

UC Davis

UC Davis Previously Published Works

Title

Thermodynamic machine learning through maximum work production

Permalink

<https://escholarship.org/uc/item/4076403t>

Journal

New Journal of Physics, 24(8)

ISSN

1367-2630

Authors

Boyd, Alexander B
Crutchfield, James P
Gu, Mile

Publication Date

2022-08-01

DOI

10.1088/1367-2630/ac4309

Peer reviewed

PAPER • OPEN ACCESS

Thermodynamic machine learning through maximum work production

To cite this article: Alexander B Boyd *et al* 2022 *New J. Phys.* **24** 083040

View the [article online](#) for updates and enhancements.

You may also like

- [Quantum Brownian motion for magnets](#)
J Anders, C R J Sait and S A R Horsley
- [Distribution of first-reaction times with target regions on boundaries of shell-like domains](#)
Denis S Grebenkov, Ralf Metzler and Gleb Oshanin
- [Assessment of weak-coupling approximations on a driven two-level system under dissipation](#)
W S Teixeira, F L Semião, J Tuorila *et al.*



PAPER

Thermodynamic machine learning through maximum work production

OPEN ACCESS

RECEIVED
1 June 2021REVISED
27 November 2021ACCEPTED FOR PUBLICATION
14 December 2021PUBLISHED
7 September 2022

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the
title of the work, journal
citation and DOI.

Alexander B Boyd^{1,2,5,*} , James P Crutchfield^{3,*}  and Mile Gu^{1,2,4,*} ¹ Complexity Institute, Nanyang Technological University, Singapore² School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore³ Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, CA 95616, United States of America⁴ Centre for Quantum Technologies, National University of Singapore, Singapore⁵ The Division of Physics, Mathematics and Astronomy, Caltech, 1200 E California Blvd, Pasadena, CA 91125, United States of America

* Authors to whom any correspondence should be addressed.

E-mail: alecboy@gmail.com, chaos@ucdavis.edu and mgu@quantumcomplexity.org**Keywords:** nonequilibrium thermodynamics, Maxwell's demon, Landauer's principle, machine learning, maximum likelihood estimation

Abstract

Adaptive systems—such as a biological organism gaining survival advantage, an autonomous robot executing a functional task, or a motor protein transporting intracellular nutrients—must somehow embody relevant regularities and stochasticity in their environments to take full advantage of thermodynamic resources. Analogously, but in a purely computational realm, machine learning algorithms estimate models to capture predictable structure and identify irrelevant noise in training data. This happens through optimization of performance metrics, such as model likelihood. If such learning is physically implemented, is there a sense in which computational models estimated through machine learning are physically preferred? We introduce the thermodynamic principle that work production is the most relevant performance measure for an adaptive physical agent and compare the results to the maximum-likelihood principle that guides machine learning. Within the class of physical agents that most efficiently harvest energy from their environment, we demonstrate that an efficient agent's model explicitly determines its architecture and how much useful work it harvests from the environment. We then show that selecting the maximum-work agent for given environmental data corresponds to finding the maximum-likelihood model. This establishes an equivalence between nonequilibrium thermodynamics and dynamic learning. In this way, work maximization emerges as an organizing principle that underlies learning in adaptive thermodynamic systems.

1. Introduction

What is the relationship, if any, between abiotic physical processes and intelligence? Addressed to either living or artificial systems, this challenge has been taken up by scientists and philosophers repeatedly over the last centuries, from the 19th century teleologists [1] and biological structuralists [2, 3] to cybernetics of the mid-20th century [4, 5] and contemporary neuroscience-inspired debates of the emergence of artificial intelligence in digital simulations [6]. The challenge remains vital today [7–10]. A key thread in this colorful and turbulent history explores issues that lie decidedly at the crossroads of thermodynamics and communication theory—of physics and engineering. In particular, what bridges the dynamics of the physical world and its immutable laws and principles to the purposeful behavior of intelligent agents? The following argues that an essential connector lies in a new thermodynamic principle: work maximization drives learning.

Perhaps unintentionally, James Clerk Maxwell laid foundations for a physics of intelligence with what Lord Kelvin (William Thomson) referred to as 'intelligent demons' [11]. Maxwell in his 1857 book *Theory*

of Heat argued that a ‘very observant’ and ‘neat fingered being’ could subvert the second law of thermodynamics [12]. In effect, his ‘finite being’ uses its intelligence (Maxwell’s word) to sort fast from slow molecules, creating a temperature difference that drives a heat engine to do useful work. The demon presented an apparent paradox because directly converting disorganized thermal energy to organized work energy is forbidden by the second law. The cleverness in Maxwell’s paradox turned on equating the thermodynamic behavior of mechanical systems with the intelligence in an agent that can accurately measure and control its environment. This established an operational equivalence between energetic thermodynamic processes, on the one hand, and intelligence, on the other.

We will explore the intelligence of physical processes, substantially updating the setting from the time of Kelvin and Maxwell, by calling on a wealth of recent results on the nonequilibrium thermodynamics of information [13, 14]. In this, we directly equate the operation of physical agents descended from Maxwell’s demon with notions of intelligence found in modern machine learning. While learning is not necessarily the only capability of a presumed intelligent being, it is certainly a most useful and interesting feature.

The root of many tasks in machine learning lies in discovering structure from data. The analogous process of creating models of the world from incomplete information is essential to adaptive organisms, too, as they must model their environment to categorize stimuli, predict threats, leverage opportunities, and generally prosper in a complex world. Most prosaically, translating *training data* into a generative *model* corresponds to *density estimation* [15–17], where the algorithm uses the data to construct a probability distribution.

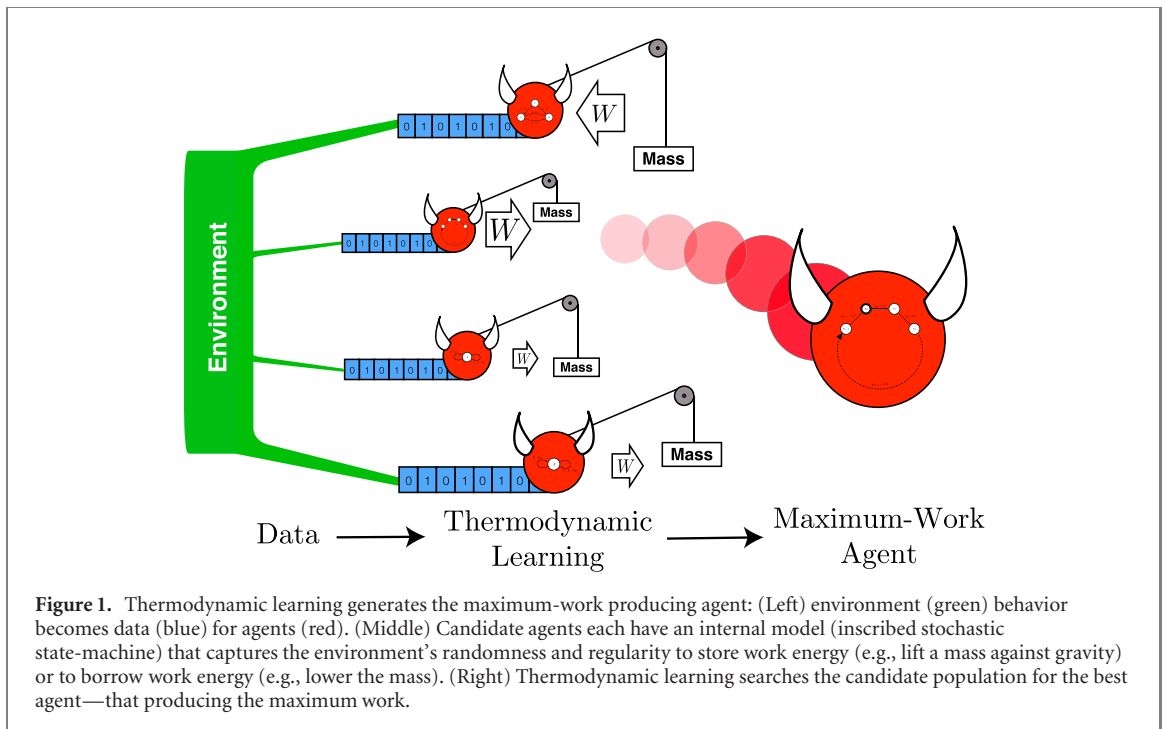
This type of model-building at first appears far afield from more familiar machine learning tasks such as categorizing pet pictures into cats and dogs or generating a novel image of a giraffe from a photo travelogue. Nonetheless, it encompasses them both [18]. Thus, by addressing thermodynamic roots of model estimation, we seek a physical foundation for a wide breadth of machine learning.

In estimating a distribution, it is common to invoke the principle of maximum-likelihood to guide intelligent learning. This is maximum likelihood estimation (MLE), whose central tenet is that, of the possible models consistent with the training data, an algorithm should select that model with maximum probability of having generated the data. Our exploration of the physics of learning asks whether a similar thermodynamic principle guides physical systems as they adapt to their environments.

The modern understanding of Maxwell’s demon no longer entertains violating the second law of thermodynamics [19]. In point of fact, the second law’s primacy has been repeatedly affirmed in modern nonequilibrium theory and experiment. That said, what has emerged is that we now understand how intelligent (demon-like) physical processes can harvest thermal energy as useful work. They do this by exploiting an information reservoir [19–21]—a storehouse of information as randomness and correlation. That reservoir is the demon’s informational environment, and the mechanism by which the demon measures and controls its environment embodies the demon’s intelligence, according to modern physics. We will show that this mechanism is directly linked to the demon’s model of its environment, which allows us to formalize the connection to machine learning.

Machine learning estimates likelihoods of different models given the same data. Analogously, in the physical setting of information thermodynamics, different demons harness different amounts of work using the same information reservoir. Leveraging this commonality, section 2 introduces *thermodynamic learning* as a physical process that infers optimal demons from environmental information. As shown in figure 1, thermodynamic learning selects demons that produce maximum work, paralleling density estimation through selection of models with maximum likelihood. Section 3 establishes background in MLE, computational mechanics, and thermodynamic computing necessary to formalize the comparison of maximum-work and maximum-likelihood learning. Our surprising result is that these two principles of maximization are the same, when compared in a common setting. This adds credence to the longstanding perspective that thermodynamics and statistical mechanics underlie many of the tools of machine learning [17, 22–28].

Section 4 formally establishes that to construct an intelligent work-harvesting demon, a probabilistic model of its environment is essential. That is, the demon’s Hamiltonian evolution is directly determined by its environmental model. This comes as a result of discarding demons that are ineffective at harnessing energy from *any* input, focusing only on a refined class of *efficient* demons that make the best use of the given data. This leads to the central result, found in section 5, that the demon’s work production from environmental ‘training data’ increases linearly with the log-likelihood of the demon’s model of its environment. Thus, if the thermodynamic training process selects the maximum-work demon for given data, it has also selected the maximum-likelihood model for that same data. Specifically, when harvesting work from a time series, the maximum-work agent is an information ratchet [29, 30] that specifies the transition and emission probabilities of a hidden Markov model (HMM).



Ultimately, our work demonstrates an *equivalence between the conditions of maximum work and maximum likelihood*. In this way, *thermodynamic learning is machine learning for thermodynamic machines*—it is a physical process that infers models in the same way a machine learning algorithm does. Thus, work itself can be interpreted as a *thermodynamic performance measure for learning*. In this framing, learning is physical, building on the long-lived narrative of the thermodynamics of organization, which we recount in section 6. While it is natural to argue that learning confers benefits, our result establishes that the benefit is fundamentally rooted in the physical link between energy and information dictated by the second law of thermodynamics.

2. Framework

While demons continue to haunt discussions of physical intelligence, the notion of a physical process trafficking in information and energy exchanges need not be limited to mysterious intelligent beings. Most prosaically, we are concerned with any physical system that, while interacting with an environment, simultaneously processes information at some energetic cost or benefit. Avoiding theological distractions, we refer to these processes as *thermodynamic agents*. In truth, any physical system can be thought of as an agent, but only a limited number of them are especially useful for or adept at commandeering information to convert between various kinds of thermodynamic resources, such as between heat and work. Here, we introduce a construction that shows how to find physical systems that are the most capable of processing information to affect thermodynamic transformations.

Consider an environment that produces information in the form of a time series of physical values at regular time intervals of length τ . We denote the particular state realized by the environment’s output at time $j\tau$ by the symbol $y_j \in \mathcal{Y}_j$. Just as the agent must be instantiated by a physical system, so too must the environment and its outputs to the agent. Specifically, \mathcal{Y}_j represents the state space of the j th output, which is a subsystem of the environment.

An agent has no access to the internals of its environment and so treats it as a black box. Thus, the agent can only access and interact with the environment’s output system \mathcal{Y}_j over each time interval $t \in (j\tau, (j+1)\tau)$. In other words, the state y_j realized by the environment’s output is also the agent’s *input* at time $j\tau$. For instance, the environment may produce realizations of a two level spin system $\mathcal{Y}_j = \{\uparrow, \downarrow\}$, which the agent is then tasked to manipulate through Hamiltonian control.

The aim, then, is to find an agent that produces as much work as possible using these black-box outputs. To do so, the agent must come to know something about the black box’s structure. This is the *principle of requisite complexity* [31]—thermodynamic advantage requires that the agent’s organization match that of its environment. We implement this by introducing a method for thermodynamic learning as shown in figure 1, that selects a specific agent from a collection of candidates.

Peeking into the internal mechanism of the black box, we wait for a time $L\tau$, receiving the L symbols $y_{0:L} = y_0 y_1 \dots y_{L-1}$. This is the agent's training data, which is copied as needed to allow a population of candidate agents to interact with it. As each agent interacts with a copy, it produces an amount of work, which it stores in the work reservoir for later use. In figure 1, the work reservoir is illustrated by a hanging mass which raises when positive work is produced, storing more energy in gravitational potential energy, and lowers when work production is negative, expending that same potential energy. However the work energy is stored, after the agents harvest work from the training data, the agent that produced the most work is selected.

Finding the maximum-work agent is 'thermodynamic learning' in the sense that it selects a device based on measuring its thermodynamic performance—the amount of work the device extracts. Ultimately, the goal is that the agent selected by thermodynamic learning continues to extract work as the environment produces new symbols. However, we leave analyzing the long-term effectiveness of thermodynamic learning to the future. Here, we concentrate on the condition of maximum-work itself, deriving and interpreting it.

Section 4 begins by describing the general class of physical agents that can harness work from symbol sequence, known as *information ratchets* [29, 30]. While these agents are sufficiently general to implement virtually any (Turing) computation, maximizing work production precludes a wide array of agents. Section 4.2 then refines our consideration to agents that waste as little work as possible and, in so doing, vastly narrow the search by thermodynamic learning. For this refined class of agents, we find that each agent's operation is exactly determined by its environment model. This leads to our final result, that the agent's work increases linearly with the model's log-likelihood.

For clarity, note that thermodynamic learning differs from physical systems that, evolving in time, dynamically adapt to their environment [26, 32, 33]. Work maximization as described here is thermodynamic in its objective, while these previous approaches to learning are thermodynamic in their mechanism.

That said, the perspectives are linked. In particular, it was suggested that physical systems spontaneously decrease work absorbed from driving [32]. Note that work absorbed by the system is opposite the work produced. And so, as they evolve over time, these physical systems appear to seek higher work production, paralleling how thermodynamic learning selects for the highest work production. And, the synchronization by which a physical system decreases work absorption is considered learning [32]. Reference [33] goes further, comparing the effectiveness of physical evolution to maximum-likelihood estimation employing an autoencoder. Notably, it reports that that form of machine learning performs markedly better than physical evolution, for the particular system considered there. By contrast, we show that the advantage of machine learning over thermodynamic learning does not hold in our framework. Simply speaking, they are synonymous.

We compare thermodynamic learning to machine learning algorithms that use maximum-likelihood to select models consistent with given data. As figure 1 indicates, each agent has an internal model of its environment; a connection section 4.5 formalizes. Each agent's work production is then evaluated for the training data. Thus, arriving at a maximum-work agent also selects that agent's internal model as a description of the environment. Moreover and in contrast with reference [33], which compares thermodynamic and machine learning methods numerically, the framework here leads to an analytic derivation of the equivalence between thermodynamic learning and MLE.

3. Preliminaries

Directly comparing thermodynamic learning and density estimation requires explicitly demonstrating that thermodynamically-embedded computing and machine learning share the framework just laid out. The following introduces what we need for this: concepts from machine learning, computational mechanics, and thermodynamic computing. (Readers preferring fuller detail should refer to appendix A.)

3.1. Maximum likelihood estimation

MLE determines, from training data, the parameters θ of a probability distribution. In the present setting, θ parameterizes a family of probabilities $\Pr(Y_{0:\infty} = y_{0:\infty} | \Theta = \theta)$ over sequences (or *words*) of any length. Here, $Y_{0:\infty} = Y_0 Y_1 \dots$ is the infinite-sequence random variable, composed of the random variables Y_j that each realize the environment's output y_j at time $j\tau$, and Θ is the random variable for the model. In other words, a given model θ predicts the probability of any sequence $y_{0:L}$ of any length L that one might see.




Figure 2. ϵ -Machine generating the phase-uncertain period-2 process: with probability 0.5, an initial transition is made from the start state s^* to state A. From there, it emits the sequence 1010 However, with probability 0.5, the start state transitions to state B and outputs the sequence 0101

For convenience, we introduce random variables Y_j^θ that define a model:

$$\Pr(Y_{0:\infty}^\theta) \equiv \Pr(Y_{0:\infty} | \Theta = \theta).$$

With training data $y_{0:L}$, the *likelihood* of model θ is given by the probability of the data given the model:

$$\begin{aligned} \mathcal{L}(\theta|y_{0:L}) &= \Pr(Y_{0:L} = y_{0:L} | \Theta = \theta) \\ &= \Pr(Y_{0:L}^\theta = y_{0:L}). \end{aligned}$$

MLE seeks to optimize the likelihood $\mathcal{L}(\theta|y_{0:L})$ [15, 17, 34]. However, the procedure that finds maximum-likelihood estimates usually employs the *log-likelihood* instead:

$$\ell(\theta|y_{0:L}) = \ln \Pr(Y_{0:L}^\theta = y_{0:L}), \quad (1)$$

since it is maximized for the same models, but converges more effectively [35].

3.2. Computational mechanics

Given that our data is a time series of arbitrary length starting with y_0 , we must choose a model class whose possible parameters $\Theta = \{\theta\}$ specify a wide range of possible distributions $\Pr(Y_{0:\infty}^\theta)$ —the *semi-infinite processes*. ϵ -machines, a class of finite-state machines introduced to describe bi-infinite processes $\Pr(Y_{-\infty:\infty}^\theta)$, provide a systematic means to do this [36]. As described in appendix A these finite-state machines comprise just such a flexible class of representations; they can describe any semi-infinite process. This follows from the fact that they are the minimal sufficient statistic for prediction explicitly constructed from the process.

A process's ϵ -machine consists of a set of hidden states \mathcal{S} , a set of output states \mathcal{Y} , a start state $s^* \in \mathcal{S}$, and conditional output-labeled transition matrix $\theta_{s \rightarrow s'}^{(y)}$ over the hidden states:

$$\theta_{s \rightarrow s'}^{(y)} = \Pr(\mathcal{S}_{j+1}^\theta = s', Y_j^\theta = y | \mathcal{S}_j^\theta = s).$$

$\theta_{s \rightarrow s'}^{(y)}$ specifies the probability of transitioning to hidden state s' and emitting symbol y given that the machine is in state s . In other words, the model is fully specified by the tuple:

$$\theta = \{\mathcal{S}, \mathcal{Y}, s^*, \{\theta_{s \rightarrow s'}^{(y)}\}_{s, s' \in \mathcal{S}, y \in \mathcal{Y}}\}.$$

As an example, figure 2 shows an ϵ -machine that generates a periodic process with initially uncertain phase.

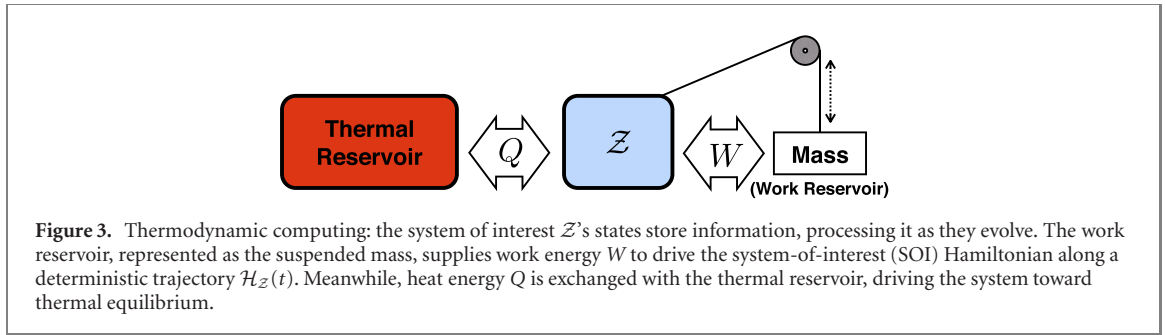
ϵ -Machines are *unifilar*, meaning that the current causal state s_j along with the next k symbols uniquely determines the following causal state through the *propagator function*:

$$s_{j+k} = \epsilon(s_j, y_{j:j+k}).$$

This yields a simple expression for the probability of any word in terms of the model parameters:

$$\Pr(Y_{0:L}^\theta = y_{0:L}) = \prod_{j=0}^{L-1} \theta_{\epsilon(s^*, y_{0:j}) \rightarrow \epsilon(s^*, y_{0:j+1})}^{(y_j)}.$$

In addition to being uniquely determined by the semi-infinite process, the ϵ -machine uniquely generates that same process. This means that our model class Θ is equivalent to the class of possible distributions over time series data. Moreover, knowledge of the causal state of an ϵ -machine at any time step j contains all information about the future that could be predicted from the past. In this sense, the causal state is



predictive of the process. These and other properties have motivated a long investigation of ϵ -machines, in which the memory cost of storing the causal states is frequently used as a measure of process structure. Appendix A gives an extended review.

3.3. Thermodynamic computing

Computation is physical—any computation takes place embedded in a physical system. Here, we refer to the substrate of the physically-embedded computation as the *system of interest* (SOI). Its states, denoted $\mathcal{Z} = \{z\}$, are taken as the underlying physical system's *information bearing degrees of freedom* [19]. The SOI's dynamic evolves the state distribution $\Pr(Z_t = z_t)$, where Z_t is the random variable describing state at time t . Computation over time interval $t \in [\tau, \tau']$ specifies how the dynamic maps the SOI from the initial time $t = \tau$ to the final time $t = \tau'$. It consists of two components:

- An initial distribution over states $\Pr(Z_\tau = z_\tau)$ at time $t = \tau$.
- Application of a Markov channel M , characterized by the conditional probability of transitioning to the final state $z_{\tau'}$ given the initial state z_τ :

$$M_{z_\tau \rightarrow z_{\tau'}} = \Pr(Z_{\tau'} = z_{\tau'} | Z_\tau = z_\tau).$$

Together, these specify the SOI's computational elements. In this, z_τ is the *input* to the physical computation, $z_{\tau'}$ is the *output*, and $M_{z_\tau \rightarrow z_{\tau'}}$ is the *logical architecture*.

Figure 3 illustrates a computation's physical implementation. SOI \mathcal{Z} is coupled to a work reservoir, depicted as a mass hanging from a string, that controls the system's Hamiltonian along a trajectory $\mathcal{H}_{\mathcal{Z}}(t)$ over the computation interval $t \in [\tau, \tau']$ [37]. This is the basic definition of a *thermodynamic agent*: an evolving Hamiltonian driving a physical system to compute at the cost of work.

In a classical system, this control determines each state's energy $E(z, t)$. As a result of the control, changes in energy due to changes in the Hamiltonian correspond to work exchanges between the SOI and work reservoir. The system \mathcal{Z} follows a *state trajectory* $z_{\tau:\tau'}$ over the time interval $t \in [\tau, \tau']$, which we can write:

$$z_{\tau:\tau'} = z_\tau z_{\tau+dt} \dots z_{\tau'-dt} z_{\tau'},$$

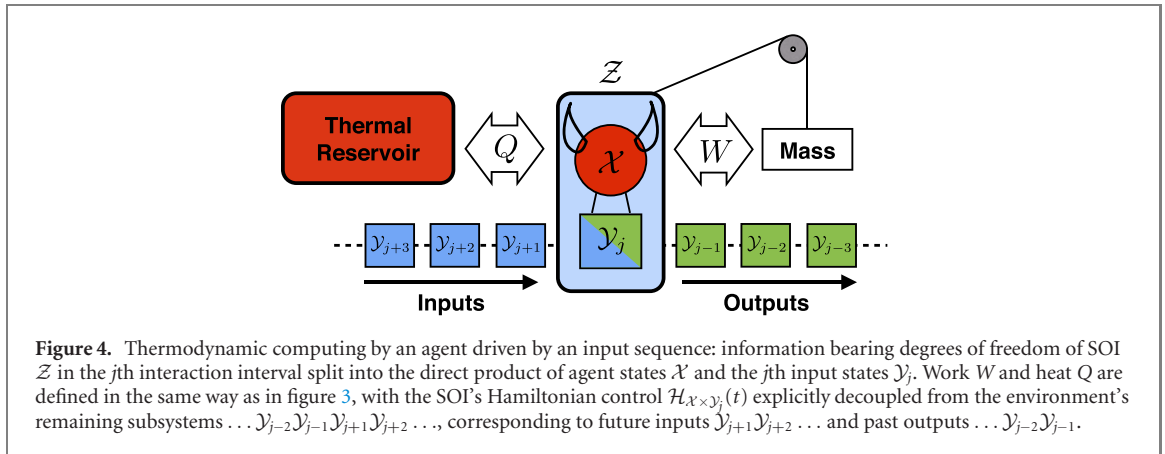
where z_t is the system state at time t . Here, we decomposed the trajectory into intervals of duration dt , taken short enough to yield infinitesimal changes in state probabilities and the Hamiltonian. The resulting *work production* for this trajectory is then the integrated change in energy due to the Hamiltonian's time dependence [37]:

$$W|_{z_{\tau:\tau'}} = - \int_{\tau}^{\tau'} dt \partial_t E(z, t)|_{z=z_t}. \quad (2)$$

This is the work energy that can be harvested in the ideal case: when the work reservoir has sufficiently large inertia it drives the SOI deterministically.

Note that while the state trajectory $z_{\tau:\tau'}$ mirrors the time series notation used for the training data $y_{0:L} = y_0 y_1 \dots y_{L-1}$, they are different objects and should not be conflated. On the one hand, the training data series $y_{0:L}$ is composed of realizations of L *separate* subsystems, each produced at different times $j\tau$, $j \in \{0, 1, 2, \dots, L-1\}$. y_j is realized in the subsystem \mathcal{Y}_j , and so it can be manipulated completely separately from any other element of $y_{0:L}$ lying outside of \mathcal{Y}_j . By contrast, z_t depends dynamically on many other elements in $z_{\tau:\tau'}$, all of which lie in the *same* system, since the time series $z_{\tau:\tau'}$ represents state evolution of the single system \mathcal{Z} over time.

While the SOI exchanges work energy with the work reservoir, as figure 3 shows, it exchanges heat Q with the thermal reservoir. Coupling to a heat reservoir adds stochasticity to the state trajectory $z_{\tau:\tau'}$. Since the SOI computes while coupled to a thermal reservoir at temperature T , *Landauer's principle* [19] relates a



computation's logical processing to its energetics. In its contemporary form, it bounds the *average* work production $\langle W \rangle$ by a term proportional to SOI's entropy change. Taking the Shannon entropy $H[Z_t] = -\sum_z \Pr(Z_t = z) \ln \Pr(Z_t = z)$ in natural units, the second law of thermodynamics implies [14]:

$$\langle W \rangle \leq k_B T (H[Z_{\tau'}] - H[Z_{\tau}]).$$

Here, the average $\langle W \rangle$ is taken over all possible microscopic trajectories. And, the energy landscape is assumed to be flat at the computation's start and end, giving no energetic preference to a particular informational state.

4. Agent energetics

We now construct the theoretical framework for how agents extract work from time-series data. This involves breaking down the agent's actions into manageable elementary components—where we demonstrate their actions can be described as repeated application of a sequence of computations. We then introduce tools to analyze work production within such general computations on finite data. We highlight the importance of the agent's model of the data in determining work production. This model-dependence emerges by refining the class of agents to those that execute their computation most efficiently. The results are finally combined, resulting in a closed-form expression for agent work production from time-series data.

4.1. Agent architecture

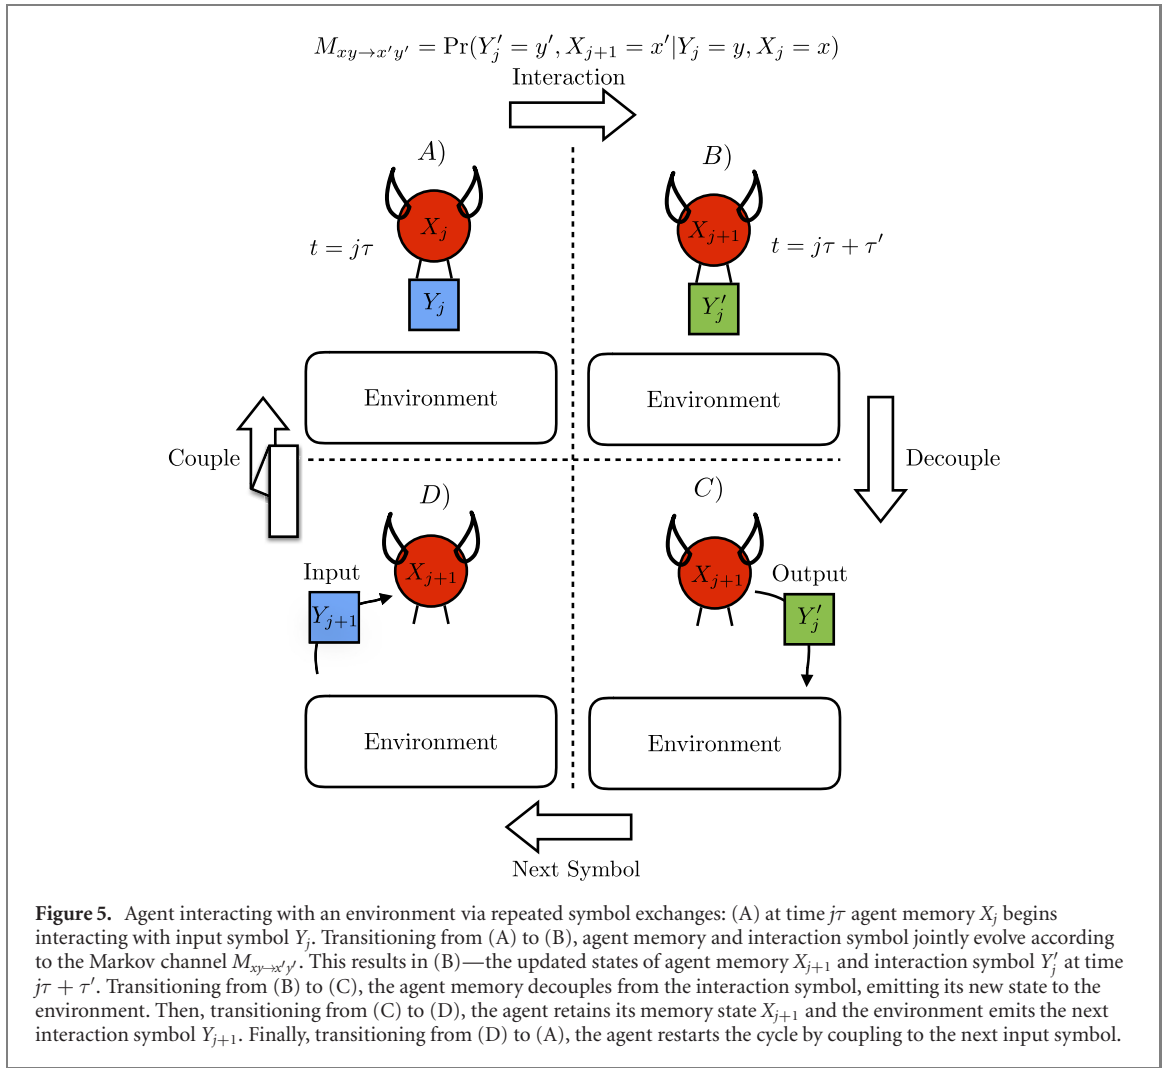
Recall from section 2 that the basic framework describes a thermodynamic agent interacting with an environment at regular time-intervals τj in state y_j . Each y_j is drawn according to a random variable Y_j , such that the sequence $Y_{0:\infty} = Y_0 Y_1 \dots$ is a semi-infinite stochastic process. The agent's task is to interact with this input string to generate useful work.

For example, consider an agent charged with extracting work from an alternating process—a sequence emitted by a degenerate two-level system that alternates periodically between symbols 0 and 1. In isolation each symbol looks random and has no free energy. Thus, an agent that interacts with each symbol the same way gains no work. However, a memoryful agent can adaptively adjust its behavior, after reading the first symbol, to exactly predict succeeding symbols and, therefore, extract meaningful work. This method of harnessing *temporal correlations* is implemented by *information ratchets* [29, 30]. They combine physical inputs with additional agent memory states that store the input's temporal correlations.

As shown in figure 4, we describe an agent's memory via an ancillary physical system \mathcal{X} . The agent then operates cyclically with duration τ , such that the j th cycle runs over the time-interval $[j\tau, (j+1)\tau)$. Each cycle involves two phases:

- Interaction:* agent memory \mathcal{X} couples to and interacts with the j th input system \mathcal{Y}_j that contains the j th input symbol y_j . This phase has duration $\tau' < \tau$, meaning the j th interaction phase occurs over the time-interval $[j\tau, j\tau + \tau')$. At the end, the agent decouples from the system \mathcal{Y}_j , passing its new state y'_j to the environment as output or exhaust.
- Rest:* during time interval $[j\tau + \tau', (j+1)\tau)$, the agent's memory \mathcal{X} sits idle, waiting for the next input \mathcal{Y}_{j+1} .

In this way, the agent transforms a series of inputs $y_{0:L}$ into a series of outputs $y'_{0:L}$.



In each cycle, all nontrivial thermodynamics occur in the interaction phase, during which the SOI consists of the joint agent-input system: i.e., $\mathcal{Z} = \mathcal{X} \otimes \mathcal{Y}_j$, as shown in figure 4. While the other subsystems $\dots \mathcal{Y}_{j-2}\mathcal{Y}_{j-1}$ and $\mathcal{Y}_{j+1}\mathcal{Y}_{j+2} \dots$ may be physically instantiated somewhere else in the environment, they do not participate in the interaction, since they are energetically decoupled from the agent in this phase. Paralleling the computation shown in figure 3, Hamiltonian control over the joint space $\mathcal{H}_{\mathcal{X} \times \mathcal{Y}_j}(t)$ results in a transformation of the agent's SOI that also requires the exchange of work and heat.

This interaction phase updates SOI states according to a Markov transition matrix M shown in figure 5:

$$M_{xy \rightarrow x'y'} = \Pr(X_{j+1} = x', Y'_j = y' | X_j = x, Y_j = y), \quad (3)$$

where X_j and X_{j+1} are the random variables for the states of the agent's memory \mathcal{X} before and after the j th interaction interval, and Y_j and Y'_j are the random variables for the system \mathcal{Y}_j before and after the same interaction interval, realizing the input and output, respectively.

As section 3.3 described, M is the logical architecture of the physical computation that transforms the agent's memory and input simultaneously. It is the central element in the agent's procedure for transforming inputs $y_{0:L}$ into associated outputs $y'_{0:L}$. The key observation is that M captures all of the agent's internal logic. The logic does not change from cycle to cycle. However, the presence of persistent internal memory between cycles implies that the agent's behavior adapts to past inputs and outputs. This motivates us to define M as the *agent architecture* since it determines how an agent stores information temporally. As we will show, M is one of two essential elements in determining the work an agent produces from a time series.

Note that prior related efforts to address agent energetics focused on ensemble-average work production [29–31, 38–42]. In contrast, here we relate work production to MLE—which involves each agent being given a specific data string $y_{0:L}$ for training. To address this shift in problem, the following determines the work production for *single-shot short-input* strings.

4.2. Energetics of computational maps

The agent architecture M specifies a physical computation as described in section 3.3 and therefore has a minimum energy cost determined by Landauer's bound. However, this is a bound on the *average* work production, which depends explicitly on the distribution of inputs. We need to determine, instead, the work produced from a *single* input y_j . To find this we return to the general case of SOI \mathcal{Z} undergoing a thermodynamic computation M .

A physical operation takes the SOI from state z_τ at time τ to state $z_{\tau'}$ at time τ' . This specifies a *computational map* $z_\tau \rightarrow z_{\tau'}$ that ignores intermediate states in the SOI state trajectory, as all information relevant to the computation's logical operation lies in the input and output. Thus, our attention turns to the question: what is the work production of a computational map $z_\tau \rightarrow z_{\tau'}$ performed by the computation M at temperature T ?

To answer this question we investigate the average work produced when conditioned on the initial and final states of the map:

$$\langle W_{|z_\tau, z_{\tau'}} \rangle = \sum_{z'_{\tau, \tau'}} W_{|z'_{\tau, \tau'}} \Pr(Z_{\tau, \tau'} = z'_{\tau, \tau'} | z_\tau, z_{\tau'}). \quad (4)$$

This depends sensitively on how the computation is implemented. However, there is an important subclass of computations for which there is a precise relationship between energy harvested and underlying: thermodynamically efficient computations.

Theorem 1. *An efficient computation that maps between equal-energy memory states with an estimated model θ of the SOI distribution produces work from the computational map $z_\tau \rightarrow z_{\tau'}$:*

$$\langle W_{|z_\tau, z_{\tau'}}^\theta \rangle = k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{\Pr(Z_{\tau'}^\theta = z_{\tau'})}. \quad (5)$$

Proof. See appendix B. □

This is the average energy harvested from a particular input–output pair. W^θ is used to denote that this model θ was used to design the energy landscape. This work depends on the *estimated* distributions $\Pr(Z_i^\theta)$ encoded in the evolving energy landscape $E(z, t)$. However, it is independent of the *actual* distributions since work production of a computational map is conditioned on the input z_τ and output $z_{\tau'}$.

Without going into the details of the proof presented in appendix B, the role that the model plays in the design of energetics $E(z, t)$ is opaque, but it is integral in relating thermodynamic learning to machine learning. To illustrate how this relationship emerges, consider a two-state SOI $\mathcal{Z} = \{\uparrow, \downarrow\}$. We describe a computation that harvests maximum energy from an initial distribution with 80% spin-up $\{\Pr(Z_\tau^\theta = \uparrow), \Pr(Z_\tau^\theta = \downarrow)\} = \{0.8, .02\}$ and maps it to a final distribution $\{\Pr(Z_{\tau'}^\theta = \uparrow), \Pr(Z_{\tau'}^\theta = \downarrow)\} = \{0.4, 0.6\}$.

According to Landauer's bound [14], when starting and ending in a flat energy landscape the maximum harvestable energy is the difference in Shannon entropies: $\langle W \rangle_{\max} = k_B T (H[Z_{\tau'}^\theta] - H[Z_\tau^\theta])$. Achieving this bound means that the computation is efficient, not wasting any thermodynamic resources.

One route to achieving Landauer's bound is to execute a three step process, as shown in figure 6:

- (a) Instantaneously quench the energy landscape from equal energies:

$$E(\uparrow, \tau) = E(\downarrow, \tau) = 0,$$

to equilibrium with the estimated input distribution, with:

$$\{E(\uparrow, \tau^+), E(\downarrow, \tau^+)\} = \{-k_B T \ln 0.8, -k_B T \ln 0.2\}.$$

- (b) Quasistatically evolve the energy landscape from equilibrium with the initial distribution:

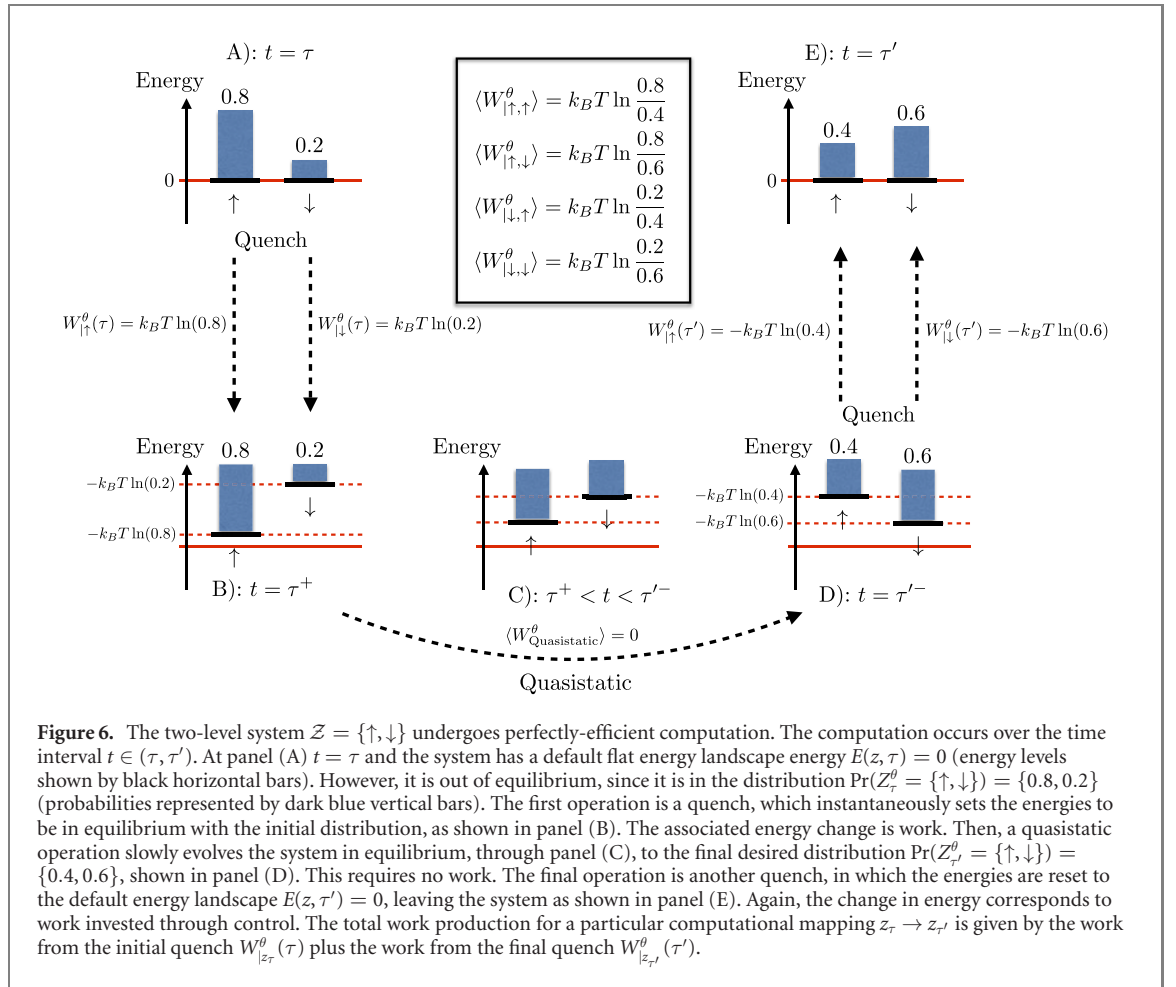
$$\{E(\uparrow, \tau^+), E(\downarrow, \tau^+)\} = \{-k_B T \ln 0.8, -k_B T \ln 0.2\},$$

to equilibrium with the final target distribution:

$$\{E(\uparrow, \tau'^-), E(\downarrow, \tau'^-)\} = \{-k_B T \ln 0.4, -k_B T \ln 0.6\}.$$

- (c) Instantaneously quench the energy landscape back from equilibrium:

$$\{E(\uparrow, \tau'^-), E(\downarrow, \tau'^-)\} = \{-k_B T \ln 0.4, -k_B T \ln 0.6\},$$



to the original equal energy landscape:

$$E(\uparrow, \tau') = E(\downarrow, \tau') = 0.$$

In executing this cycle in the energy landscape, we have explicitly encoded the model θ and the associated input distribution Z_τ^θ and output distribution $Z_{\tau'}^\theta$ through the Hamiltonian control:

$$E(z, \tau^+) = -k_B T \ln \Pr(Z_\tau^\theta = z)$$

$$E(z, \tau'^-) = -k_B T \ln \Pr(Z_{\tau'}^\theta = z).$$

The protocol shown in figure 6 allows us to calculate the work produced W^θ both on average and for particular computational maps. Using the work definition in equation (2), computational-map work production breaks into the three protocol substages:

(a) The initial quench contributes:

$$\begin{aligned} W_{|z_\tau}^\theta(\tau) &= E(z_\tau, \tau^+) - E(z_\tau, \tau) \\ &= k_B T \ln \Pr(Z_\tau^\theta = z_\tau). \end{aligned}$$

(b) The quasistatic evolution contributes minus the change in equilibrium free energy, which is zero in this case:

$$\langle W_{\text{Quasistatic}}^\theta \rangle = 0.$$

(See appendix E for proof.)

(c) The final quench contributes:

$$\begin{aligned} W_{|z_{\tau'}}^\theta(\tau') &= E(z_{\tau'}, \tau') - E(z_{\tau'}, \tau'^-) \\ &= -k_B T \ln \Pr(Z_{\tau'}^\theta = z_{\tau'}). \end{aligned}$$

The work production definition dictates that instantaneous changes in energy levels come from work, providing the contributions in steps (a) and (c). The net result for a computational map $z_\tau \rightarrow z_{\tau'}$ is the sum of the three terms: $W_{z_\tau}^\theta(\tau) + \langle W_{\text{Quasistatic}}^\theta \rangle + W_{|z_{\tau'}}^\theta(\tau') = k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{\Pr(Z_{\tau'}^\theta = z_{\tau'})}$. This exactly reproduces equation (5) of theorem 1. The simple computational protocol shown here demonstrates how the model θ , used to design an efficient computation, can determine energy production.

While we reproduced equation (5) for a particular one-bit computation, appendix E shows how to construct *any* stochastic input–output map such that it achieves the maximum work predicted by Landauer’s bound on average and by equation (5) for particular computational maps. Again, the strategy is to implement a three step sequence: (1) initial quench, (2) slow quasistatic operations, and (3) final quench. This provides an intuitive mechanistic basis for theorem 1 in addition to the explicit fluctuation theorem proof provided in appendix B.

Thermodynamic learning concerns agents that maximize work production from their input data. *Maximum-work agents* zero out the average entropy production $\langle \Sigma \rangle = -\langle W \rangle - \Delta F^{\text{neq}}$ so that the work production of a computational map satisfies equation (5). As a result, we see that pursuing maximum work production results in this strict relationship between the energy produced from a computational map and the model θ .

From here on out, when we refer to *efficient* agents we refer to those that maximize work production from the available change in nonequilibrium free energy. Equation (5) links an efficient agent’s thermodynamic performance to model performance, which is an essential step in determining a model through thermodynamic learning.

4.3. Thermally efficient agents

With the work production of a maximally-efficient computational map established, we are poised to determine the work production for thermodynamically-efficient agents. Specifically, consider an agent parameterized by its logical architecture M and model parameters θ . As described by the agent architecture in section 4.1, the agent uses its memory X_j to map inputs Y_j to outputs Y'_j and to update its memory to X_{j+1} . In stochastically mapping the SOI $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, the model parameter θ provides an estimate of the distribution over the current initial state (X_j^θ, Y_j^θ) as well as the final state $(X_{j+1}^\theta, Y_{j+1}^\theta)$. Assuming the agent’s logical architecture M is executed optimally, direct application of equation (5) then says that the expected work production of the computational map $x_j y_j \rightarrow x_{j+1} y'_j$ is:

$$\langle W_{j;x_j y_j \rightarrow x_{j+1} y'_j}^\theta \rangle \equiv k_B T \ln \frac{\Pr(X_j^\theta = x_j, Y_j^\theta = y_j)}{\Pr(X_{j+1}^\theta = x_{j+1}, Y_{j+1}^\theta = y'_j)}. \quad (6)$$

In this, the estimated final distribution comes from the logical architecture updating the initial distribution:

$$\Pr(X_{j+1}^\theta = x', Y_{j+1}^\theta = y') = \sum_{x,y} \Pr(X_j^\theta = x, Y_j^\theta = y) M_{xy \rightarrow x' y'}.$$

Equation (6)’s expression for work establishes that all functional aspects (logical operation and energetics) of an efficient agent are determined by two factors:

- The logical architecture M that specifies how the agent manipulates inputs and updates its own memory.
- The estimated input distribution $\Pr(X_j^\theta, Y_j^\theta)$ for which the agents’ execution of M is optimized to minimize dissipation.

Thus, we define a *thermally-efficient agent* by the ordered pair $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$.

So defined, we can calculate the work produced when such agents act on a particular input sequence

$y_{0:L}$.

Theorem 2. *The average work produced by an efficient agent with logical architecture M and estimated input distributions $\Pr(X_j^\theta, Y_j^\theta)$ when it transduces the input sequence $y_{0:L}$ is:*

$$\langle W_{|y_{0:L}}^\theta \rangle = k_B T \sum_{x_{0:L+1}, y'_{0:L}} \Pr(X_0 = x_0) \prod_{k=0}^{L-1} M_{x_k, y_k \rightarrow x_{k+1}, y'_k} \ln \prod_{j=0}^{L-1} \frac{\Pr(X_j^\theta = x_j, Y_j^\theta = y_j)}{\Pr(X_{j+1}^\theta = x_{j+1}, Y_{j+1}^\theta = y'_j)}. \quad (7)$$

Proof. First consider the work production of a particular sequence of agent memory states $x_{0:L+1}$ and outputs $y'_{0:L}$:

$$\langle W_{|y_{0:L}, y'_{0:L}, x_{0:L+1}}^\theta \rangle = \sum_{j=0}^{L-1} \langle W_{j;x_j y_j \rightarrow x_{j+1} y'_j}^\theta \rangle = k_B T \ln \prod_{j=0}^{L-1} \frac{\Pr(X_j^\theta = x_j, Y_j^\theta = y_j)}{\Pr(X_{j+1}^\theta = x_{j+1}, Y_{j+1}^\theta = y'_j)}.$$

Then, to obtain the average work produced from a particular input sequence $y_{0:L}$, we must average over all possible hidden-state sequences and $x_{0:L+1}$ and output sequences $y'_{0:L}$:

$$\left\langle W_{y_{0:L}}^\theta \right\rangle = \sum_{x_{0:L+1}, y'_{0:L}} \Pr(Y'_{0:L} = y'_{0:L}, X_{0:L+1} = x_{0:L+1} | Y_{0:L} = y_{0:L}) \left\langle W_{y_{0:L}, y'_{0:L}, x_{0:L+1}}^\theta \right\rangle.$$

The probability of hidden states and outputs is determined from the agent's logical architecture and input sequence

$$\Pr(Y'_{0:L} = y'_{0:L}, X_{0:L+1} = x_{0:L+1} | Y_{0:L} = y_{0:L}) = \Pr(X_0 = x_0) \prod_{k=0}^{L-1} M_{x_k, y_k \rightarrow x_{k+1}, y'_k},$$

yielding equation (7). □

On its own, equation (7)'s work production is a deeply interesting quantity. In point of fact, since our agents are stochastic Turing machines [43], this is the work production for any general form of computation that maps inputs to output distributions $\Pr(Y'_{0:L} | Y_{0:L} = y_{0:L})$ [44]. Thus, equation (7) determines the possible *work benefit for universal thermodynamic computing*. Differing critically from prior explorations of information ratchets [29, 30, 42, 45] that evaluate average energy production, this expression is for a *particular input sequence*.

Given this general expression for work production, one might conclude that the next step for thermodynamic learning is to search for the agent tuple $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$ that maximizes the work production. However, this strategy comes with two issues. First, it requires a wider search than necessary. Second, it does not draw a direct connection to the underlying model θ of the time series. Recall that we are considering ϵ -machine models θ of the input sequence that give the probability estimate $\Pr(Y_{0:L}^\theta = y_{0:L})$ for any L .

We address both of these issues by refining the search space to agents whose anticipated inputs $\Pr(X_j^\theta, Y_j^\theta)$ are explicitly determined by their initial state distribution $\Pr(X_0^\theta)$ and estimated input process $\Pr(Y_{0:\infty}^\theta)$:

$$\Pr(X_j^\theta = x_j, Y_j^\theta = y_j) = \sum_{x_{0:j}, y_{0:j}, y'_{0:j}} \Pr(X_0^\theta = x_0) \Pr(Y_{0:j+1}^\theta = y_{0:j+1}) \prod_{k=0}^{j-1} M_{x_k, y_k \rightarrow x_{k+1}, y'_k}.$$

We use this estimate for the initial state in equation (7), since it is maximally efficient, dissipating as little as possible if the agent architecture M receives the input distribution $\Pr(Y_{0:\infty}^\theta)$. As a result of its efficiency, the resulting computation performed by the agent produces the maximum possible work, given its logical architecture. This simplifies our search for maximum-work agents by directly tying the estimated inputs to the model θ . However, it still leaves one piece of the agent undefined: its logical architecture M . Fortunately, as we discuss now, the thermodynamics of modularity further simplifies the search.

4.4. Thermodynamics of modularity

An agent transduces inputs to outputs through a series of modular operations. The Hamiltonian $\mathcal{H}_{\mathcal{X} \times \mathcal{Y}_j}(t)$ that governs the evolution of the j th operation is decoupled from the other elements of the time series $\mathcal{Y}_0 \times \mathcal{Y}_1 \dots \mathcal{Y}_{j-1} \times \mathcal{Y}_{j+1} \times \dots$. As a result of this modular computational architecture, the correlations lost between the agent and the rest of the information reservoir are irreversibly dissipated, producing entropy. This is an energetic cost associated directly with the agent's logical architecture, known as the *modularity dissipation* [45]. It sets a minimum for the work dissipated in a complex computation composed of many elementary steps.

To continue our pursuit of maximum-work, we must design the agent's logical architecture to minimize dissipated work. Past analysis of agents that harness energy from a pattern $\Pr(Y_{0:\infty})$ showed that the modularity dissipation is only minimized when the agent's states are *predictive* of the pattern [45]. This means that to maximize work extracted, an agent's state must contain all information about the past relevant to the future $\Pr(Y_{j+1:L} | X_j) = \Pr(Y_{j+1:L} | Y_{0:j})$. That is, for maximal work extraction agent states must be sufficient statistics for predicting the future.

Moreover, the ϵ -machines introduced in section 3.2 are constructed with hidden states S_j that are a *minimal* predictor of their output process. This is why they are referred to as *causal states*. And so, the ϵ -machine is a minimal sufficient statistic for prediction. Transitions among causal states trigger outputs according to:

$$\theta_{s \rightarrow s'}^{(y)} = \Pr(Y_j^\theta = y, S_{j+1}^\theta = s' | S_j^\theta = s),$$

which is the probability that an ϵ -machine in internal state s transitions to s' and emits output y . ϵ -Machines have the additional property that they are unifilar, meaning that the next causal state s' is uniquely determined by the current state s and output y via the propagator function $s' = \epsilon(s, y)$.

In short, to produce maximum work agent memory must store at least the causal states of the environment's own ϵ -machine. Appendix F describes how the logical architecture of the maximum-work minimum-memory agent is determined by its estimated ϵ -machine model θ of its inputs. The agent states are chosen to be the same as the causal states $\mathcal{X} = \mathcal{S}$, and they update according to the propagator function ϵ :

$$M_{xy \rightarrow x'y'} = \frac{1}{|\mathcal{Y}_j|} \times \begin{cases} \delta_{x', \epsilon(x, y)}, & \text{if } \sum_{x'} \theta_{x \rightarrow x'}^{(y)} \neq 0, \\ \delta_{x', x}, & \text{otherwise.} \end{cases} \quad (8)$$

This guarantees that the agent's internal states follow the causal states of its estimated input process $\Pr(Y_{0:\infty}^\theta)$. In turn, it prevents the agent from dissipating temporal correlations within that process.

4.5. Agent-model equivalence

In the pursuit of maximum work, we refined the structure of candidate agents considerably, limiting consideration to those that minimize the modularity dissipation. As a result, they store the predictive states of their estimated input and their logical architecture is explicitly determined by an ϵ -machine model. Moreover, to be efficient, the candidate agent should begin in the ϵ -machine's start state s^* such that the model θ uniquely determines the second piece of an efficient agent. This is the estimated initial distribution over its memory state and input:

$$\Pr(Y_j^\theta = y_j, X_j^\theta = s_j) = \sum_{y_{0:j}, s_{0:j}, s_{j+1}} \delta_{s_0, s^*} \prod_{k=0}^j \theta_{s_k \rightarrow s_{k+1}}^{(y_k)}. \quad (9)$$

Conversely, maximum-work agents, characterized by their logical architecture and anticipated input distributions $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$, also specify the ϵ -machine's model of their estimated distribution:

$$\begin{aligned} \theta_{s \rightarrow s'}^{(y)} &= \Pr(Y_j^\theta = y | S_j^\theta = s) \delta_{s', \epsilon(s, y)} \\ &= \Pr(Y_j^\theta = y | X_j^\theta = s) |\mathcal{Y}_j| M_{sy \rightarrow s'y'}. \end{aligned} \quad (10)$$

Through the ϵ -machine, the agent also specifies its estimated input process. In this way, we arrive at a class of agents that are uniquely determined by their environment model.

Figure 7 explicitly lays this out. It presents an agent that estimates a period-2 process with uncertain phase, such that $\Pr(Y_{0:\infty}^\theta = 0101 \dots) = \Pr(Y_{0:\infty}^\theta = 1010 \dots) = 0.5$. The middle column shows the ϵ -machine that is uniquely determined for that process, characterized by the model parameters $\theta_{s \rightarrow s'}^{(y)}$. The right column shows the unique minimal agent that harnesses as much work as possible from that process. All three, the estimated process, the estimated ϵ -machine model, and the minimum-memory maximum-work agent are equivalent. By *equivalent*, we mean that each can be determined from another. This holds true for *any* estimated input process $\Pr(Y_{0:\infty}^\theta)$:

Theorem 3. *The following are equivalent for an input process, up to relabeling of hidden/causal states:*

- The distribution over inputs $\Pr(Y_{0:L}^\theta)$.
- The ϵ -machine model θ that predicts the process.
- The agent that harvests maximum work from the process $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$.

Proof.

- (b) \implies (a): as stated in section 3.2

$$\Pr(Y_{0:L}^\theta = y_{0:L}) = \prod_{j=0}^{L-1} \theta_{\epsilon(s^*, y_{0:j}) \rightarrow \epsilon(s^*, y_{0:j+1})}^{(y_j)}.$$

- (a) \implies (b): the ϵ -machine θ , which is unique up to causal state relabeling, can be determined from applying a causal equivalence relation to $\Pr(Y_{0:L}^\theta)$ as described in appendix A.
- (c) \implies (b): equation (10) determines θ from $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$.
- (b) \implies (c): equations (9) and (8) determine $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$ from θ .

□

Under the equivalence of model θ and agent operation, when we monitor the agent's thermodynamic performance through its work production, we also measure the predictive performance of its underlying model.

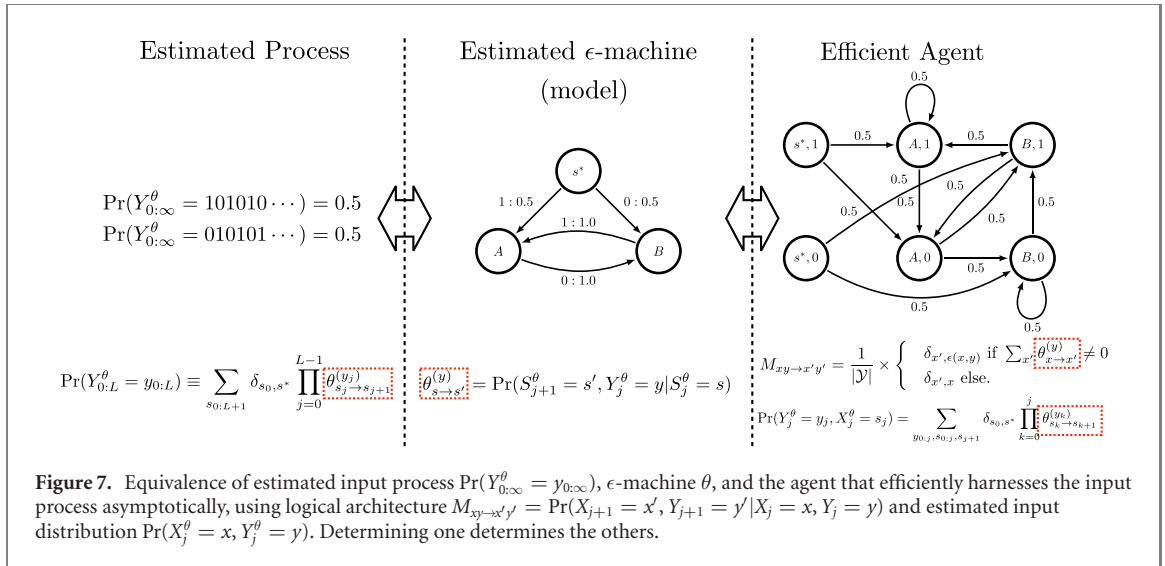


Figure 7. Equivalence of estimated input process $\Pr(Y_{0:\infty}^\theta = y_{0:\infty})$, ϵ -machine θ , and the agent that efficiently harnesses the input process asymptotically, using logical architecture $M_{xy \rightarrow x'y'} = \Pr(X_{j+1} = x', Y_{j+1} = y' | X_j = x, Y_j = y)$ and estimated input distribution $\Pr(X_j^\theta = x, Y_j^\theta = y)$. Determining one determines the others.

This completes the thermodynamic learning framework laid out in figure 1. There, the model an agent holds affects its interaction with the symbol sequence $y_{0:L}$ and, ultimately, its work production. And so, from this point forward, when discussing an estimated process or an ϵ -machine that generates that guess, we are also describing the unique thermodynamic agent designed to produce maximal work from the estimated process. We can now turn to explore how such agents' work production ties to their underlying models. A direct comparison to MLE can now be drawn.

5. Work-likelihood correspondence for agent design

We are now ready to return to our core objective—exploring work production as a performance measure for a model estimated from a time series $y_{0:L}$. In comparison to the expression for general computing in equation (7), using efficiently-designed predictive agents leads to a much simpler expression for work production:

Theorem 4. *The work produced by a maximum-work agent with estimated input process $\Pr(Y_{0:\infty}^\theta)$ when it receives the particular sequence $y_{0:L}$ is:*

$$\langle W_{|y_{0:L}}^\theta \rangle = k_B T (\ln \Pr(Y_{0:L}^\theta = y_{0:L}) + L \ln |\mathcal{Y}|). \tag{11}$$

Proof. See appendix G. □

The mechanism behind this vast simplification arises from unifilarity—a property of prediction machines that guarantees a single state trajectory on $x_{0:L}$ for each input string $y_{0:L}$.

This expression directly captures the relationship between work production and the agent's underlying model of the data. From it, we arrive at the central result of thermodynamic learning.

Theorem 5. *Maximizing the work produced from the time series $y_{0:L}$ also maximizes the likelihood of the agent's model θ .*

Proof. $\Pr(Y_{0:L}^\theta = y_{0:L})$ is simply the probability that the candidate model will output $y_{0:L}$. Therefore, the log-likelihood $\ell(\theta|y_{0:L})$ of MLE coincides with $\ln \Pr(Y_{0:L}^\theta = y_{0:L})$, and we can express the work production of an efficient agent:

$$\langle W_{|y_{0:L}}^\theta \rangle = k_B T \ell(\theta|y_{0:L}) + k_B T L \ln |\mathcal{Y}|. \tag{12}$$

One concludes that work production is maximized precisely when the log-likelihood $\ell(\theta|y_{0:L})$ is maximized. □

Thus, the criterion for *creating a good model of an environment* is the same as that for *extracting maximal work*. An analogous result, linking maximum work to maximum likelihood, was recently found for a many-particle Szilard engine [46].

This link is made concrete via the simple example presented in appendix H. It goes through an explicit description of the Hamiltonian control required to implement a memoryless agent that harvests work from a sequence of up spins \uparrow and down spins \downarrow that compose the time series $y_{0:L}$. The agent's internal

memoryless model results in equation (12)'s work production. And, we find that the maximum-work agent has learned about the input sequence. Specifically, the agent learns the frequency of spins \uparrow and \downarrow , confirming the basic principle of maximum-work thermodynamic learning. However, the learning presented in appendix H precludes the possibility of learning temporal structure in the spin sequence, since the agents and their internal models have no memory [31]. To learn about the temporal correlations within the sequence, one must use agents with multiple memory states. We leave thermodynamic learning among memoryful agents for later investigation. Stepping back, we see the relationship between machine learning and information thermodynamics more clearly. In MLE we have:

- (a) Data $y_{0:L}$ that provides a window into a black box.
- (b) A model θ of the black box that determines an estimated distribution over the data $\Pr(Y_{0:L}^\theta)$.
- (c) A performance measure for the model of the data, given by the log-likelihood

$$\ell(\theta|y_{0:L}) = \ln \Pr(Y_{0:L}^\theta = y_{0:L}).$$

The parallel in thermodynamic learning is exact, with:

- (a) Data $y_{0:L}$ physically stored in systems $\mathcal{Y}_0 \times \mathcal{Y}_1 \times \dots \times \mathcal{Y}_{L-1}$ output from the black box.
- (b) An agent $\{M, \Pr(X_j^\theta, Y_j^\theta)\}$ that is entirely determined by the model θ .
- (c) The agent's thermodynamic performance, given by its work production $\langle W_{|y_{0:L}}^\theta \rangle$, increases linearly with the log-likelihood $\ell(\theta|y_{0:L})$.

In this way, we see that thermodynamic learning through work maximization is *equivalent* to MLE.

Intuitively, the natural world is replete with complex learning systems—an observation seemingly at odds with thermodynamics and its second law which dictates that order inevitably decays into disorder. However, our results are tantamount to a contravening physical principle that drives the emergence of order through learning: work maximization. We showed, in point of fact, that work maximization and learning are equivalent processes. At a larger remove, this hints of general physical principles of emergent organization.

6. Searching for principles of organization

Introducing an *equivalence of maximum work production and optimal learning* comes at a late stage of a long line of inquiry into what kinds of thermodynamic constraints and laws govern the emergence of organization and, for that matter, biological life. So, let us historically place the seemingly-new principle. In fact, it enters a crowded field.

Within statistical physics the paradigmatic principle of organization was found by Kirchhoff [47]: in electrical networks current distributes itself so as to dissipate the least possible heat for the given applied voltages. Generalizations, for equilibrium states, are then found in Gibbs' variational principle for entropy for heterogeneous equilibrium [48], Maxwell's principles of minimum-heat [49, pp 407–408], and Onsager's minimizing the 'rate of dissipation' [50].

Close to equilibrium, Prigogine introduced minimum entropy production [51], identifying *dissipative structures* whose maintenance requires energy [52]. However, far from equilibrium the guiding principles can be quite the opposite. And so, the effort continues today, for example, with recent applications of nonequilibrium thermodynamics to pattern formation in chemical reactions [53]. That said, statistical physics misses at least two, related, but key components: dynamics of and information in thermal states.

Dynamical systems theory takes a decidedly mechanistic approach to the emergence of organization, analyzing the geometric structures in a system's state space that amplify fluctuations and eventually attenuate them into macroscopic behaviors and patterns. This was eventually articulated by pattern formation theory [54–56]. A canonical example is fluid turbulence [57]—a dynamical explanation for its complex organizations occupied much of the 70s and 80s. Landau's original theory of incommensurate oscillations was superseded by the mathematical discovery in the 1950s of chaotic attractors [58, 59]. This approach, too, falls short of leading to a principle of emergent organization. Patterns emerge, but what exactly are they and what complex behavior do they exhibit?

Answers to this challenge came from a decidedly different direction—Shannon's theory of noisy communication channels and his measures of information [60, 61], appropriately extended [62]. While adding an important new perspective—that organized systems store and transmit information—this, also, did not go far enough as it side-stepped the content and meaning of information [63]. In-roads to these appeared in the theory of computation inaugurated by Turing [64]. The most direct and ambitious approach to the role of information in organization, though, appeared in Wiener's cybernetics [4, 65].

While it eloquently laid out the goals to which principles should strive, it ultimately never harnessed the mathematical foundations and calculational tools needed. Likely, the earliest overt connection between statistical mechanics and information, though, appeared with Jaynes' maximum entropy [66] and minimum entropy production principles [67].

So, what is new today is the synthesis of statistical physics, dynamics, and information. This, finally, allows one to answer the question, how do physical systems store and process information? The answer is that they *intrinsically compute* [36]. With this, one can extract from behavior a system's information processing, even going so far as to discover the effective equations of motion [68–71]. One can now frame questions about how a physical system reacts to, controls, and adapts to its environment.

All such systems, however, are embedded in the physical world and require resources to operate. More to the point, what energetic resources underlie computation? Initiated by Brillouin [72] and Landauer and Bennett [19, 73], today there is a nascent *physics of information* [14, 74]. Resource constraints on computing by thermodynamic systems are now expressed in a suite of new principles.

For example, the *thermodynamics of predicting* complex patterns links formal notions of inference to optimal energy manipulation [75]. Along this line of investigation, the *principle of requisite complexity* [31] dictates that maximally-efficient interactions require an agent's internal organization match the environment's organization. This stems from the thermodynamic resource costs that arise from the modularity of an agent's architecture [45]. Pushing the search for organization further, the preceding established a principle for the thermodynamics of learning. And this, looking forward, gives a basis for a thermodynamics of adaptive organization.

To fully appreciate organization in natural processes, though, one must also address dynamics of agent populations, first on the time scale of agent life cycles and second on the scale of many generations. In fact, tracking the complexity of individuals reveals that selection pressures spontaneously emerge in purely-replicating populations [76] and replication itself necessarily dissipates energy [77].

As these pieces assembled, a picture has come into focus. Intelligent, adaptive systems learn to harness resources from their environment, expending energy to live and reproduce. Taken altogether, the historical perspective suggests we are moving close to realizing Wiener's cybernetics [4].

7. Conclusion

We introduced thermodynamic machine learning—a physical process that trains intelligent agents by maximizing work production from complex environmental stimuli supplied as time-series data. This involved constructing a framework to describe thermodynamics of computation at the single-shot level, enabling us to evaluate the work an agent can produce from individual data realizations. Key to the framework is its generality—applicable to agents exhibiting arbitrary adaptive input–output behavior and implemented within any physical substrate.

In the pursuit of maximum work, we refined this general class of agents to those that are best able to harness work from temporally-correlated inputs. We found that the performance of such maximum-work agents increases proportionally to the log-likelihood of the model they use for predicting their environment. As a consequence, our results show that thermodynamic learning exactly mimics MLE in machine learning. Thus, work is a thermodynamic performance measure for physically-embedded learning. This result further solidifies the connections between agency, intelligence, and the thermodynamics of information—hinting that energy harvesting and learning may be two sides of the same coin.

These connections suggest a number of exciting future directions. From the technological perspective, they hint at a natural method for designing intelligent energy harvesters—establishing that our present tools of machine learning can be directly mapped to automated design of efficient information ratchets and pattern engines [31, 42, 78]. Meanwhile, recent results indicate that quantum systems can generate complex adaptive behaviors using fewer resources than classical counterparts [79–81]. Does this suggest there are new classes of quantum-enhanced energy harvesters and learners?

Ultimately, energy is an essential currency for life. This highlights the question, to what extent is work optimization a natural tendency of driven physical systems? Indeed, recent results indicate complex many-body systems can evolve to increase work production [32, 33], opening a fascinating possibility. Could the equivalence between work production and learning then indicate that the Universe itself naturally learns? The fact that complex intelligent life emerged from the lifeless soup of the Universe might be considered a continuing miracle: a string of unfathomable statistical anomalies strung together over eons. It would certainly be extraordinary if this evolution then has a physical basis—hidden laws of thermodynamic organization that guide the Universe to create entities capable of extracting maximal work. Such a law has yet to be revealed, but the equivalence we established between work maximization and learning extends the

relevance of such thermodynamic predictions to intelligent learning systems. Technically, it gives a foundation to *embodied learning* by including physical energetics in machine learning.

Acknowledgments

The authors thank the Telluride Science Research Center for hospitality during visits and the participants of the Information Engines Workshops there. ABB thanks Wesley Boyd for useful conversations and JPC similarly thanks Adam Rupe. JPC acknowledges the kind hospitality of the Santa Fe Institute, Institute for Advanced Study at the University of Amsterdam, and California Institute of Technology for their hospitality during visits. MG thanks FQXi for the hospitality in workshop ‘Mind Matters: Intelligence and Agency in the Physical World’. This project is supported by the is supported by National Research Foundation (NRF) Singapore under its NRFF Fellow program (Award No. NRF-NRFF2016-02), Singapore Ministry of Education Tier 1 Grant Nos. RG146/20 and RG77/22, Singapore Ministry of Education Tier 2 Grant T2EP50221-0014, Grant Nos. FQXi-RFP-IPW-1902 and FQXi-RFP-1809 from the Foundational Questions Institute and Fetzer Franklin Fund (a donor-advised fund of Silicon Valley Community Foundation), the Templeton World Charity Foundation Power of Information fellowship TWCF0337, TWCF0560 and TWCF0570, and US Army Research Laboratory and the US Army Research Office under Grants W911NF-18-1-0028 and W911NF-21-1-0048. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation or Ministry of Education, Singapore. During submission we became aware of related work: L Touzo, M Marsili, N Merhav, and E Roldan. Optimal work extraction and the minimum description length principle. arXiv:2006.04544.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Appendix A. Extended background

Developing the principle of maximum work production and calling out the physical benefits of an agent modeling its environment drew heavily from the areas of computational mechanics, nonequilibrium thermodynamics, and machine learning. Due to the variety of topics addressed, the following provides a more detailed notational and conceptual summary. This should help the development to be more self-contained, hopefully providing a common language across the areas and a foundation for further exploration. While we make suggestive comparisons by viewing foundations of each area side-by-side, it may be most appropriate for readers already familiar with them and concerned with novel results to skip, using the review to clarify unfamiliar notation.

A.1. Machine learning and generative models

Thermodynamically, what is a good model of the data with which an agent interacts? Denote the data’s state space as $\mathcal{Z} = \{z\}$. If we have many copies of \mathcal{Z} , all initially prepared in the same way, then as we observe successive realizations $\vec{z} = \{z_0, z_1, \dots, z_N\}$ from an ensemble, the frequency of an observed state z approaches the *actual* probability distribution $\Pr(Z = z)$, where Z is the random variable that realizes states $z \in \mathcal{Z}$. However, with only a finite number N of realizations, the best that can be done is to characterize the environment with an *estimated* distribution $\Pr(Z^\theta = z)$. Estimating models that agree with finite data is the domain of statistical inference and machine learning of generative models [15, 34, 82].

At first blush, estimating a probability distribution appears distinct from familiar machine learning challenges, such as image classification and the inverse problem of artificially generating exemplar images from given data. However, both classification and prediction can be achieved through a form of unsupervised learning [18]. For instance, if the system is a joint variable over both the pixel images and the corresponding label $\mathcal{Z} = \text{pixels} \times \{\text{cat}, \text{dog}\}$, then our estimated distribution $\Pr(Z^\theta = z)$ gives both a means of choosing a label for an image $\Pr(\text{label}^\theta = \text{cat} | \text{pixels}^\theta = \text{image})$ and a means of choosing an image for a label $\Pr(\text{pixels}^\theta = \text{image} | \text{label}^\theta = \text{cat})$.

A generative model is specified by a set of parameters θ from which the model produces the estimated distribution $\Pr(Z^\theta = z) = \Pr(Z = z | \Theta = \theta)$. The procedure of arriving at this estimated model is parametric density estimation [15, 82]. However, we take the random variable Z^θ for the estimated distribution to denote the model for notational and conceptual convenience.

The Shannon entropy [61]:

$$H[Z] \equiv - \sum_z \Pr(Z = z) \ln \Pr(Z = z)$$

measures uncertainty in nats, a ‘natural’ unit for thermodynamic entropies. The Shannon entropy easily extends to joint probabilities and all information measures that come from their composition (conditional and mutual informations). For instance, if the environment is composed of two correlated subcomponents $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, the probability and entropy are expressed:

$$\begin{aligned} \Pr(Z = z) &= \Pr(X = x, Y = y) \\ &\neq \Pr(X = x) \Pr(Y = Y) \quad \text{and} \\ H[Z] &= H[X, Y] \\ &\neq H[X] + H[Y], \end{aligned}$$

respectively.

While there are many other ways to create parametric models θ —from polynomial functions with a small number of parameters to neural networks with thousands [15]—the goal is to match as well as possible the estimated distribution $\Pr(Z^\theta)$ to the actual distribution $\Pr(Z)$.

One measure of success in this is the probability that the model generated the data—the *likelihood*. The likelihood of the model θ given a data point z is the same as the likelihood of Z^θ :

$$\begin{aligned} \mathcal{L}(\theta|z) &= \Pr(Z = z|\Theta = \theta) \\ &= \Pr(Z^\theta = z) \\ &= \mathcal{L}(Z^\theta|z). \end{aligned}$$

Given a set $\vec{z} = \{z_1, z_2, \dots, z_N\}$ of training data and assuming independent samples, then the likelihood of the model is the product:

$$\mathcal{L}(Z^\theta|\vec{z}) = \prod_{i=1}^N \mathcal{L}(Z^\theta|z_i). \quad (\text{A1})$$

This is a commonly used performance measure in machine learning, where algorithms search for models with maximum likelihood [15]. However, it is common to use the *log-likelihood* instead, which is maximized for the same models:

$$\begin{aligned} \ell(\theta|\vec{z}) &= \ln \mathcal{L}(\theta|\vec{z}) \\ &= \sum_{i=1}^N \ln \Pr(Z^\theta = z_i). \end{aligned} \quad (\text{A2})$$

If the model Z^θ were specified by a neural network, the log-likelihood could be determined through stochastic gradient descent back-propagation [34, 83], for instance. The intention is that the procedure will converge on a network model that produces the data with high probability.

A.2. Thermodynamics of information

Learning from data translates information in an environment into a useful model. What makes that model useful? In a physical setting, recalling from Landauer that ‘information is physical’ [84], the usefulness one can extract from thermodynamic processes is work. Figure 3 shows a basic implementation for physical computation. Such an information-storing physical system $\mathcal{Z} = \{z\}$, in contact with a thermal reservoir, can execute useful computations by drawing energy from a work reservoir. Energy flowing from the system \mathcal{Z} into the thermal reservoir is positive heat Q . When energy flows from the system \mathcal{Z} to the work reservoir, it is positive work W production. Work production quantifies the amount of energy that is stored in the work reservoir available for later use. And so, in this telling, it represents a natural and physically-motivated measure of thermodynamic performance. In the framework for thermodynamic computation of figure 3, work is extracted via controlling the system’s Hamiltonian.

Specifically, the system’s informational states are controlled via a time-dependent Hamiltonian—energy $E(z, t)$ of state z at time t . For state trajectory $z_{\tau:\tau'} = z_\tau z_{\tau+\text{dt}} \dots z_{\tau'-\text{dt}} z_{\tau'}$ over time interval $t \in [\tau, \tau']$, the

resulting work extracted by Hamiltonian control is the temporally-integrated change in energy [37]:

$$W_{|z_{\tau:\tau'}} = - \int_{\tau}^{\tau'} dt \partial_t E(z, t)|_{z=z_t}.$$

Heat $Q_{|z_{\tau:\tau'}} = E(z_{\tau}, \tau) - E(z_{\tau'}, \tau') - W_{|z_{\tau:\tau'}}$ flows into the thermal reservoir, increasing its entropy:

$$\Delta S_{|z_{\tau:\tau'}}^{\text{reservoir}} = \frac{Q_{|z_{\tau:\tau'}}}{T}, \quad (\text{A3})$$

where the thermal reservoir is at temperature T . The second law of thermodynamics states that, on average, any processing on the informational states can only yield nonnegative entropy production of the Universe (reservoir and system \mathcal{Z}):

$$\begin{aligned} \langle \Sigma \rangle &= \langle \Delta S^{\text{reservoir}} \rangle + \langle \Delta S^{\mathcal{Z}} \rangle \\ &\geq 0. \end{aligned} \quad (\text{A4})$$

This constrains the energetic cost of computations performed within the system \mathcal{Z} .

A computation over time interval $t \in [\tau, \tau']$ has two components:

- An initial distribution over states $\Pr(Z_{\tau} = z_{\tau})$, where Z_t is the random variable of system \mathcal{Z} at time t .
- A Markov channel that transforms it, specified by the conditional probability of the final state $z_{\tau'}$ given the initial input z_{τ} :

$$M_{z_{\tau} \rightarrow z_{\tau'}} = \Pr(Z_{\tau'} = z_{\tau'} | Z_{\tau} = z_{\tau}).$$

This specifies, in turn, the final distribution $\Pr(Z_{\tau'} = z_{\tau'})$ that allows direct calculation of the system-entropy change [85]:

$$\Delta S_{|z_{\tau:\tau'}}^{\mathcal{Z}} = k_B \ln \frac{\Pr(Z_{\tau'} = z_{\tau'})}{\Pr(Z_{\tau} = z_{\tau})}.$$

Adding this to the information reservoir's entropy change yields the entropy production of the Universe. This can also be expressed in terms of the work production:

$$\begin{aligned} \Sigma_{|z_{\tau:\tau'}} &\equiv \Delta S_{|z_{\tau:\tau'}}^{\text{reservoir}} + \Delta S_{|z_{\tau:\tau'}}^{\mathcal{Z}} \\ &= \frac{-W_{|z_{\tau:\tau'}} + \phi(z_{\tau}, \tau) - \phi(z_{\tau'}, \tau')}{T}. \end{aligned}$$

Here, $\phi(z, t) = E(z, t) + k_B T \ln \Pr(Z_t = z)$ is the *pointwise nonequilibrium free energy*, which becomes the *nonequilibrium free energy* when averaged: $\langle \phi(z, t) \rangle_{\Pr(Z_t = z)} = F^{\text{neq}}(t)$ [86].

Note that the entropy production is also proportional to the additional work that could have been extracted if the computation was efficient. This is referred to as the *dissipated work*:

$$W_{|z_{\tau:\tau'}}^{\text{diss}} = T \Sigma_{|z_{\tau:\tau'}}. \quad (\text{A5})$$

Turning back to the second law of thermodynamics, we see that the *average* work extracted is bounded by the change in nonequilibrium free energy:

$$\langle W \rangle \geq F^{\text{neq}}(\tau') - F^{\text{neq}}(\tau).$$

When the system starts and ends as an information reservoir, with equal energies for all states $E(z, \tau) = E(z', \tau')$ [37], this reduces to Landauer's familiar principle for erasure [19]—work production must not exceed the change in state uncertainty:

$$\langle W \rangle \leq k_B T (H[Z_{\tau'}] - H[Z_{\tau}]),$$

where $H[Z_t] = -\sum_{z \in \mathcal{Z}} \Pr(Z_t = z) \ln \Pr(Z_t = z)$ is the Shannon entropy of the system at time t measured in nats. This is the starting point for determining the work production that agents can extract from data.

A.3. Computational mechanics

When describing thermodynamics and machine learning, data was taken from the state space \mathcal{Z} all at once. However, what if we consider a state space composed of L identical components $\mathcal{Z} = \mathcal{Y}^L$ that are received in sequence. Our model of the time series $y_{0:L}$ of realizations is described by an estimated distribution $\Pr(Y_{0:L}^{\theta} = y_{0:L})$. However, for L large enough, this object becomes impossible to store, due to the exponential increase in the number of sequences. Fortunately, there are ways to generally characterize an arbitrarily long time-series distribution using a finite model.

A.3.1. Generative machines

A hidden Markov model (HMM) is described by a set of hidden states \mathcal{S} , a set of output states \mathcal{Y} , a conditional output-labeled matrix that gives the transition probabilities between the states:

$$\theta_{s \rightarrow s'}^{(y)} = \Pr(S_{j+1}^\theta = s', Y_j^\theta = y | S_j^\theta = s)$$

for all j , and a start state $s^* \in \mathcal{S}$. (Generally, one also specifies an initial state distribution, with selecting a start state being a special case.) We label the transition probabilities with the model parameter θ , since these are the actual parameters that must be stored to generate probabilities of time series. For instance, figure 2 shows an HMM that generates a periodic process with uncertain phase. Edges between hidden states s and s' are labeled $y: \theta_{s \rightarrow s'}^{(y)}$, where y is the output symbol and $\theta_{s \rightarrow s'}^{(y)}$ is the probability of emitting that symbol on that transition.

If $\{\theta_{s \rightarrow s'}^{(y)}\}$ is the model for the estimated input $\Pr(Y^\theta = y_{0:L})$, then the probability of any word is calculated by taking the product of transition matrices and summing over internal states:

$$\Pr(Y_{0:L}^\theta = y_{0:L}) = \sum_{s_0, s^*} \delta_{s_0, s^*} \prod_{j=0}^{L-1} \theta_{s_j \rightarrow s_{j+1}}^{(y_j)}.$$

Beyond generating length- L symbol strings, these HMMs generate distributions over semi-infinite strings $\Pr(Y_{0:\infty}^\theta = y_{0:\infty})$. As such, they allow us to anticipate more than just the first L symbols from the same source. Once we have a model $\theta = \{\theta_{s \rightarrow s'}^{(y)}, s, s', y\}_{s, s', y}$ from our training data $y_{0:L}$, we can calculate probabilities of longer words $\Pr(Y_{0:L'}^\theta = y'_{0:L'})$ and, thus, the probability of symbols following the training data $\Pr(Y_{L:L'}^\theta = y'_{L:L'} | Y_{0:L}^\theta = y_{0:L})$. Distributions over semi-infinite strings $\Pr(Y_{0:\infty}^\theta = y_{0:\infty})$ are similar to processes, which are distributions over bi-infinite strings $\Pr(Y_{-\infty:\infty}^\theta = y_{-\infty:\infty})$. While not insisting on stationarity, and so allowing for flexibility in using subwords of length L , we can mirror computational mechanics' construction of unifilar HMMs from time series, where the hidden states S_j^θ are minimal sufficient statistics of the past $Y_{0:j}^\theta$ about the future $Y_{j:\infty}^\theta$ [87]. In other words, the hidden states are perfect predictors.

Given a semi-infinite process $\Pr(Y_{0:\infty}^\theta = y_{0:\infty})$, we construct a minimal predictor through a causal equivalence relation $y_{0:k} \sim y'_{0:j}$ that says two histories $y_{0:k}$ and $y'_{0:j}$ are members of the same equivalence class if and only if they have the same semi-infinite future distribution:

$$\Pr(Y_{j:\infty}^\theta = y''_{0:\infty} | Y_{0:j}^\theta = y'_{0:j}) = \Pr(Y_{k:\infty}^\theta = y''_{0:\infty} | Y_{0:k}^\theta = y_{0:k}).$$

An equivalence class of histories is a *causal state*. Causal states also induce a map $\epsilon(\cdot)$ from histories $y_{0:k}$ to states s_i :

$$\begin{aligned} s_i &= \{y'_{0:j} | y_{0:k} \sim y'_{0:j}\} \\ &\equiv \epsilon(y_{0:k}). \end{aligned}$$

This guarantees that a causal state is a sufficient statistic of the past about the future, such that we can track it as a perfect predictor:

$$\Pr(Y_{k:\infty}^\theta | Y_{0:k}^\theta = y_{0:k}) = \Pr(Y_{k:\infty}^\theta | S_k^\theta = \epsilon(y_{0:k})).$$

In fact, the causal states are minimal sufficient statistics. Constructing causal states reveals a number of properties of stochastic processes and models. One of these is *unifilarity*, which means that if the current causal state $s_k = \epsilon(y_{0:k})$ is followed by any sequence $y_{k:j}$, then the resulting causal state $s_k = \epsilon(y_{0:k})$ is uniquely determined. And, we can expand our use of the ϵ function to include updating a causal state:

$$\begin{aligned} s_j &= \epsilon(s_k, y_{k:j}) \\ &\equiv \{y'_{0:l} | \exists y_{0:k} \ni s_k = \epsilon(y_{0:k}) \text{ and } y'_{0:l} \sim y_{0:j}\}. \end{aligned}$$

This is the set of histories $y'_{0:l}$ that predict the same future as a history $y_{0:j}$ which leads to causal state s_k via the initial sequence $y_{0:k}$ and then follows with the sequence $y_{k:j}$.

Unifilarity is key to deducing several useful properties of the HMM $\theta_{s \rightarrow s'}^{(y)}$, which we will refer to as a *nonstationary* ϵ -machine. First, unifilarity implies that for any causal state s followed by a symbol y , there is a unique next state $s' = \epsilon(s, y)$, meaning that the symbol-labeled transition matrix can be written:

$$\theta_{s \rightarrow s'}^{(y)} = \theta_{s \rightarrow \epsilon(s, y)}^{(y)} \delta_{s', \epsilon(s, y)}.$$




Figure 8. The delay channel ϵ -transducer: the last input symbol is stored in its memory (states). If the last symbol was 1, then the corresponding transitions, labeled $y'|1 : 1.0$, update the hidden state to A. Then, all outputs from A are symbol 1. Similarly, input 0 leads to state B, whose corresponding outputs are all 0. In this way, the delay channel outputs the previous input symbol.

Moreover, the ϵ -machine's form is uniquely determined from the semi-infinite process:

$$\theta_{s \rightarrow s'}^{(y)} = \Pr(Y_j^\theta = y | S_j^\theta = s) \delta_{s', \epsilon(s, y)},$$

where the conditional probability is determined from the process:

$$\Pr(Y_j^\theta = y | S_j^\theta = \epsilon(y_{0:i})) = \Pr(Y_j^\theta = y | Y_{0:j}^\theta = y_{0:j}).$$

Once constructed, the ϵ -machine allows us to reconstruct word probabilities via the simple product:

$$\Pr(Y_{0:L}^\theta = y_{0:L}) = \prod_{j=0}^{L-1} \theta_{\epsilon(s^*, y_{0:j}) \rightarrow \epsilon(s^*, y_{0:j+1})}^{(y_j)},$$

where $y_{0:0}$ denotes the null word, taking a causal state to itself under the causal update $\epsilon(s, y_{0:0}) = s$.

Allowing for arbitrarily-many causal states, our class of models (nonstationary ϵ -machines) is so general that it can represent *any* semi-infinite process and, thus, any distribution over sequences \mathcal{Y}^L . One concludes that computational mechanics provides an ideal class of generative models to fit to data $y_{0:L}$. *Bayesian structural inference* implements just this [88].

In these ways, computational mechanics already had solved (and several decades prior) the unsupervised learning challenge recently posed by reference [89] to create an 'AI physicist': a machine that learns regularities in time series to make predictions of the future from the past [90].

A.3.2. Input–output machines

This way one constructs a predictive HMM that generates a desired semi-infinite process $\Pr(Y_{0:L}^\theta = y_{0:L})$. The generalization to ϵ -transducers allows for an *input* as well as an *output* process—a transformation between processes [44]. The transducer at the i th time step is described by transitions among the hidden states $X_i \rightarrow X_{i+1}$, that are conditioned on the input Y_i , emitting the output Y'_i :

$$M_{x \rightarrow x'}^{(y'|y)} = \Pr(Y'_i = y', X_{i+1} = x' | Y_i = y, X_i = x).$$

ϵ -Transducer state-transition diagrams label the edges of transitions between hidden states $y'|y : M_{x \rightarrow x'}^{(y'|y)}$. As figure 8 shows this is to be read as the probability $M_{x \rightarrow x'}^{(y'|y)}$ of output y' and next hidden state x' given input y and current hidden state x .

These devices are memoryful channels, with their memory encoded in the hidden states X_i . They implement a wide variety of functional operations. Figure 8 shows the *delay channel*. With sufficient memory, though, an ϵ -transducer can implement a universal Turing machine [43]. Moreover, if the input and output alphabets are the same, then they represent the form of a physical information ratchet, which have energetic requirements that arise from the thermodynamics of their operation [29, 30]. Since these physically-implementable information processors are so general in their ability to compute, they represent a very broad class of physical agents. As such, we use the framework of information ratchets to explore the functionality of agents that process information as a fuel.

Appendix B. Proof of computational trajectory work

To determine the work produced for a computational trajectory $z_\tau \rightarrow z_{\tau'}$, we first prove a useful relation between the entropy and work production for a particular state trajectory $z_{\tau:\tau'}$. Specifically, let $W_{|z_{\tau:\tau'}}$ and $\Sigma_{|z_{\tau:\tau'}}$ denote the work and total entropy production along this trajectory, respectively. Let $E(z_t, t)$ denote the system energy when it is in state z_t at time t . Now, consider the *pointwise nonequilibrium free energy*:

$$\phi(z_t, t) = E(z_t, t) + k_B T \ln \Pr(Z_t = z_t). \quad (\text{B1})$$

More familiarly, note that its time-averaged quantity is the *nonequilibrium free energy* [86]:

$$F^{\text{neq}} = \langle \phi(z, t) \rangle_{\text{Pr}(Z_t=z)}.$$

We can then show that the *entropy production* Σ can be expressed:

$$\Sigma|_{z_\tau, \tau'} = \frac{-W|_{z_\tau, \tau'} + \phi(z_\tau, \tau) - \phi(z_{\tau'}, \tau')}{T}. \quad (\text{B2})$$

This follows by noting that the total entropy produced from thermodynamic control is the sum of the entropy change in the system [85]:

$$\Delta S|_{z_\tau, \tau'}^{\mathcal{Z}} = k_B \ln \frac{\text{Pr}(Z_\tau = z_\tau)}{\text{Pr}(Z_{\tau'} = z_{\tau'})}$$

and that of the thermal reservoir:

$$\Delta S|_{z_\tau, \tau'}^{\text{reservoir}} = \frac{Q|_{z_\tau, \tau'}}{T}.$$

Equation (B2) follows by summing up these contributions to the total entropy production $\Sigma = \Delta S^{\text{reservoir}} + \Delta S^{\mathcal{Z}}$ and noting that the system of interest's (SOI's) change in energy obeys the first law of thermodynamics $\Delta E^{\mathcal{Z}} = -W - Q$.

Since only the SOI's initial and final states matter to the logical operation of the computational map, we take a statistical average of all trajectories beginning in z_τ and ending in $z_{\tau'}$. This results in the work production:

$$\langle W|_{z_\tau, z_{\tau'}} \rangle = \sum_{z'_{\tau: \tau'}} W|_{z'_{\tau: \tau'}} \text{Pr}(Z_{\tau: \tau'} = z'_{\tau: \tau'} | z_\tau, z_{\tau'}), \quad (\text{B3})$$

for the computational map $z_\tau \rightarrow z_{\tau'}$. This determines how much energy is stored in the work reservoir on average when a computation results in this particular input–output pair.

Similarly, taking the same average of the entropy production shown in equation (B2), conditioned on inputs and outputs, gives:

$$\begin{aligned} T \langle \Sigma|_{z_\tau, z_{\tau'}} \rangle &= -\langle W|_{z_\tau, z_{\tau'}} \rangle + \phi(z_\tau, \tau) - \phi(z_{\tau'}, \tau'), \\ &= -\langle W|_{z_\tau, z_{\tau'}} \rangle - \Delta \phi|_{z_\tau, z_{\tau'}}. \end{aligned}$$

This suggestively relates computational-mapping work and the change in pointwise nonequilibrium free energy $\phi(z, t)$.

This relation between work and free energy simplifies for thermodynamically-efficient computations. In such scenarios, the average total entropy production over all trajectories vanishes. Appendix D shows that zero average entropy production, combined with the Crooks fluctuation theorem [91, 92], implies that entropy production along any individual trajectory $z_{\tau: \tau'}$ produces zero entropy: $\Sigma|_{z_\tau, \tau'} = 0$. This is expected from linear response [93].

Thus, substituting zero entropy production into equation (B2), we arrive at our result: *work production for thermodynamically-efficient computations is the change in pointwise nonequilibrium free energy*:

$$W|_{z_\tau, \tau'}^{\text{eff}} = -\Delta \phi|_{z_\tau, z_{\tau'}}.$$

Substituting equation (B1) then gives:

$$W|_{z_\tau, \tau'}^{\text{eff}} = -\Delta E_{\mathcal{Z}} + k_B T \ln \frac{\text{Pr}(Z_\tau = z_\tau)}{\text{Pr}(Z_{\tau'} = z_{\tau'})},$$

where $\Delta E_{\mathcal{Z}} = E(z_{\tau'}, \tau') - E(z_\tau, \tau)$. This is what we would expect in quasistatic computations, where the system energies $E(z, t)$ are varied slowly enough that the system \mathcal{Z} remains in equilibrium for the duration. We should note, though, that it is possible to implement efficient computations rapidly and out of equilibrium [94].

This also holds if we average over intermediate states of the SOI's state trajectory, yielding the work production of a computational map:

$$\langle W|_{z_\tau, z_{\tau'}}^{\text{eff}} \rangle = -\Delta E_{\mathcal{Z}} + k_B T \ln \frac{\text{Pr}(Z_\tau = z_\tau)}{\text{Pr}(Z_{\tau'} = z_{\tau'})}. \quad (\text{B4})$$

The energy required to perform efficient computing is independent of intermediate properties. It depends only on the probability and energy of initial and final states. This measures the energetic gains from a single data realization as it transforms during a computation, as opposed to the ensemble average.

Appendix C. Energetics of estimates

SOI state probabilities feature centrally in the expression for nonequilibrium free energy and, thus, for the work production of efficient agents. However, the *actual* input distribution $\Pr(Z_\tau)$ may vary while the agent, defined by its Hamiltonian $\mathcal{H}_Z(t)$ over the computation interval, remains fixed. Moreover, since the work production $\langle W_{|z_\tau, z_{\tau'}}^{\text{eff}} \rangle$ of a computational map explicitly conditions on the initial and final SOI state, this work cannot explicitly depend on the input distribution. At first blush, this is a contradiction: work that simultaneously does and does not depend on the input distribution.

This is resolved once one recognizes the role that estimates play in thermodynamics. As indicated in figure 1, we claim that an agent has an estimated model of its environment that it uses to predict the SOI. This model is, in one form or another, encoded in the evolving Hamiltonian $\mathcal{H}_Z(t)$ that determines both the agent's energetic interactions and its logical architecture. If an agent's estimated model of SOI \mathcal{Z} is encoded as parameters θ , then the agent estimates that the SOI state z at time t has probability:

$$\Pr(Z_t^\theta = z_t) = \Pr(Z_t = z_t | \Theta = \theta).$$

The physical relevance of the estimated distribution comes from insisting that *the agent dissipates as little work as possible from a SOI whose distribution matches its own estimate*. In essence, the initial estimated distribution $\Pr(Z_\tau^\theta)$ must be one of the distributions that minimizes the average entropy production [95]:

$$\Pr(Z_\tau^\theta) \in \arg \min_{\Pr(Z_\tau)} \langle \Sigma [\Pr(Z_\tau)] \rangle.$$

Estimated probabilities $\Pr(Z_t^\theta)$ at later times $t > \tau$ are determined by updating the initial estimate via the stochastic dynamics that result from the Hamiltonian $\mathcal{H}_Z(t)$ interacting with the thermal bath.

Thus, since an efficient agent produces zero entropy when the SOI follows the minimum dissipation distribution $\Pr(Z_t^\theta)$, the work it produces from a computational map is:

$$\langle W_{|z_\tau, z_{\tau'}}^\theta \rangle = -\Delta E_Z + k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{\Pr(Z_{\tau'}^\theta = z_{\tau'})}. \quad (\text{C1})$$

In this, we replaced the superscript 'eff' with ' θ ' to emphasize that the agent is designed to be thermodynamically efficient for that particular estimated model. Specifying the estimated model is essential, since misestimating the input distribution leads to dissipation and entropy production [95, 96]. Returning to thermodynamic learning, this is how the model θ factors into the ratchet's operation: estimated distributions explicitly determine the work production of computational maps.

Appendix E gives a concrete quasistatic mechanism for implementing *any* computation M and achieving the work given by equation (C1). This directly demonstrates how the model θ is built into the evolving energy landscape $\mathcal{H}_Z(t)$ that implements M . The model θ determines the initial and final change in state energies: $\Delta E(z, \tau) = -k_B \ln \Pr(Z_\tau^\theta = z)$ and $\Delta E(z, \tau') = k_B \ln \Pr(Z_{\tau'}^\theta = z)$. This quasistatic protocol operates and produces the same work for a particular input–output pair regardless of the *actual* input distribution.

Since we focus on the energetic benefits derived from information itself rather than those from changing energy levels, the example implementations we use also start and end with the same flat energy landscape. Restricting to such information-driven agents, we consider cases where $\Delta E_Z = 0$, whereby:

$$\langle W_{|z_\tau, z_{\tau'}}^\theta \rangle = k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{\Pr(Z_{\tau'}^\theta = z_{\tau'})}. \quad (\text{C2})$$

Appendix D. Proof of zero entropy production of trajectories

Perfectly-efficient agents dissipate zero work and generate zero entropy $\langle \Sigma \rangle = 0$. The Crooks fluctuation theorem and other detailed fluctuation theorems say that entropy production is proportional to the log-ratio of probabilities [91, 92]:

$$\Sigma_{|z_\tau, z_{\tau'}} = k_B \ln \frac{\rho_F(\Sigma_{|z_\tau, z_{\tau'}})}{\rho_R(-\Sigma_{|z_\tau, z_{\tau'}})},$$

where $\rho_F(\Sigma|_{z_\tau, \tau'})$ is the probability of the entropy production under the protocol that controls the ratchet and $\rho_R(-\Sigma|_{z_\tau, \tau'})$ is the probability of minus that same entropy production if the control protocol is reversed. Thus, the average entropy production is proportional to the relative entropy between these two distributions [61]:

$$\begin{aligned} \langle \Sigma \rangle &= k_B \sum_{z_\tau, \tau'} \rho_F(\Sigma|_{z_\tau, \tau'}) \ln \frac{\rho_F(\Sigma|_{z_\tau, \tau'})}{\rho_R(-\Sigma|_{z_\tau, \tau'})} \\ &\equiv k_B D_{KL}(\rho_F(\Sigma|_{z_\tau, \tau'}) || \rho_R(-\Sigma|_{z_\tau, \tau'})). \end{aligned}$$

If the control is thermodynamically efficient, this relative entropy vanishes [97], implying the necessary and sufficient condition that $\rho_F(\Sigma) = \rho_R(-\Sigma)$. This then implies that all paths produce zero entropy:

$$\Sigma|_{z_\tau, \tau'} = 0.$$

Entropy fluctuations vanish as the entropy production goes to zero.

Appendix E. Thermodynamically efficient Markov channels

Given a physical system $\mathcal{Z} = \{z\}$, a computation on its states is given by a Markov channel $M_{z \rightarrow z'} = \Pr(Z_{\tau'} = z' | Z_\tau = z)$ and an input distribution $\Pr(Z_\tau = z)$. The following describes a quasistatic thermodynamic control that implements this computation efficiently if the input distribution matches the estimated distribution $\Pr(Z^\theta = z)$. This means that the work production is equal to the change in pointwise nonequilibrium free energy:

$$\begin{aligned} \langle W_{|z_\tau, z_\tau'}^\theta \rangle &= \phi(z_\tau, \tau) - \phi(z_\tau', \tau') \\ &= -\Delta E_{\mathcal{Z}} + k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{\Pr(Z_\tau^\theta = z_\tau')}. \end{aligned} \quad (E1)$$

Note that, while $\Pr(Z_\tau^\theta = z)$ is the input distribution for which the computation is efficient, it is possible that other input distributions $\Pr(Z_\tau = z)$ yield zero entropy production as well. They are only required to minimize $D_{KL}(Z_\tau || Z_\tau^\theta) - D_{KL}(Z_{\tau'} || Z_{\tau'}^\theta) = 0$ [95].

The physical setting that we take for thermodynamic control is overdamped Brownian motion with a controllable energy landscape. This is described by detailed-balanced rate equations. However, if our physical state-space is limited to \mathcal{Z} , then not all channels can be implemented with continuous-time rate equations [94]. Fortunately, this can be circumvented by additional ancillary or hidden states [94, 98]. And so, to implement any possible channel, we add an ancillary copy of our original system \mathcal{Z}' , such that our entire physical system is $\mathcal{Z}_{\text{total}} = \mathcal{Z} \times \mathcal{Z}'$. However, since the ancillary copy starts and ends in the same uniform uncorrelated distribution, it does not affect the energy and entropy calculations.

Prescriptions have been given that efficiently implement any computation, specified by a Markov channel $M_{z_\tau \rightarrow z_{\tau'}} = \Pr(Z_{\tau'} = z_{\tau'} | Z_\tau = z_\tau)$, using quasistatic manipulation of the \mathcal{Z} 's energy levels and an ancillary copy \mathcal{Z}' [42, 45]. However, these did not determine the work production for individual computational maps $z_\tau \rightarrow z_{\tau'}$ during the computation interval (τ, τ') .

The following implements an analogous form of quasistatic computation that allows us to easily calculate the energy associated with implementing the computation $M_{z_\tau \rightarrow z_{\tau'}}$, assuming the subsystem \mathcal{Z} started in z_τ and ends in $z_{\tau'}$. Due to detailed balance, the rate equation dynamics over the computational system and its ancillary copy $\mathcal{Z}_{\text{total}} = \mathcal{Z} \times \mathcal{Z}'$ are partially specified by the energy $E(z, z', t)$ of system state z and ancillary state z' at time t . This also uniquely specifies the equilibrium distribution:

$$\Pr(Z_t^{\text{eq}} = z, Z_t^{\text{eq}} = z') = \frac{e^{-E(z, z', t)/k_B T}}{\sum_{z, z'} e^{-E(z, z', t)/k_B T}}.$$

The normalization constant $\sum_{z, z'} e^{-E(z, z', t)/k_B T}$ is the partition function that determines the equilibrium free energy:

$$F^{\text{eq}}(t) = -k_B T \ln \left(\sum_{z, z'} e^{-E(z, z', t)/k_B T} \right).$$

The equilibrium free energy adds to the system energy. It is constant over the states:

$$E(z, z', t) = F^{\text{eq}}(t) - k_B T \ln \Pr(Z_t^{\text{eq}} = z, Z_t^{\text{eq}} = z').$$

We leverage the relationship between energy and equilibrium probability to design a protocol that achieves the work production given by equation (E2) for a Markov channel M . The estimated distribution over the whole space assumes that the initial distribution of the ancillary variable is uncorrelated and uniformly distributed:

$$\Pr(Z_\tau^\theta, Z_{\tau'}^\theta) = \frac{\Pr(Z_\tau^\theta)}{|\mathcal{Z}|}.$$

Assuming the default energy landscape is constant initially and finally— $E(z, z', \tau) = E(z, z', \tau') = \xi$ —the maximally efficient protocol over the interval $[\tau, \tau']$ decomposes into five epochs, see figure 9:

- (a) Quench: $[\tau, \tau^+]$,
- (b) Quasistatic transformation: $(\tau, \tau_1]$,
- (c) Swap: $(\tau_1, \tau_2]$,
- (d) Quasistatic transformation: $(\tau_2, \tau']$, and
- (e) Reset: $[\tau'^-, \tau']$.

For all protocol epochs, except for epoch 3 during which the two subsystems are swapped, \mathcal{Z} is held fixed while the ancillary system \mathcal{Z}' follows the local equilibrium distribution. Let us detail these in turn.

1. Quench: instantaneously quench the energy from $E(z, z', \tau) = \xi$ to $E(z, z', \tau^+) = k_B T \ln(|\mathcal{Z}| / \Pr(Z_\tau = z))$ over the infinitesimal time interval $[\tau, \tau^+]$ such that, if the distribution was as we expect, it would be in equilibrium $\Pr(Z_\tau^{\text{eq}}, Z_\tau^{\text{eq}}) = \Pr(Z_\tau^\theta) / |\mathcal{Z}|$.

If the system started in z_τ , then the associated work produced is opposite the energy change:

$$\begin{aligned} \left\langle W_{|z_\tau, z_{\tau'}}^{\theta, 1} \right\rangle &= E(z_\tau, z', \tau) - E(z_\tau, z', \tau^+) \\ &= \xi + k_B T \ln \frac{\Pr(Z_\tau^\theta = z_\tau)}{|\mathcal{Z}|}. \end{aligned}$$

$\left\langle W_{|z_\tau, z_{\tau'}}^{\theta, 1} \right\rangle$ denotes that the work is produced in epoch 1, conditioned on the estimated distributions Z_τ^θ and $Z_{\tau'}^\theta$, and initial and final states z_τ and $z_{\tau'}$. Note that we also condition on $Z_{\tau'}^\theta = z_{\tau'}$, since work production in this phase is unaffected by the computation's end state.

2. Quasistatic transformation: quasistatically evolve the energy landscape over a third of total time interval $(\tau, \tau_1]$ such that the joint system remains in equilibrium and the ancillary system \mathcal{Z}' is determined by the Markov channel M applied to the system \mathcal{Z} :

$$\begin{aligned} \Pr(Z_{\tau_1} = z, Z'_{\tau_1} = z') &= \Pr(Z_\tau = z) M_{z \rightarrow z'} \\ E(z, z', \tau_1) &= -k_B T \ln \Pr(Z_\tau^\theta = z) M_{z \rightarrow z'}. \end{aligned}$$

Also, hold the energy barriers between states in \mathcal{Z} high, preventing probability flow between states and preserving the distribution $\Pr(Z_t) = \Pr(Z_\tau)$ for all $t \in (\tau, \tau_1]$.

Given that the system started in $Z_\tau = z_\tau$, the work production during this epoch corresponds to the average change in energy:

$$\begin{aligned} \left\langle W_{|z_\tau, z_{\tau'}}^{\theta, 2} \right\rangle &= - \sum_{z, z'} \int_{\tau^+}^{\tau_1} dt \Pr(Z_t' = z', Z_t = z | Z_\tau = z_\tau) \partial_t E(z, z', t). \end{aligned}$$

As the system \mathcal{Z} remains in z_τ over the interval:

$$\Pr(Z_t' = z', Z_t = z | Z_\tau = z_\tau) = \Pr(Z_t' = z' | Z_t = z) \delta_{z, z_\tau}$$

and the work production simplifies to:

$$\left\langle W_{|z_\tau, z_{\tau'}}^{\theta, 2} \right\rangle = - \sum_{z'} \int_{\tau^+}^{\tau_1} dt \Pr(Z_t' = z' | Z_t = z_\tau) \partial_t E(z_\tau, z', t).$$



We can express the energy in terms of the estimated equilibrium probability distribution:

$$E(z_\tau, z', t) = -k_B T \ln \Pr(Z'_t = z' | Z_t = z_\tau) \Pr(Z_t^0 = z_\tau).$$

And, since the distribution over the system \mathcal{Z} is fixed during this interval:

$$\Pr(Z'_t = z' | Z_t = z_\tau) \Pr(Z_t^0 = z_\tau) = \Pr(Z'_t = z' | Z_t = z_\tau) \Pr(Z_\tau^0 = z_\tau).$$

Plugging these into the expression for the work production, we find that the evolution happens without energy exchange:

$$\begin{aligned}
\langle W_{|z_\tau, z_{\tau'}}^{\theta, 2} \rangle &= -k_B T \int_{\tau^+}^{\tau_1} dt \sum_{z'} \Pr(Z'_t = z' | Z_t = z_\tau) \partial_t \ln \Pr(Z'_t = z' | Z_t = z_\tau) \Pr(Z'_\tau = z_\tau) \\
&= -k_B T \int_{\tau^+}^{\tau_1} dt \sum_{z'} \Pr(Z'_t = z' | Z_t = z_\tau) \frac{\Pr(Z'_\tau = z_\tau) \partial_t \Pr(Z'_t = z' | Z_t = z_\tau)}{\Pr(Z'_t = z' | Z_t = z_\tau) \Pr(Z'_\tau = z_\tau)} \\
&= -k_B T \int_{\tau^+}^{\tau_1} dt \sum_{z'} \partial_t \Pr(Z'_t = z' | Z_t = z_\tau) \\
&= -k_B T \int_{\tau^+}^{\tau_1} dt \partial_t \sum_{z'} \Pr(Z'_t = z' | Z_t = z_\tau) \\
&= -k_B T \int_{\tau^+}^{\tau_1} dt \partial_t 1 \\
&= 0.
\end{aligned}$$

After this state, at time τ_1 the resulting joint distribution over the ancillary and primary system matches the desired computation:

$$\Pr(Z_{\tau_1} = z, Z'_{\tau_1} = z') = \Pr(Z_\tau = z) M_{z \rightarrow z'}. \quad (\text{E2})$$

3. Swap: over time interval $(\tau_1, \tau_2]$, efficiently swap the two systems $\mathcal{Z} \leftrightarrow \mathcal{Z}'$, such that $\Pr(Z_{\tau_1} = z, Z'_{\tau_1} = z') = \Pr(Z_{\tau_2} = z', Z'_{\tau_2} = z)$ and $E(z, z', \tau_1) = E(z', z, \tau_2)$. This operation requires zero work as well, as it is reversible, regardless of where the system starts or ends:

$$\langle W_{|z_\tau, z_{\tau'}}^{\theta, 3} \rangle = 0.$$

Such an efficient swap operation has been demonstrated in finite time [94]. The resulting joint distribution over the ancillary and primary system matches a flip of the desired computation:

$$\Pr(Z_{\tau_2} = z, Z'_{\tau_2} = z') = \Pr(Z_\tau = z') M_{z' \rightarrow z}.$$

4. Quasistatically evolve: over time interval (τ_2, τ') , quasistatically evolve the energy landscape from:

$$E(z, z', \tau_2) = -k_B T \ln \Pr(Z'_\tau = z') M_{z' \rightarrow z}$$

to

$$\begin{aligned}
E(z, z', \tau'^-) &= -k_B T \ln \frac{\sum_{z''} \Pr(Z'_\tau = z'') M_{z'' \rightarrow z}}{|\mathcal{Z}|} \\
&= -k_B T \ln \frac{\Pr(Z'_\tau = z)}{|\mathcal{Z}|}.
\end{aligned}$$

We keep the primary system \mathcal{Z} fixed as in epoch 2. And, as in epoch 2, there is zero work production:

$$\langle W_{|z_\tau, z_{\tau'}}^{\theta, 4} \rangle = 0.$$

The result is that the primary system is in the desired final distribution:

$$\Pr(Z_{\tau'} = z) \equiv \sum_{z'} \Pr(Z_\tau = z') M_{z' \rightarrow z},$$

having undergone a mapping from its original state at time τ , while the ancillary system has returned to an uncorrelated uniform distribution.

5. Reset: finally, over time interval $[\tau'^-, \tau']$ instantaneously change the energy to the default flat landscape $E(z, z', \tau') = \xi$. Given that the system ends in state $z_{\tau'}$, the associated work production is:

$$\begin{aligned} \langle W_{z_{\tau}, z_{\tau'}}^{\theta, 5} \rangle &= E(z'_{\tau}, z', \tau'^-) - E(z'_{\tau}, z', \tau') \\ &= -\xi - k_B T \ln \frac{\Pr(Z_{\tau'}^{\theta} = z_{\tau'})}{|\mathcal{Z}|}. \end{aligned}$$

The net work production given the initial state z_{τ} and final state $z_{\tau'}$ is then:

$$\begin{aligned} \langle W_{z_{\tau}, z_{\tau'}}^{\theta} \rangle &= \sum_{i=1}^5 \langle W_{z_{\tau}, z_{\tau'}}^{\theta, i} \rangle \\ &= k_B T \ln \frac{\Pr(Z_{\tau}^{\theta} = z_{\tau})}{\Pr(Z_{\tau'}^{\theta} = z_{\tau'})}. \end{aligned}$$

Thus, when we average over all possible inputs and outputs and the estimated an actual distributions are the same $Z_{\tau}^{\theta} = Z_{\tau}$, we see that this protocol achieves the thermodynamic Landauer's bound:

$$\begin{aligned} \langle W^{\theta} \rangle &= \sum_{z_{\tau}, z_{\tau'}} \Pr(Z_{\tau} = z_{\tau}, Z_{\tau'} = z_{\tau'}) \langle W_{z_{\tau}, z_{\tau'}}^{\theta} \rangle \\ &= k_B T \ln 2(H[Z_{\tau'}] - H[Z_{\tau}]). \end{aligned}$$

One concludes that this is a thermodynamically-efficient method for computing any Markov channel.

Appendix F. Designing agents

To design a predictive thermodynamic agent, its hidden states must match the states of the ϵ -machine at every time step $X_i = S_i^{\theta}$. To do this, the agent states and causal states occupy the same space $\mathcal{X} = \mathcal{S}$, and the transitions within the agent M are directly drawn from causal equivalence relation:

$$M_{xy \rightarrow x'y'} = \frac{1}{|\mathcal{Y}|} \times \begin{cases} \delta_{x', \epsilon(x, y)} & \text{if } \sum_{x'} \theta_{x \rightarrow x'}^{(y)} \neq 0, \\ \delta_{x', x} & \text{otherwise.} \end{cases}$$

The factor $1/|\mathcal{Y}|$ maximizes work production by mapping to uniform outputs.

The second term on the right is the probability of the next agent state given the current input and current hidden state $\Pr(X_{i+1} = x' | Y_i = y, X_i = x)$. The top case $\delta_{x', \epsilon(x, y)}$ gives the probability that the next causal state is $S_{i+1}^{\theta} = x'$ given that the current causal state is $S_i^{\theta} = x$ and output of the ϵ -machine is $Y_i = y$. This is contingent on the probability of seeing y given causal state x being nonzero. If it is, then the transitions among the agent's hidden states match the transitions of the ϵ -machine's causal states.

In this way, if $y_{0:L}$ is a sequence that could be produced by the ϵ -machine, we have designed the agent to stay synchronized to the causal state of the input $X_i = S_i^{\theta}$, so that the ratchet is predictive of the process $\Pr(Y_{0:\infty}^{\theta})$ and produces maximal work by fully randomizing the outputs:

$$\Pr(Y'_i = y'_i | \cdot) = \frac{1}{|\mathcal{Y}|}.$$

It can be the case that the ϵ -machine cannot produce y from the causal state x . This corresponds to a disallowed transition of our model $\theta_{x \rightarrow x'}^{(y)} = 0$. In these cases, we arbitrarily choose the next state to be the same $\delta_{x, x'}$. There are many possible choices, though—such as resetting to the start state s^* . However, it is physically irrelevant, since these transitions correspond zero estimated probability and, thus, to infinite work dissipation, drowning out all other details of the model. However, this particular choice for when y cannot be generated from the causal state x preserves unifilarity and allows the agent to wait in its current state until it receives an input that it can accept.

Modulo disallowed, infinitely-dissipative transitions, we now have a direct mapping between our estimated input process $\Pr(Y_{0:\infty}^{\theta})$ and its ϵ -machine θ to the logical architecture M of a maximum work-producing agent.

As yet, this does not fully specify agent behavior, since it leaves out the estimated input distribution $\Pr(Y_i^{\theta} = y, X_i^{\theta} = x)$. This distribution must match the actual distribution $\Pr(Y_i = y, X_i = x)$ for the agent to be locally efficient, not accounting for temporal correlations. Fortunately, since agent states are designed to

match the ϵ -machine's causal states, we know that the agent state distribution matches the causal-state distribution and inputs:

$$\Pr(Y_i^\theta = y_i, X_i^\theta = s_i) = \Pr(Y_i^\theta = y_i, S_i^\theta = s_i),$$

if the ratchet is driven by the estimated input. The joint distribution over causal states and inputs is also determined by the ϵ -machine, since the construction assumes starting in the state $s_0 = s^*$. To start, note that the joint probability trajectory distribution is given by:

$$\begin{aligned} & \Pr(Y_{0:i+1}^\theta = y_{0:i+1}, S_{0:i+2}^\theta = s_{0:i+2}) \\ &= \delta_{s_0, s^*} \prod_{j=0}^i \theta_{s_j \rightarrow s_{j+1}}^{(y_j)}. \end{aligned}$$

Summing over the variables besides Y_i^θ and S_i^θ , we obtain an expression for the estimated agent distribution in terms of just the ϵ -machine's HMM:

$$\Pr(Y_i^\theta = y_i, X_i^\theta = s_i) = \sum_{y_{0:i}, s_{0:i}, s_{i+1}} \delta_{s_0, s^*} \prod_{j=0}^i \theta_{s_j \rightarrow s_{j+1}}^{(y_j)}.$$

Thus, an agent $\{M, \Pr(X_i^\theta, Y_i^\theta)\}$ designed to be *globally efficient* for the estimated input process $\Pr(Y_{0:L}^\theta)$ can be derived from the estimated input process through its ϵ -machine $\theta_{s \rightarrow s'}$.

Appendix G. Work production of optimal transducers

The work production of an arbitrary transducer M driven by an input $y_{0:L}$ can be difficult to calculate, as shown in equation (7). However, when the transducer is designed to harness an input process with ϵ -machine T , such that:

$$M_{xy \rightarrow x'y'} = \frac{1}{|\mathcal{Y}|} \times \begin{cases} \delta_{x', \epsilon(x, y)} & \text{if } \sum_{x'} \theta_{x \rightarrow x'}^{(y)} \neq 0, \\ \delta_{x', x} & \text{else,} \end{cases}$$

the work production simplifies. To see this, we express equation (7) in terms of the estimated distribution $\Pr(Y_{0:L}^\theta)$, ratchet M , and input $y_{0:L}$, assuming that the word $y_{0:L}$ can be produced from every initial hidden state with nonzero probability $\Pr(Y_{0:L}^\theta = y_{0:L} | S_0 = s_0) \Pr(X_0 = s_0) \neq 0$, which guarantees that $\sum_{x_{i+1}} \theta_{x_i \rightarrow x_{i+1}}^{(y_i)} \neq 0$. If this constraint is not satisfied, the agent will dissipate infinite work, as it implies $\Pr(X_i^\theta = x_i, Y_i^\theta = x_i) = 0$ for some i . Thus, we use the expression $M_{xy \rightarrow x'y'} = \delta_{x', \epsilon(x, y)} / |\mathcal{Y}|$ in the work production:

$$\begin{aligned} & \langle W_{y_{0:L}}^\theta \rangle \\ &= k_B T \sum_{x_{0:L+1}, y'_{0:L}} \Pr(X_0 = x_0) \prod_{j=0}^{L-1} M_{x_j, y_j \rightarrow x_{j+1}, y'_j} \ln \prod_{i=0}^{L-1} \frac{\Pr(X_i^\theta = x_i, Y_i^\theta = y_i)}{\Pr(X_{i+1}^\theta = x_{i+1}, Y_i^\theta = y'_i)} \\ &= k_B T \sum_{x_{0:L+1}, y'_{0:L}} \Pr(X_0 = x_0) \prod_{j=0}^{L-1} \frac{\delta_{x_{j+1}, \epsilon(x_j, y_j)}}{|\mathcal{Y}|} \ln \prod_{i=0}^{L-1} \frac{\Pr(X_i^\theta = x_i, Y_i^\theta = y_i)}{\Pr(X_{i+1}^\theta = x_{i+1}) / |\mathcal{Y}|} \\ &= k_B T \ln |\mathcal{Y}| + k_B T \sum_{x_{0:L+1}} \Pr(X_0 = x_0) \prod_{j=0}^{L-1} \delta_{x_{j+1}, \epsilon(x_j, y_j)} \left(\ln \prod_{i=0}^{L-1} \Pr(Y_i^\theta = y_i | X_i^\theta = x_i) + \ln \prod_{i=0}^{L-1} \frac{\Pr(X_i^\theta = x_i)}{\Pr(X_{i+1}^\theta = x_{i+1})} \right). \end{aligned}$$

Note that $\prod_{j=0}^{L-1} \delta_{x_{j+1}, \epsilon(x_j, y_j)}$ vanishes unless each element of the hidden state trajectory x_i corresponds to the resulting state of the ϵ -machine when the initial state x_0 is driven by the first i inputs $y_{0:i}$, which is $\epsilon(x_0, y_{0:i})$. The fact that the agent is driven into a unique state is guaranteed by the ϵ -machine's unifilarity. Thus, we rewrite:

$$\prod_{j=0}^{L-1} \delta_{x_{j+1}, \epsilon(x_j, y_j)} = \prod_{j=1}^L \delta_{x_j, \epsilon(x_0, y_{0:j})}.$$

This engenders a simplification of the work production:

$$\begin{aligned}
 \langle W_{|y_{0:L}}^\theta \rangle &= k_B T \ln |\mathcal{Y}| + k_B T \sum_{x_{0:L+1}} \Pr(X_0 = x_0) \prod_{j=1}^L \delta_{x_j, \epsilon(x_0, y_{0:j})} \ln \prod_{i=0}^{L-1} \frac{\Pr(Y_i^\theta = y_i | X_i^\theta = x_i) \Pr(X_i^\theta = x_i)}{\Pr(X_{i+1}^\theta = x_{i+1})} \\
 &= k_B T \ln |\mathcal{Y}| + k_B T \sum_{x_0} \Pr(X_0 = x_0) \ln \prod_{i=0}^{L-1} \Pr(Y_i^\theta = y_i | X_i^\theta = \epsilon(x_0, y_{0:i})) \\
 &\quad + k_B T \sum_{x_0} \Pr(X_0 = x_0) \ln \frac{\Pr(X_0^\theta = x_0)}{\Pr(X_L^\theta = \epsilon(x_0, y_{0:L}))}, \tag{G1}
 \end{aligned}$$

where $y_{0:0}$ denotes the null input, which leaves the state fixed under the ϵ -map $\epsilon(x_0, y_{0:0}) = x_0$.

This brings us to an easily calculable work production, especially when the system is initialized in the start state s^* . Recognizing that if we initiate the ϵ -machine in its start state s^* , such that $\Pr(X_0 = x_0) = \delta_{x_0, s^*}$, then $X_0 = S_0$ is predictive. By extension, every following agent state is predictive and equivalent to the causal state $X_i = S_i$ yielding:

$$\begin{aligned}
 \Pr(Y_i = y_i | X_i = \epsilon(s^*, y_{0:i})) &= \Pr(Y_i = y_i | S_i = \epsilon(s^*, y_{0:i})) \\
 &= \Pr(Y_i = y_i | Y_{0:i} = y_{0:i}).
 \end{aligned}$$

Thus, the work production simplifies a sum of terms that includes the log-likelihood of the finite input:

$$\begin{aligned}
 \langle W_{|y_{0:L}}^\theta \rangle &= k_B T \sum_{x_0} \delta_{x_0, s^*} \ln \prod_{i=0}^{L-1} \Pr(Y_i^\theta = y_i | Y_{0:i} = y_{0:i}) + k_B T L \ln |\mathcal{Y}| + k_B T \sum_{x_0} \delta_{x_0, s^*} \ln \frac{\delta_{x_0, s^*}}{\Pr(X_L^\theta = \epsilon(x_0, y_{0:L}))} \\
 &= k_B T (\ln \Pr(Y_{0:L}^\theta = y_{0:L}) + L \ln |\mathcal{Y}| - \ln \Pr(X_L^\theta = \epsilon(s^*, y_{0:L}))).
 \end{aligned}$$

The first $\log(\cdot)$ in the last line is the log-likelihood of the model generating $y_{0:L}$ —a common performance measure for machine learning algorithms. If an input has zero probability this leads to $-\infty$ work production, and all other features are drowned out by the log-likelihood term. Thus, the additional terms that come into play when the input probability vanishes become physically irrelevant: the agent is characterized by the ϵ -machine. From a machine learning perspective, the model is also characterized by the ϵ -machine $\theta_{s \rightarrow y}^{(y)}$, for the process $\Pr(Y_{0:\infty}^\theta)$. The additional term $k_B T L \ln |\mathcal{Y}|$ is the work production that comes from exhausting fully randomized outputs and does not change depending on the underlying model.

The final term $-k_B T \ln \Pr(X_L^\theta = \epsilon(s^*, y_{0:L}))$ does directly depend on the model. $\Pr(X_L^\theta = x)$ is the distribution over agent states \mathcal{X} at time $L\tau$ if the agent is driven by the estimated input distribution $Y_{0:L}^\theta$. This component of the work production is larger, on average, for agents with high state uncertainty, since this leads, on-average, to smaller values of $\Pr(X_L^\theta)$. This contribution to the work production comes from the state space expanding from the start state s^* to the larger (recurrent) subset of agent states, and so it provides additional work. This indicates that we are neglecting the cost of resetting to the start state while harnessing the energetic benefit of starting in it.

If the machine is designed to efficiently harness inputs again after it operates on one string, it must be reset to the start state s^* . This can be implemented with an efficient channel that anticipates the input distribution $\Pr(X_L^\theta = x)$, outputs the distribution $\Pr(X_{L+1}^\theta = x) = \delta_{x, s^*}$, and so costs:

$$W_{|y_{0:L}}^{\theta, \text{reset}} = k_B T \ln \Pr(X_L^\theta = \epsilon(s^*, y_{0:L})).$$

Thus, when we add the cost of resetting the agent to the start state at X_{L+1} , the work production is dominated by the log-likelihood:

$$\langle W_{|y_{0:L}}^\theta \rangle = k_B T (\ln \Pr(Y_{0:L}^\theta = y_{0:L}) + L \ln |\mathcal{Y}|). \tag{G2}$$

Appendix H. Training simple agents

We now outline a case study of thermodynamic learning that is experimentally implementable using a controllable two-level system. We first introduce a straightforward method to implement the simplest possible efficient agent. Second, we show that this physical process achieves the general



Figure 10. Memoryless model of binary data consisting of a single state A and the probability of outputting a \uparrow and a \downarrow , denoted $\theta_{A \rightarrow A}^{(\uparrow)}$ and $\theta_{A \rightarrow A}^{(\downarrow)}$, respectively.

maximum-likelihood result arrived at in the main development. Last, we find the agent selected by thermodynamic learning along with its corresponding model. As expected, we find that this maximum-work producing agent learns its environment's predictive features.

H.1. Efficient computational trajectories

The simplest possible information ratchets have only a single internal state A and receive binary data y_j from a series of two-level systems $\mathcal{Y}_j = \{\uparrow, \downarrow\}$. These agents' internal models correspond to memoryless ϵ -machines, as shown in figure 10. The model's parameters are the probabilities of emitting \uparrow and \downarrow , denoted $\theta_{A \rightarrow A}^{(\uparrow)}$ and $\theta_{A \rightarrow A}^{(\downarrow)}$, respectively.

Our first step is to design an efficient computation that maps an input distribution $\Pr(Z_{j\tau})$ to an output distribution $\Pr(Z_{j\tau+\tau'})$ over the j th interaction interval $[j\tau, j\tau + \tau']$. The agent corresponds to the Hamiltonian evolution $\mathcal{H}_Z(t) = \mathcal{H}_{\mathcal{X} \times \mathcal{Y}_j}(t)$ over the joint space of the agent memory and j th input symbol. The resulting energy landscape $E(z, t)$ is entirely specified by the energy of the two input states $E(A \times \uparrow, t)$ and $E(A \times \downarrow, t)$.

Appropriately designing this energy landscape allows us to implement the efficient computation. The procedure for an efficient computation on two states is laid out in figure 6. It can be easily generalized to include the single memory state A . The thermodynamic evolution there instantaneously quenches the energy landscape into equilibrium with the estimated distribution at the beginning of the interaction interval $\Pr(Z_{j\tau}^\theta)$, then quasistatically evolves the system in equilibrium to the estimated final distribution $\Pr(Z_{j\tau+\tau'}^\theta)$, and, finally, quenches back to the default energy landscape. The system undergoes a cycle, starting and ending with the same flat energy landscape, such that $\Delta E_Z = 0$.

We control the transformation over time interval $t \in (j\tau, j\tau + \tau')$ such that the time scale of equilibration in the SOI is much shorter than the interval length τ' . This slow-moving quasistatic control means that the states are in equilibrium with the energy landscape over the interval. In this case, the state distribution becomes the Boltzmann distribution:

$$\Pr(Z_t = z) = e^{(E^{\text{EQ}}(t) - E(z, t))/k_B T}.$$

To minimize dissipation for the estimated distribution, the state distribution must be the estimated distribution $\Pr(Z_t = z) = \Pr(Z_t^\theta = z)$. And so, we set the two-level-system energies to be in equilibrium with the estimates:

$$E(z, t) = -k_B T \ln \Pr(Z_t^\theta = z).$$

The resulting process produces zero work:

$$\begin{aligned} W^{\text{quasistatic}} &= - \int_{z=j\tau}^{j\tau+\tau'} dt \sum_z \Pr(Z_t^\theta = z) \partial_t E(z, t) \\ &= 0 \end{aligned}$$

and maps $\Pr(Z_{j\tau}^\theta)$ to $\Pr(Z_{j\tau+\tau'}^\theta)$ without dissipation.

With the quasistatic transformation producing zero work, the total work produced from the initial joint state $x \times y$ is exactly opposite the change in energy during the initial quench:

$$E(x \times y, j\tau) - E(x \times y, j\tau^+) = k_B T \ln \Pr(Z_{j\tau}^\theta = x \times y)$$

minus the change in energy of the final joint state $x' \times y'$ during the final quench:

$$E(x' \times y', j\tau + \tau'^-) - E(x' \times y', j\tau + \tau') = -k_B T \ln \Pr(Z_{j\tau+\tau'}^\theta = x' \times y').$$

The two-level system's state is fixed during the instantaneous energy changes. Thus, if the joint state follows the computational mapping $x \times y \rightarrow x' \times y'$ the work production is, as expected, directly connected to the estimated distributions:

$$\langle W_{|x \times y, x' \times y'} \rangle = k_B T \ln \frac{\Pr(Z_{j\tau}^\theta = x \times y)}{\Pr(Z_{j\tau+\tau'}^\theta = x' \times y')}. \quad (\text{H1})$$

Recall from section 4.1 that the ratchet system variable $Z_{j\tau}^\theta = X_j^\theta \times Y_j^\theta$ splits into the random variable X_j^θ for the j th agent memory state and the j th input Y_j^θ . Similarly, $Z_{j\tau+\tau'}^\theta = X_{j+1}^\theta \times Y_j^\theta$ splits into the $(j+1)$ th agent memory state X_{j+1}^θ and j th output Y_j^θ . This work production achieves the efficient limit for a model θ ($\langle W_{|x \times y, x' \times y'} \rangle = \langle W_{|x \times y, x' \times y'}^\theta \rangle$) described in equation (5).

Appendix E generalized the thermodynamic operation above to *any* computation $M_{z_\tau \rightarrow z_{\tau'}}$. While it requires an ancillary copy of the system \mathcal{Z} to execute the conditional dependencies in the computation, it is conceptually identical in that it uses a sequence of quenching, evolving quasistatically, and then quenching again. This appendix extends the strategies outlined in references [42, 45] to computational-mapping work calculations.

H.2. Efficient information ratchets

With the method for efficiently mapping inputs to outputs in hand, we can design a series of such computations to implement a simple information ratchet that produces work from a series $y_{0:L}$. As prescribed in equation (8) of section 5, to produce the most work from estimated model θ , the agent's logical architecture should randomly map every state to all others:

$$\begin{aligned} M_{xy \rightarrow x'y'} &= \frac{1}{|\mathcal{Y}_j|} \\ &= \frac{1}{2}, \end{aligned}$$

since there is only one causal state A . In conjunction with equation (10), we find that the estimated joint distribution of the agent and interaction symbol at the start of the interaction is equivalent to the parameters of the model:

$$\begin{aligned} \Pr(Z_{j\tau}^\theta = x \times y) &= \Pr(X_j^\theta = x, Y_j^\theta = y) \\ &= \Pr(Y_j^\theta = y | X_j^\theta = A) \Pr(X_j^\theta = A) \\ &= \theta_{A \rightarrow A}^{(y)}, \end{aligned}$$

where we again used the fact that A is the only causal state. In turn, the estimated distribution after the interaction is:

$$\begin{aligned} \Pr(Z_{j\tau+\tau'}^\theta = x' \times y') &= \sum_{xy} \Pr(X_j^\theta = x, Y_j^\theta = y) M_{xy \rightarrow x'y'} \\ &= \frac{1}{2}. \end{aligned}$$

Thus, assuming the agent has model θ built-in, then equation (H1) determines that the work production for mapping $A \times y$ to output $A \times y'$ for a particular symbol y is:

$$\langle W_{|A \times y, A \times y'} \rangle = k_B T \left(\ln 2 + \ln \theta_{A \rightarrow A}^{(y)} \right).$$

Since A is the only memory state and work does not depend on the output symbol y' , the average work produced from an input y is:

$$\langle W_{|y} \rangle = \langle W_{|A \times y, A \times y'} \rangle. \quad (\text{H2})$$

With the work production expressed for a single input y_j , we can now consider how much work our designed agent harvests from the training data $y_{0:L}$. Summing the work production of each input yields

a simple expression in terms of the model θ :

$$\begin{aligned}\langle W_{|y_{0:L}} \rangle &= \sum_{j=0}^{L-1} \langle W_{|y_j} \rangle \\ &= \sum_{j=0}^{L-1} k_B T \left(\ln 2 + \ln \theta_{A \rightarrow A}^{(y_j)} \right) \\ &= k_B T \left(L \ln 2 + \ln \prod_{j=0}^{L-1} \theta_{A \rightarrow A}^{(y_j)} \right).\end{aligned}$$

Due to the single causal state, the product within the logarithm simplifies to the probability of the word given the model $\prod_{j=0}^{L-1} \theta_{A \rightarrow A}^{(y_j)} = \Pr(Y_{0:L}^\theta = y_{0:L})$. So, the resulting work production depends on the familiar log-likelihood:

$$\begin{aligned}\langle W_{|y_{0:L}} \rangle &= k_B T (L \ln 2 + \ell(\theta|y_{0:L})) \\ &= \langle W_{|y_{0:L}}^\theta \rangle,\end{aligned}$$

again, achieving efficient work production, as expected.

H.3. Maximizing work for memoryless models

Leveraging the explicit construction for efficient information ratchets, we can search for the agent that maximizes work from the input string $y_{0:L}$. To infer a model through work maximization, we label the frequency of \uparrow states in this sequence with $f(\uparrow)$ and the frequency of \downarrow with $f(\downarrow)$. The corresponding log-likelihood of the model is:

$$\begin{aligned}\ell(\theta|y_{0:L}) &= \ln \left(\theta_{A \rightarrow A}^{(\uparrow)} \right)^{Lf(\uparrow)} \left(\theta_{A \rightarrow A}^{(\downarrow)} \right)^{Lf(\downarrow)} \\ &= Lf(\uparrow) \ln \left(\theta_{A \rightarrow A}^{(\uparrow)} \right) + Lf(\downarrow) \ln \left(\theta_{A \rightarrow A}^{(\downarrow)} \right).\end{aligned}$$

Thus, for the corresponding agent, the work production is:

$$\begin{aligned}\langle W_{|y_{0:L}}^\theta \rangle &= k_B T \ell(\theta|y_{0:L}) + k_B T L \ln 2 \\ &= k_B T L \left(\ln 2 + f(\uparrow) \ln \theta_{A \rightarrow A}^{(\uparrow)} + f(\downarrow) \ln \theta_{A \rightarrow A}^{(\downarrow)} \right).\end{aligned}$$

Selecting from all possible memoryless agents, the model parameters θ maximizing work production are given by the frequency of symbols in the input: $f(\uparrow) = \theta_{A \rightarrow A}^{(\uparrow)}$ and $f(\downarrow) = \theta_{A \rightarrow A}^{(\downarrow)}$. The resulting work production is:

$$\langle W_{|y_{0:L}}^\theta \rangle = k_B T L (\ln 2 - H[f(\uparrow)]),$$

where $H[f(\uparrow)]$ is the Shannon entropy of binary variable Y with $\Pr(Y = \uparrow) = f(\uparrow)$ measured in nats.

This simple example of learning statistical bias serves to explicitly lay out the stages of thermodynamic machine learning. The class of models is too simple, though, to illustrate the full power of the new learning method. That said, it does confirm that thermodynamic work maximization leads to useful models of data in the simplest case. As one would expect, the simple agent found by thermodynamic machine learning discovers the frequency of zeros in the input and, thus, it learns about its environment. The corresponding work production is the same as energetic gain of randomizing L bits distributed according to the frequency $f(\uparrow)$.

However, this neglects the substantial thermodynamic benefits possible with temporally-correlated environments [38]. To illustrate how to extract this additional energy, a sequel designs and analyzes memoryful agents.

ORCID iDs

Alexander B Boyd  <https://orcid.org/0000-0003-1092-922X>

James P Crutchfield  <https://orcid.org/0000-0003-4466-5410>

Mile Gu  <https://orcid.org/0000-0002-5459-4313>

References

- [1] Cuvier G 1818 *Essay on the Theory of the Earth* (New York: Kirk and Mercein)
- [2] Bergson H 1907 *Creative Evolution* (New York: Henry Holt and Company)
- [3] Thompson D W 1917 *On Growth and Form* (Cambridge: Cambridge University Press)
- [4] Wiener N 1948 *Cybernetics: Or Control and Communication in the Animal and the Machine* (Cambridge, MA: MIT Press)
- [5] Wiener N 1988 *The Human Use of Human Beings: Cybernetics and Society* (Cambridge: Da Capo Press)
- [6] Dennett D C 1991 *Consciousness Explained* (New York: Little, Brown and Co.)
- [7] Gould S J and Lewontin R 1979 The spandrels of san marco and the panglossian paradigm: a critique of the adaptationist programme *Proc. R. Soc. B* **205** 581–98
- [8] Dennett D C 1995 *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon and Schuster)
- [9] Maynard-Smith J and Szathmari E 1998 *The Major Transitions in Evolution* reprint edn (Oxford: Oxford University Press)
- [10] Wagner G P 2014 *Homology, Genes, and Evolutionary Innovation* (Princeton, NJ: Princeton University Press)
- [11] Thomson W 1874 Kinetic theory of the dissipation of energy *Nature* **9** 441–4
- [12] Maxwell J C 1888 *Theory of Heat* 9th edn (London: Longmans Green)
- [13] Sagawa T 2012 Thermodynamics of information processing in small systems *Prog. Theor. Phys.* **127** 1–56
- [14] Parrondo J M R, Horowitz J M and Sagawa T 2015 Physics of information *Nat. Phys.* **11** 131
- [15] Shalev-Shwartz S and Ben-David S 2014 *Understanding Machine Learning: From Theory to Algorithms* (Cambridge: Cambridge University Press)
- [16] Hastie T, Tibshirani R and Friedman J 1818 *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (*Springer Series in Statistics*) 2nd edn (Berlin: Springer)
- [17] Mehta P, Bukov M, Wang C-H, Day A G R, Richardson C, Fisher C K and Schwab D J 2019 A high-bias, low-variance introduction to machine learning for physicists *Phys. Rep.* **810** 1–124
- [18] Lin H W, Tegmark M and Rolnick D 2017 Why does deep and cheap learning work so well? *J. Stat. Phys.* **168** 1223–47
- [19] Landauer R 1961 Irreversibility and heat generation in the computing process *IBM J. Res. Dev.* **5** 183–91
- [20] Bennett C H 1987 Engines and the second law *Sci. Am.* **257** 108–16
- [21] Szilard L 1929 On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings *Z. Phys.* **53** 840–56
- [22] Watkin T L H, Rau A and Biehl M 1993 The statistical mechanics of learning a rule *Rev. Mod. Phys.* **65** 499–556
- [23] Engel A and Van den Broeck C 2001 *Statistical Mechanics of Learning* (Cambridge: Cambridge University Press)
- [24] Bell A 2011 Learning out of equilibrium *Santa Fe Institute Workshop on Randomness, Structure, and Causality* (9–13 January 2011) ed J P Crutchfield and J Machta vol 21
- [25] Sohl-Dickstein J, Weiss E A, Maheswaranathan N and Ganguli S 2015 Deep unsupervised learning using nonequilibrium thermodynamics (arXiv:1503.03585)
- [26] Goldt S and Seifert U 2017 Stochastic thermodynamics of learning *Phys. Rev. Lett.* **118** 010601
- [27] Alemi A A and Fischer I 2018 TherML: thermodynamics of machine learning (arXiv:1807.04162)
- [28] Bahri Y, Kadmon J, Pennington J, Schoenholz S S, Sohl-Dickstein J and Ganguli S 2020 Statistical mechanics of deep learning *Annu. Rev. Condens. Matter Phys.* **11** 501–28
- [29] Mandal D and Jarzynski C 2012 Work and information processing in a solvable model of Maxwell's demon *Proc. Natl Acad. Sci.* **109** 11641–5
- [30] Boyd A B, Mandal D and Crutchfield J P 2016 Identifying functional thermodynamics in autonomous Maxwellian ratchets *New J. Phys.* **18** 023049
- [31] Boyd A B, Mandal D and Crutchfield J P 2016 Leveraging environmental correlations: the thermodynamics of requisite variety *J. Stat. Phys.* **167** 1555–85
- [32] Gold J M and England J L 2019 Self-organized novelty detection in driven spin glasses (arXiv:1911.07216)
- [33] Zhong W, Gold J M, Marzen S, England J L and Halpern N Y 2020 Learning about learning by many-body systems (arXiv:2004.03604 [cond-mat.stat-mech])
- [34] Jimenez Rezende D, Mohamed S and Wierstra D 2014 Stochastic backpropagation and approximate inference in deep generative models (arXiv:1401.4082)
- [35] Dempster A P, Laird N M and Rubin D B 1977 Maximum likelihood from incomplete data via the EM algorithm *J. R. Stat. Soc. B* **39** 1–22
- [36] Crutchfield J P 2012 Between order and chaos *Nat. Phys.* **8** 17–24
- [37] Deffner S and Jarzynski C 2013 Information processing and the second law of thermodynamics: an inclusive, Hamiltonian approach *Phys. Rev. X* **3** 041003
- [38] Boyd A B, Mandal D and Crutchfield J P 2017 Correlation-powered information engines and the thermodynamics of self-correction *Phys. Rev. E* **95** 012152
- [39] Boyd A B, Mandal D, Riechers P M and Crutchfield J P 2017 Transient dissipation and structural costs of physical information transduction *Phys. Rev. Lett.* **118** 220602
- [40] Merhav N 2015 Sequence complexity and work extraction *J. Stat. Mech.* P06037
- [41] Merhav N 2017 Relations between work and entropy production for general information-driven, finite-state engines *J. Stat. Mech.* 1–20
- [42] Garner A J P, Thompson J, Vedral V and Gu M 2017 Thermodynamics of complexity and pattern manipulation *Phys. Rev. E* **95** 042140
- [43] Brookshear J G 1989 *Theory of Computation: Formal Languages, Automata, and Complexity* (New York: Benjamin-Cummings)
- [44] Barnett N and Crutchfield J P 2015 Computational mechanics of input–output processes: structured transformations and the ϵ -transducer *J. Stat. Phys.* **161** 404–51
- [45] Boyd A B, Mandal D and Crutchfield J P 2018 Thermodynamics of modularity: structural costs beyond the Landauer bound *Phys. Rev. X* **8** 031036
- [46] Touzo L, Marsili M, Merhav N and Roldán É 2020 Optimal work extraction and the minimum description length principle *J. Stat. Mech.* 093403
- [47] Kirchhoff G R 1848 *Ann. Phys.* **75** 1891
- [48] Gibbs J W 1906 *The Scientific Papers of J. Willard Gibbs* (London: Longmans Green)
- [49] Maxwell J C 1954 *A Treatise on Electricity and Magnetism* 3rd edn vol 1–2 (New York: Dover)
- [50] Onsager L 1931 Reciprocal relations in irreversible processes. I *Phys. Rev.* **37** 405–26

- [51] Prigogine I 1945 Modération et transformations irréversibles des systèmes ouverts *Bull. Class. Sci. Acad. R. Belg.* **31** 600–6
- [52] Prigogine I and Glansdorff P 1971 *Thermodynamic Theory of Structure, Stability and Fluctuations* (New York: Interscience (Wiley-Interscience))
- [53] Falasco G, Rao R and Esposito M 2018 Information thermodynamics of Turing patterns *Phys. Rev. Lett.* **121** 108301
- [54] Turing A M 1952 The chemical basis of morphogenesis *Trans. R. Soc. B* **237** 37
- [55] Hoyle R 2006 *Pattern Formation: An Introduction to Methods* (Cambridge: Cambridge University Press)
- [56] Cross M and Greenside H 2009 *Pattern Formation and Dynamics in Nonequilibrium Systems* (Cambridge: Cambridge University Press)
- [57] Heisenberg W 1967 Nonlinear problems in physics *Phys. Today* **20** 23–33
- [58] Ruelle D and Takens F 1971 On the nature of turbulence *Commun. Math. Phys.* **20** 167–92
- [59] Brandstätter A, Swift J, Swinney H L, Wolf A, Farmer J D, Jen E and Crutchfield J P 1983 Low-dimensional chaos in a hydrodynamic system *Phys. Rev. Lett.* **51** 1442
- [60] Shannon C E 1948 A mathematical theory of communication *Bell Syst. Tech. J.* **27** 623–56
- [61] Cover T M and Thomas J A 2006 *Elements of Information Theory* 2nd edn (New York: Interscience (Wiley-Interscience))
- [62] James R G, Ellison C J and Crutchfield J P 2011 Anatomy of a bit: information in a time series observation *CHAOS* **21** 037109
- [63] Shannon C 1956 The bandwagon *IRE Trans. Inf. Theory* **2** 3
- [64] Turing A 1937 On computable numbers, with an application to the Entscheidungsproblem *Proc. Math. Soc.* **42** 544–6
- [65] Turing A 1937 On computable numbers, with an application to the Entscheidungsproblem *Proc. Math. Soc.* **43** 230–65
- [66] Ross Ashby W 1960 *An Introduction to Cybernetics* 2nd edn (New York: Wiley)
- [67] Jaynes E T 1957 Information theory and statistical mechanics *Phys. Rev.* **106** 620–30
- [68] Jaynes E T 1980 The minimum entropy production principle *Annu. Rev. Phys. Chem.* **31** 579–601
- [69] Kolmogorov A N 1958 A new metric invariant of transient dynamical systems and automorphisms in Lebesgue spaces *Dokl. Akad. Nauk SSSR* **119** 861 (Russian) *Math. Rev.* vol 21
- [70] Packard N H, Crutchfield J P, Farmer J D and Shaw R S 1980 Geometry from a time series *Phys. Rev. Lett.* **45** 712
- [71] Takens F 1981 Detecting strange attractors in fluid turbulence *Symp. Dynamical Systems and Turbulence* ed D A Rand and L S Young vol 898 (Berlin: Springer) p 366
- [72] Crutchfield J P and McNamara B S 1987 Equations of motion from a data series *Complex Syst.* **1** 417–52
- [73] Brillouin L 1962 *Science and Information Theory* 2nd edn (New York: Academic)
- [74] Bennett C H 1982 Thermodynamics of computation—a review *Int. J. Theor. Phys.* **21** 905
- [75] Sagawa T and Ueda M 2013 Information thermodynamics: Maxwell’s demon in nonequilibrium dynamics *Nonequilibrium Stat. Phys. Small Syst. Fluct. Relations Beyond* **181–211**
- [76] Still S, Sivak D A, Bell A J and Crooks G E 2012 Thermodynamics of prediction *Phys. Rev. Lett.* **109** 120604
- [77] Crutchfield J P and Görnerup O 2006 Objects that make objects: the population dynamics of structural complexity *J. R. Soc. Interface.* **3** 345–9
- [78] England J L 2013 Statistical physics of self-replication *J. Chem. Phys.* **139** 121923
- [79] Serreli V, Lee C-F, Kay E R and Leigh D A 2007 A molecular information ratchet *Nature* **445** 523–7
- [80] Thompson J, Garner A J P, Vedral V and Gu M 2017 Using quantum theory to simplify input–output processes *njp Quantum Inf.* **3** 1–8
- [81] Loomis S and Crutchfield J P 2020 Thermal efficiency of quantum memory compression *Phys. Rev. Lett.* **125** 020601
- [82] Woods M P, Silva R, Putz G, Stupar S and Renner R 2018 Quantum clocks are more accurate than classical ones (arXiv:1806.00491)
- [83] Dayan P, Hinton G E, Neal R M and Zemel R S 1995 The Helmholtz machine *Neural Comput.* **7** 889–904
- [84] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [85] Landauer R 1991 Information is physical *Phys. Today* **44** 23–9
- [86] Esposito M and van den Broeck C 2011 Second law and Landauer principle far from equilibrium *Europhys. Lett.* **95** 40004
- [87] Deffer S and Lutz E 2012 Information free energy for nonequilibrium states (arXiv:1201.3888)
- [88] Crutchfield J P and Feldman D P 2003 Regularities unseen, randomness observed: levels of entropy convergence *CHAOS* **13** 25–54
- [89] Strelhoff C C and Crutchfield J P 2014 Bayesian structural inference for hidden processes *Phys. Rev. E* **89** 042119
- [90] Wu T and Tegmark M 2019 Toward an artificial intelligence physicist for unsupervised learning *Phys. Rev. E* **100** 033311
- [91] Crutchfield J P 1994 The calculi of emergence: computation, dynamics, and induction *Physica D* **75** 11–54
- [92] Crooks G E 1999 Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences *Phys. Rev. E* **60** 2721
- [93] Jarzynski C 2000 Hamiltonian derivation of a detailed fluctuation theorem *J. Stat. Phys.* **98** 77–102
- [94] Speck T and Seifert U 2004 Distribution of work in isothermal nonequilibrium processes *Phys. Rev. E* **70** 066112
- [95] Ray K J, Wimsatt G W, Boyd A B and Crutchfield J P 2020 Non-Markovian momentum computing: universal and efficient (arXiv:2010.01152)
- [96] Kolchinsky A and Wolpert D H 2017 Dependence of dissipation on the initial distribution over states *J. Stat. Mech.* **083202**
- [97] Riechers P M and Gu M 2020 Initial-state dependence of thermodynamic dissipation for any quantum process (arXiv:2002.11425)
- [98] Jarzynski C 2011 Equalities and inequalities: irreversibility and the second law of thermodynamics at the nanoscale *Annu. Rev. Condens. Matter Phys.* **2.1** 329–51
- [99] Owen J A, Kolchinsky A and Wolpert D H 2019 Number of hidden states needed to physically implement a given conditional distribution *New J. Phys.* **21** 013022