T\log 2$. The average work is in units of $k_{\mathrm{B}}T$. Insets show representative trajectory ensembles at iteration 1, 1000, and 2000. After 2000 iterations, almost all trajectory ends near $x = +1$.

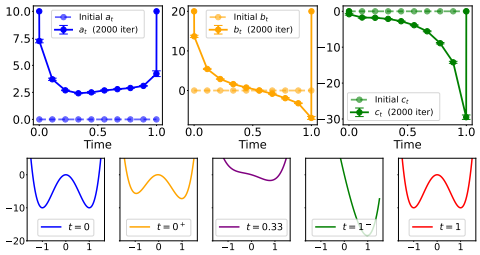

Figure 2: Optimal Erasure Protocol: For each protocol value $\{a_t, b_t, c_t\}$ in the double well potential, we use 10 breaking points for parameterization. All initial breaking points are set to be 0. We trained 2000 iterations. In the bottom 5 panels, we show the jumps of the potential at $t = 0^+$ and $t = 1^-$.

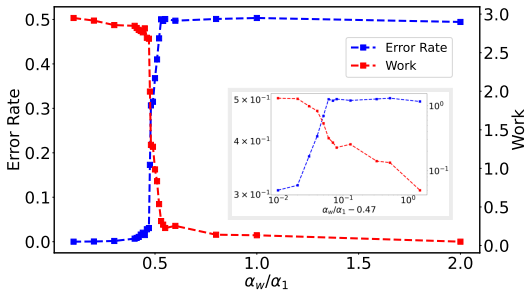

Figure 3: Error rate and work dependence on the hyperparameter ratio in erasure: Around $\alpha_w/\alpha_1 = 0.47$, we observe a sharp rise in error rate. The inset is error rate and work versus $(\alpha_w/\alpha_1 - 0.47)$ in log scale, indicating a second order phase transition.

infeasible. In contrast, one-dimension underdamped dynamics has additional degree of freedom through momentum, allowing us to distinguish and thus implement bit-flip operations effectively. Consider a one-dimension underdamped dynamics:

$$dx = vdt\ ,\ mdv = -\gamma v - \partial_x U(x,a)dt + \sigma dw\ , \tag{7}$$

where the white noise variance σ is related to the thermal bath temperature T by $\frac{1}{2}\sigma^2 = \gamma k_{\mathrm{B}}T$. In underdamped dynamics, FRR (Eq. 4) still holds. The aforementioned bit-flip is also called NOT operation. In bit flip task, we do not change the distribution in the system, the work cost in ideal bit flip is $0\ k_{\mathrm{B}}T$. Consider a double well potential $U(x,t)$ with the form of $U(x,t) = a_t x^4 - b_t x^2$. The boundary conditions are still $a_0 = a_\tau = a > 0$, $b_0 = b_\tau = b > 0$ and $2a = b$. The error loss function is $\mathcal{L}_{\text{error}} = (x_\tau - \text{target}(x_0))^2$ where $\text{target}(x_0) = +1$ if $x_0 < 0$ and $\text{target}(x_0) = -1$ if $x_0 > 0$. In Figs. 4 and 5, we show the flip results.

5 Discussion and applications

Our results demonstrate that FRR can serve as efficient gradient estimators for optimizing protocols in stochastic systems. By embedding FRR into a machine learning–based optimization loop, we can directly construct loss functions that balance energetic cost against computational accuracy, offering a principled framework for designing thermodynamically efficient computation. Our framework has wide applicability. In quantum systems, FRR-based optimization could guide the design of quantum gate protocols. In chemical and biochemical networks, our methods may be used to optimize reaction pathways or nanomachines.

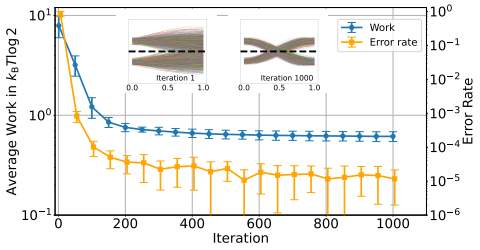

Figure 4: Bit flip optimization: The ideal work cost of bit erasure is $0 k_B T$. Insets show representative trajectory ensembles at iteration 1 (left) and iteration 1000 (right).

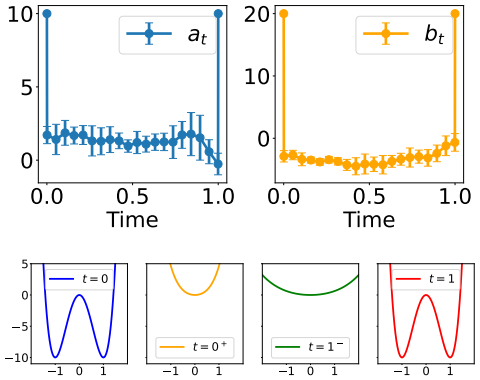

Figure 5: Optimal Erasure Protocol: For each protocol value $\{a_t,\ b_t\}$ in the double well potential, we use 20 breaking points for parameterization.

References

[1] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM journal of research and development*, 5(3):183–191, 1961.

[2] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on progress in physics*, 75(12):126001, 2012.

[3] Tim Schmiedl and Udo Seifert. Optimal finite-time processes in stochastic thermodynamics. *Phys. Rev. Lett.*, 98(10):108301, 2007.

[4] Alex Gomez-Marin, Tim Schmiedl, and Udo Seifert. Optimal protocols for minimal work processes in underdamped stochastic thermodynamics. *The Journal of chemical physics*, 129(2), 2008.

[5] Patrick R Zulkowski and Michael R DeWeese. Optimal finite-time erasure of a classical bit. *Physical Review E*, 89(5):052140, 2014.

[6] Andreas Dechant and Yohei Sakurai. Thermodynamic interpretation of wasserstein distance. *arXiv preprint arXiv:1912.08405*, 2019.

[7] Adrianne Zhong and Michael R DeWeese. Limited-control optimal protocols arbitrarily far from equilibrium. *Physical Review E*, 106(4):044135, 2022.

[8] Karel Proesmans, Jannik Ehrich, and John Bechhoefer. Finite-time landauer principle. *Physical Review Letters*, 125(10):100602, 2020.

[9] Salambô Dago and Ludovic Bellon. Logical and thermodynamical reversibility: Optimized experimental implementation of the not operation. *Physical Review E*, 108(2):L022101, 2023.

[10] Megan C Engel, Jamie A Smith, and Michael P Brenner. Optimal control of nonequilibrium systems through automatic differentiation. *Physical Review X*, 13(4):041032, 2023.

[11] Stephen Whitelam. Demon in the machine: learning to extract work and absorb entropy from fluctuating nanosystems. *Physical Review X*, 13(2):021005, 2023.

[12] Ralph Sabbagh, Olga Movilla Miangolarra, and Tryphon T Georgiou. Wasserstein speed limits for langevin systems. *Physical Review Research*, 6(3):033308, 2024.

[13] Jérémie Klinger and Grant M Rotskoff. Minimally dissipative multi-bit logical operations. *arXiv preprint arXiv:2506.24021*, 2025.

[14] Sebastian Deffner and Christopher Jarzynski. Information processing and the second law of thermodynamics: An inclusive, hamiltonian approach. *Physical Review X*, 3(4):041003, 2013.

[15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.