

The thermodynamics of prediction

Susanne Still,¹ David A. Sivak,² Anthony J. Bell,³ and Gavin E. Crooks²

¹*University of Hawaii at Mānoa, Information and Computer Sciences, sstill@hawaii.edu*

²*Physical Biosciences Division, Lawrence Berkeley National Laboratory*

³*Redwood Center for Theoretical Neuroscience, University of California at Berkeley*

A system responding to a stochastic driving signal can be interpreted as computing, by means of its dynamics, an implicit model of the environmental variables. The system's state retains information about past environmental fluctuations, and a fraction of this information is predictive of future ones. The remaining nonpredictive information reflects model complexity that does not improve predictive power, and thus represents the ineffectiveness of the model. We expose the fundamental equivalence between this model inefficiency and thermodynamic inefficiency, measured by the energy dissipated during the interaction between system and environment. Our results hold arbitrarily far from thermodynamic equilibrium and are applicable to a wide range of systems, including biomolecular machines. They highlight a profound connection between the effective use of information and efficient thermodynamic operation: any system constructed to keep memory about its environment *and* to operate with maximal energetic efficiency *has to be predictive*.

All systems perform computations by means of responding to their environment. In particular, living systems compute, on a variety of length- and time-scales, future expectations based on their prior experience. Most biological computation is fundamentally a nonequilibrium process, because a preponderance of biological machinery in its natural operation is driven far from thermodynamic equilibrium. For example, many molecular machines (such as the microtubule-associated motor kinesin) are driven by ATP hydrolysis, which liberates ~ 500 meV per molecule [32]. This energy is large compared with ambient thermal energy, $1 k_B T \approx 25$ meV (k_B is Boltzmann's constant and the temperature is $T \sim 300$ Kelvin). In general, such large energetic inputs drive the operative degrees of freedom of biological machines away from equilibrium averages.

Recently, significant progress has been made in describing driven systems far from equilibrium [2], per-

haps most notably Jarzynski's work relation [33] generalizing Clausius' Inequality; the further generalization embodied in fluctuation theorems [34, 35]; and the extension of these relations to calculating potentials of mean force [36]. These advances have allowed researchers to measure equilibrium quantities, such as free energy changes, by observing how a system reacts to being driven out of equilibrium, e.g. [39, 40].

This literature typically assumes that the experiment is known, i.e. that the exact time course of the driving signal is given. However, systems that are embedded in realistic environments, for example, a biological macromolecule operating under natural conditions, are exposed to a stochastic drive. Here, we therefore study driven systems for which the changes in the driving signal(s) are governed by some probability density P_X . This can be any stochastic process, and the results we derive require neither that P_X has specific properties, nor that it is known by the system. The dissipation, averaged not only over the system's path through its state space, but also over driving protocols, then quantifies the system's energetic inefficiency. We assume that there is no feedback from the system to the driving signal.

The system dynamics perform a computation by changing the system's state, as a function of the driving signal. As a result, the new system state contains some memory about the driving signal. The dynamics of the system can be interpreted as computing a model: past environmental influences are mapped onto the current state of the system, which through its correlation with forthcoming environmental fluctuations implicitly contains a prediction of the future.


In this paper, we ask how the quality of this (implicit) model is related to thermodynamic inefficiency. But how do we measure the quality of a model? A useful model has to have predictive power (see e.g. [41–44], and refs.

DISCLAIMER: This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

therein), meaning it must capture predictive information [45–48], while not being overly complicated. In other words, the model should contain as little dispensable *nonpredictive* information as possible.

Our central contribution is the demonstration of a fundamental equivalence between the instantaneous nonpredictive information carried by the system and the dissipation of energy.

Problem setup.— Let s_t denote the state of the system at time t , while x_t denotes the driving signal. The dynamics of the system are modeled by discrete time Markovian conditional state-to-state transition probabilities, $p(s_t|s_{t-1}, x_t)$. The external drive is governed by $P_X = p(x_0, \dots, x_\tau)$. We assume that at time $t = 0$, the system is in thermodynamic equilibrium, in contact with a heat bath with inverse temperature $\beta := 1/k_B T$. A change in the external driving signal $x_0 \rightarrow x_1$ forces the system out of equilibrium. The system responds by changing its state $s_0 \rightarrow s_1$, according to the transition probability $p(s_1|s_0, x_1)$. The external signal subsequently changes again $x_1 \rightarrow x_2$, and the process is repeated until time $t = \tau$. The system remains in thermal contact with the heat bath during the entire protocol $x_0, \dots, x_\tau$, as in [49].



Work is done during a ‘work step’, as the external signal changes from x_{t-1} to x_t [49, 53],

$$W[s_{t-1}; x_{t-1} \rightarrow x_t] := E(s_{t-1}, x_t) - E(s_{t-1}, x_{t-1}) . \quad (1)$$

In response to this change, the system relaxes from s_{t-1} to s_t in a ‘relaxation step.’ The total work over the course of a driving protocol is $W = \sum_{t=1}^{\tau} W[s_{t-1}; x_{t-1} \rightarrow x_t]$. The total change in energy, $\Delta E := E(s_\tau, x_\tau) - E(s_0, x_0) = W + Q$, equals the total work plus the total heat, $Q = \sum_{t=1}^{\tau} [E(s_t, x_t) - E(s_{t-1}, x_t)]$, flowing into the system during the relaxation steps.

For now, we assume that the kernel which describes the dynamics, $p(s_t|s_{t-1}, x_t)$, is fixed. However, the conditional distribution over states after the work step but before the system relaxes, $p(s_{t-1}|x_t)$, changes as a function of time, as does the conditional distribution over states after the relaxation step, $p(s_t|x_t)$. In general, these distributions are not the same, and neither of them is an equilibrium distribution. Under Markovian system dynamics, the probability before relaxation simplifies to

$$p(s_{t-1}|x_t) = \left\langle \left\langle \left\langle \dots \left\langle p(s_{t-1}|s_{t-2}, x_{t-1}) \right\rangle_{p(s_{t-2}|s_{t-3}, x_{t-2})} \dots \right\rangle_{p(s_1|s_0, x_0)} \right\rangle_{p(s_0|x_0)} \right\rangle_{p(x_0, \dots, x_{t-1}|x_t)} , \quad (2)$$

and the distribution after relaxation is given by

$$p(s_t|x_t) = \langle p(s_t|s_{t-1}, x_t) \rangle_{p(s_{t-1}|x_t)} . \quad (3)$$

Angled brackets with a subscripted probability p denote an average over p .

The equilibrium distribution is the same function, before and after relaxation, $p_{\text{eq}}(s|x_t) := e^{-\beta(E(s, x_t) - F_t)}$, where s refers to the state of the system with energy $E(s, x_t)$, and $F_t := F[x_t]$ denotes equilibrium free energy. The probability of a specific path through the system’s state space, given the protocol, is

$$P_{S|X} = p_{\text{eq}}(s_0|x_0) \prod_{t=1}^{\tau} p(s_t|s_{t-1}, x_t) , \quad (4)$$

and the joint probability is

$$\begin{aligned} P_{S,X} &:= p(s_0, \dots, s_\tau, x_0, \dots, x_\tau) \\ &= p(x_0) p_{\text{eq}}(s_0|x_0) \prod_{t=1}^{\tau} p(x_t|x_0, \dots, x_{t-1}) p(s_t|s_{t-1}, x_t) . \end{aligned} \quad (5)$$

In the following, unless otherwise clear from the context, angled brackets without a subscript denote an average over the distribution $P_{S,X}$.

Dissipation out of equilibrium.— After the conclusion of the protocol, the probability over system states is given by $p(s_\tau|x_\tau)$, in general not an equilibrium distribution. Then, in addition to the equilibrium free energy, F_τ , there is another contribution to the free energy, because the system is not in thermodynamic equilibrium. This additional free energy would be dissipated as heat to the environment if the system were to relax to thermodynamic equilibrium. For markovian system dynam-

ics, the additional nonequilibrium contribution is proportional [45] to the relative entropy (Kullback-Leibler divergence) between the actual distribution $p(s_\tau|x_\tau)$ at the end of the protocol and the equilibrium distribution,

$$F_\tau^{\text{add}}[p(s_\tau|x_\tau)] = k_B T D_{\text{KL}}[p(s_\tau|x_\tau)||p_{\text{eq}}(s_\tau|x_\tau)] . \quad (6)$$

The non-negative Kullback-Leibler (KL) divergence between distributions $p(x)$ and $q(x)$ is defined as

$$D_{\text{KL}}[p(x)||q(x)] := \left\langle \ln \left[\frac{p(x)}{q(x)} \right] \right\rangle_{p(x)} \geq 0 . \quad (7)$$

The sum of both contributions to the free energy constitutes the overall *nonequilibrium* free energy, $F_{\text{neq}}[p(s_\tau|x_\tau)] = F_\tau + F_\tau^{\text{add}}[p(s_\tau|x_\tau)]$. Here, nonequilibrium free energy is defined in analogy to the standard equilibrium free energy as a functional of the probability distribution, but applied to any probability distribution [50–52], that is to any $p(s|x)$:

$$F_{\text{neq}}[p(s|x)] := \langle E(s, x) \rangle_{p(s|x)} + k_B T \langle \ln[p(s|x)] \rangle_{p(s|x)} . \quad (8)$$

The average work *irretrievably* lost over the course of a driving protocol,

$$\langle W_{\text{diss}} \rangle_{P_{S|X}} := \langle W \rangle_{P_{S|X}} - \Delta F_{\text{neq}} , \quad (9)$$

equals the average work performed on the system minus the nonequilibrium free energy change $\Delta F_{\text{neq}} := F_{\text{neq}}[p(s_\tau|x_\tau)] - F_{\text{neq}}[p(s_0|x_0)]$. We can compare this to the average *excess work* for a given protocol, $\langle W_{\text{ex}} \rangle_{P_{S|X}} := \langle W \rangle_{P_{S|X}} - \Delta F$, the total work done on the system in excess of the equilibrium free energy change $\Delta F := F_\tau - F_0$ which would be the work done if the driving signal changed quasistatically (infinitely slowly), and hence the system remained in thermodynamic equilibrium throughout the protocol. This excess work equals the dissipated work only when the protocol includes a final equilibration of the system.

Since the system starts in equilibrium, the total change in nonequilibrium free energy is the equilibrium free energy change plus the abovementioned additional contribution to the free energy, $\Delta F_{\text{neq}} = \Delta F + F_\tau^{\text{add}}[p(s_\tau|x_\tau)]$. The dissipation is then the excess work minus this additional nonequilibrium contribution

$$\langle W_{\text{diss}} \rangle_{P_{S|X}} = \langle W_{\text{ex}} \rangle_{P_{S|X}} - F_\tau^{\text{add}}[p(s_\tau|x_\tau)] \leq \langle W_{\text{ex}} \rangle_{P_{S|X}} . \quad (10)$$

Later, we derive a lower bound on the dissipation and excess work averaged over all protocols, denoted by $\langle W_{\text{diss}} \rangle$, and $\langle W_{\text{ex}} \rangle$, respectively.

Each of the incremental work steps, $x_t \rightarrow x_{t+1}$, is accompanied by a nonequilibrium free energy change given by $\Delta F_{\text{neq}}[x_t \rightarrow x_{t+1}] := F_{\text{neq}}[p(s_t|x_{t+1})] - F_{\text{neq}}[p(s_t|x_t)]$, so that the average dissipation during each work step is

$$\begin{aligned} \langle W_{\text{diss}}[x_t \rightarrow x_{t+1}] \rangle &:= \langle W[s_t; x_t \rightarrow x_{t+1}] \rangle_{p(s_t, x_t, x_{t+1})} \\ &\quad - \langle \Delta F_{\text{neq}}[x_t \rightarrow x_{t+1}] \rangle_{p(x_t, x_{t+1})} . \end{aligned} \quad (11)$$

The nonequilibrium free energy change during each relaxation step is

$$\Delta F_{\text{neq}}[s_{t-1} \rightarrow s_t] = F_t^{\text{add}}[p(s_t|x_t)] - F_{t+1}^{\text{add}}[p(s_{t-1}|x_t)] , \quad (12)$$

which equals the change in KL divergence from the equilibrium distribution.

Predictive power, memory, and dissipation.— The system state and the external signal are random variables that potentially share information, $I[s_t, x_t] := \langle \ln[p(s_t, x_t)/p(s_t)p(x_t)] \rangle_{p(s_t, x_t)}$ [54, 55], where $p(s_t) = \langle p(s_t|x_t) \rangle_{p(x_t)}$. Mutual information measures the reduction in uncertainty about the outcome of a random variable upon learning the identity of another variable, and it is symmetric: $I[s_t, x_t] = H[s_t] - H[s_t|x_t] = H[x_t] - H[x_t|s_t]$. Uncertainty is quantified by the entropy, $H[x_t] := -\langle \ln p(x_t) \rangle_{p(x_t)}$, and the conditional entropy, $H[x_t|s_t] := -\langle \ln p(x_t|s_t) \rangle_{p(x_t, s_t)}$, respectively.

The system transition probability, $p(s_t|s_{t-1}, x_t)$, is assumed to depend on the current signal value x_t and system state s_{t-1} . These two dependencies are sufficient to induce correlations between the system's current state and previous signal values. The memory the system keeps about the external signal can then be quantified by the information that the system state s_t retains about a trajectory $\{x_{t-\tau_m}, \dots, x_t\}$. In general, there are temporal correlations in the input signal, and hence, there can be correlations between s_t and future signal values. That is, some of the memory retained in the system's state is information about the future trajectory $\{x_{t+1}, \dots, x_{t+\tau_f}\}$. Here we focus on the *instantaneous* memory, $I_{\text{mem}}(t) := I[s_t, x_t]$, and the *instantaneous* predictive power [43], $I_{\text{pred}}(t) := I[s_t, x_{t+1}] = H[x_{t+1}] - H[x_{t+1}|s_t]$.

The implicit model computed by the system's dynamics consists of the probabilistic map $p(s_t|x_t)$, which maps signal x_t onto state s_t , and $p(x_{t+1}|s_t) = p(s_t|x_{t+1})p(x_{t+1})/p(s_t)$, which represents the prediction of x_{t+1} , given the value of s_t .

The instantaneous *nonpredictive information* is defined as the difference between instantaneous memory and predictive power, $I_{\text{mem}}(t) - I_{\text{pred}}(t)$. It represents useless nostalgia and provides a measure for the ineffectiveness of the model.

Averaging the nonequilibrium free energy over protocols allows us to write

$$\beta \langle F_{\text{neq}}[p(s|x)] \rangle_{p(x)} = \beta \langle E(s, x) \rangle_{p(s, x)} - H[s|x] . \quad (13)$$

Combining this with Eqs. (1, 11) [24], we arrive at our first result: the instantaneous nonpredictive information is proportional to the average work dissipated while the signal changes from x_t to x_{t+1} ,

$$\beta \langle W_{\text{diss}}[x_t \rightarrow x_{t+1}] \rangle = I_{\text{mem}}(t) - I_{\text{pred}}(t) . \quad (14)$$

In summary, the unwarranted retention of past information is fundamentally equivalent to energetic inefficiency.

Lower bound on total dissipation.— We now relate the total average dissipated work during the entire protocol, averaged over all protocols, $\langle W_{\text{diss}} \rangle$, to the total nostalgia, $I_{\text{mem}} - I_{\text{pred}}$, given by the difference between the total instantaneous memory, $I_{\text{mem}} := \sum_{t=0}^{\tau-1} I_{\text{mem}}(t)$, and the total instantaneous predictive power, $I_{\text{pred}} := \sum_{t=0}^{\tau-1} I_{\text{pred}}(t)$. To that end we need to sum Eq. (11) over all time steps. This sum includes a sum over changes in nonequilibrium free energy, which can be expressed as

$$\left\langle \sum_{t=0}^{\tau-1} \Delta F_{\text{neq}}[x_t \rightarrow x_{t+1}] \right\rangle = \langle \Delta F_{\text{neq}} - \Delta F_{\text{neq}}^{\text{relax}} \rangle \quad (15)$$

in terms of the total nonequilibrium free energy change, $\langle \Delta F_{\text{neq}} \rangle$, and the sum of nonequilibrium free energy changes during relaxation steps:

$$\langle \Delta F_{\text{neq}}^{\text{relax}} \rangle := \left\langle \sum_{t=0}^{\tau-1} \Delta F_{\text{neq}}[s_t \rightarrow s_{t+1}] \right\rangle \leq 0. \quad (16)$$

The total dissipation then becomes (using Eq. (14))

$$\beta \langle W_{\text{diss}} \rangle = I_{\text{mem}} - I_{\text{pred}} - \beta \langle \Delta F_{\text{neq}}^{\text{relax}} \rangle. \quad (17)$$

On average, during relaxation steps the system evolves toward equilibrium. The total nostalgia therefore provides a lower bound on the total average dissipation, and also, due to Eq. (10), on the total average excess work,

$$I_{\text{mem}} - I_{\text{pred}} \leq \beta \langle W_{\text{diss}} \rangle \leq \beta \langle W_{\text{ex}} \rangle. \quad (18)$$

We can use this result to refine Landauer's principle [56], which states that any erasure of information must be balanced by an increase in entropy elsewhere. The information erased during a protocol, such as the reset protocol of Landauer [56], is the entropy change $\mathcal{I}_e := H[s_0|x_0] - H[s_\tau|x_\tau]$. Note that the information erased here is not mutual information about the driving signal, but rather information that could have potentially been extracted from the system by some measurement process. Landauer pointed out that the erasure of information requires heat to flow out of the system, which [using the First Law and Eqs. (8, 9)] obeys

$$-\beta \langle Q \rangle = \mathcal{I}_e + \beta \langle W_{\text{diss}} \rangle \geq \mathcal{I}_e \quad (19)$$

Substituting our result, Eq. (14), into Eq. (19) yields the new relation

$$-\beta \langle Q \rangle = \mathcal{I}_e + I_{\text{mem}} - I_{\text{pred}} - \beta \langle \Delta F_{\text{neq}}^{\text{relax}} \rangle. \quad (20)$$

Thus [using Eq. (16)] we obtain a refinement of Landauer's principle,

$$-\beta \langle Q \rangle \geq \mathcal{I}_e + I_{\text{mem}} - I_{\text{pred}}, \quad (21)$$

where the bound is augmented by the total nostalgia. The system dynamics of a computing device that retains memory therefore must be maximally predictive to approach Landauer's limit.

Discussion.— The dynamics, $p(s_t|x_{t-1}, x_t)$, have been assumed fixed for any given system. However, biological systems are typically adapted to their environment. One can then ask if there is a simple principle underlying the process producing this adaptation. If such a principle exists, then it may reflect the importance of energetic efficiency, because of the resulting competitive advantage for reproducing organisms. While other criteria play a role, such as robustness and sensitivity, energetic efficiency is of fundamental relevance. This is exemplified by biological molecules that harness environmental fluctuations to accomplish energetically-costly downstream tasks. The more efficiently such a molecule can operate, the more it can accomplish through coupling to endergonic chemical reactions or mechanical actions. For example, with more efficient coupling to the environment, the molecular motor kinesin can carry larger cargos. Likewise, with greater efficiency cytochrome c oxidase complex, an enzyme that pumps protons across a membrane, can create stronger electrochemical gradients. Evidence for the importance of energetic efficiency is furthermore found in biomolecular machines that approach 100% efficiency when driven in a natural fashion: the stall torque for the F₁-ATPase [26] and the stall force for Myosin V [27] are near the maximal values possible given the free energy liberated by ATP hydrolysis and the sizes of their respective rotations and steps.

These and many other biological functions require some correspondence between the environment and the systems that implement them. Therefore the memory of their instantiating systems must be nonzero. We have shown that any such system with nonzero memory must conduct predictive inference, at least implicitly, to approach maximal energetic efficiency.

A substantial amount of work has sought to relate emerging biological functions and behaviors to efficient energy usage. Examples range from animal behavior (e.g. [28]) to single neurons, where recently researchers have proposed that the minimization of energy expenditure, subject to constraints given by the desired function, may be “a unifying principle governing neuronal biophysics” [29]. On the other hand, there is much research on optimal information processing in neurons. For a recent review see e.g. [30], which proposes that the extraction of predictive information in biological signal processing may constitute, or at least lead to, a general principle. By directly relating memory and predictive power to dissipation out of equilibrium, the results we have derived here indicate that these two important paradigms are deeply connected.

While it is perhaps intuitive that neurons and organisms should have to implement predictive inference to function well, our results have the striking implication that on all scales energetic efficiency calls for predictive inference. This includes the smallest biological units, such as molecular machines.

Our results also specify the required kind of predictive inference: maximization of predictive power at a desired level of system memory. This connects with work on optimal predictive inference algorithms discussed in [43, 44, 46, 57], and references therein. We envision implementing these algorithms in fast and efficient hardware. The results we have derived here could then be used to choose the energetically most efficient implementation among the many possible choices.

Conclusion.— We argued that *dissipation* far from thermodynamic equilibrium is given by average work minus *nonequilibrium* free energy change. We also argued that the nonpredictive part of a system’s memory provides a natural measure for the inefficiency of a system’s implicit model of its environment.

We showed that instantaneous nonpredictive information is proportional to the energy dissipated when an external driving signal changes by an incremental amount, thereby doing work on the system. This result demonstrates the intimate connection between prediction and energetic efficiency. Summed over the entire protocol, the total nonpredictive information provides a lower bound on the total dissipation.

These results imply that any system which is built to have nonzero memory *has to be predictive* in order to allow for minimal possible dissipation, i.e. *to operate at maximal energetic efficiency*. Our results furthermore allow for a refinement of Landauer’s principle, applied to systems away from thermodynamic equilibrium.

We have provided a connection between nonequilibrium thermodynamics and learning theory, by making precise how two important aspects of life are fundamentally related: making a predictive model of the environment and using available energy efficiently.

Acknowledgements.— SS thanks William Bialek and Robert S. Shaw for many inspiring discussions, and Matteo Marsili for helpful technical comments. SS and AJB are thankful for conversations with James P. Crutchfield and Michael R. DeWeese. DAS and GEC were funded by the Office of Basic Energy Sciences of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. AJB was funded by NSF grants SBE 0542013 and SMA 1041755 to the Temporal Dynamics of Learning Center, an NSF Science of Learning Center.

-
- [32] H. Lodish, A. Berk, S. Zipursky, P. Matsudaira, D. Baltimore, and J. Darnell, *Molecular Cell Biology* (W.H. Freeman, 2000).
- [2] C. Jarzynski, Annu. Rev. Condens. Matter Phys., **2**, 329 (2011).
- [33] C. Jarzynski, Phys. Rev. Lett., **78**, 2690 (1997).
- [34] D. J. Evans, E. G. D. Cohen, and G. P. Morriss, Phys. Rev. Lett., **71**, 2401 (1993).
- [35] G. E. Crooks, Phys. Rev. E, **60**, 2721 (1999).
- [36] G. Hummer and A. Szabo, P. Natl. Acad. Sci. USA, **98**, 3658 (2001).
- [39] J. T. Liphardt, S. Dumont, S. B. Smith, I. Tinoco, and C. Bustamante, Science, **296**, 1832 (2002).
- [40] D. Collin, F. Ritort, C. Jarzynski, S. B. Smith, I. Tinoco, and C. Bustamante, Nature, **437**, 231 (2005).
- [41] H. Jeffreys, *Theory of Probability*, 3rd ed. (Oxford University Press, 1998).
- [42] W. Bialek, in *Physics of bio-molecules and cells, Ecole de physique theorique Les Houches Session LXXV* (Springer-Verlag, 2001) pp. 485–577.
- [43] S. Still, Europhys. Lett., **85**, 28005 (2009).
- [44] S. Still, J. P. Crutchfield, and C. Ellison, Chaos, **20**, 037111 (2010).
- [45] R. Shaw, *The dripping faucet as a model chaotic system* (Aerial Press, 1984).
- [46] J. P. Crutchfield and K. Young, Phys. Rev. Lett., **63**, 105 (1989).
- [47] W. Bialek, I. Nemenman, and N. Tishby, Neural Comput., **13**, 2409 (2001).
- [48] J. P. Crutchfield and D. P. Feldman, Chaos, **13**, 25 (2003).
- [49] G. E. Crooks, J. Stat. Phys., **90**, 1481 (1998).
- [53] A. Imparato, L. Peliti, G. Pesce, G. Rusciano, and A. Sasso, Phys. Rev. E, **76**, 050101 (2007).
- [50] B. Gaveau and L. S. Schulman, Phys. Lett. A, **229**, 347 (1997).
- [51] G. E. Crooks, Phys. Rev. E, **75**, 041119 (2007).
- [52] D. A. Sivak and G. E. Crooks, Phys. Rev. Lett., **108**, 150601 (2012).
- [54] C. E. Shannon, Bell Syst. Tech. J., **27**, 379 (1948).
- [55] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley-Interscience, 2006).
- [24] $\beta \langle W_{\text{diss}}[x_t \rightarrow x_{t+1}] \rangle$
 $= \beta \left(\langle E(s_t, x_{t+1}) \rangle_{p(s_t, x_{t+1})} - \langle E(s_t, x_t) \rangle_{p(s_t, x_t)} \right)$
 $- \beta \left(\langle F_{\text{neq}}[p(s|x_{t+1})] \rangle_{p(x_{t+1})} + \langle F_{\text{neq}}[p(s|x_t)] \rangle_{p(x_t)} \right)$
 $= H[s_t|x_{t+1}] - H[s_t|x_t] = I_{\text{mem}}(t) - I_{\text{pred}}(t).$
- [56] R. Landauer, IBM J. Res. Develop., **5**, 183 (1961).
- [26] K. Kinoshita, R. Yasuda, H. Noji, and K. Adachi, Philos. T. Roy. Soc. B, **355**, 473 (2000).
- [27] G. Cappello, P. Pierobon, C. Symonds, L. Busoni, J. C. M. Gebhardt, M. Rief, and J. Prost, P. Natl. Acad. Sci. USA, **104**, 15328 (2007).
- [28] R. M. Alexander, *Locomotion of Animals* (Blackie, Glasgow and London, 1982); *Optima for Animals* (Edward Arnold, London, 1981).
- [29] A. Hasenstaub, S. Otte, E. Callaway, and T. J. Sejnowski, P. Natl. Acad. Sci. USA, **107**, 12329 (2010).
- [30] W. Bialek, “Biophysics: Searching for principles,” To be published by Princeton University Press.
- [57] S. Still and D. Precup, to appear in Theory in Biosciences, GSO-2009 Special Issue. (2011), available at <http://www2.hawaii.edu/~sstill/StillPrecup2011.pdf>.
- [32] H. Lodish, A. Berk, S. Zipursky, P. Matsudaira, D. Baltimore, and J. Darnell, *Molecular Cell Biology* (W.H. Freeman, 2000).
- [33] C. Jarzynski, Phys. Rev. Lett., **78**, 2690 (1997).
- [34] D. J. Evans, E. G. D. Cohen, and G. P. Morriss, Phys. Rev. Lett., **71**, 2401 (1993).
- [35] G. E. Crooks, Phys. Rev. E **60**, 2721 (1999).

- [36] G. Hummer and A. Szabo, P. Natl. Acad. Sci. USA **98**, 3658 (2001).
- [37] R. Kawai, J. M. R. Parrondo, and C. Van den Broeck, Phys. Rev. Lett. **98**, 080602 (2007).
- [38] C. Jarzynski, Eur. Phys. J. B **64**, 331 (2008).
- [39] J. T. Liphardt, S. Dumont, S. B. Smith, I. Tinoco, and C. Bustamante, Science **296**, 1832 (2002).
- [40] D. Collin, F. Ritort, C. Jarzynski, S. B. Smith, I. Tinoco, and C. Bustamante, Nature **437**, 231 (2005).
- [41] H. Jeffreys, *Theory of Probability* (Oxford University Press, 1998), 3rd ed.
- [42] W. Bialek, in *Physics of bio-molecules and cells, Ecole de physique theorique Les Houches Session LXXV* (Springer-Verlag, 2001), pp. 485–577.
- [43] S. Still, Europhys. Lett. **85**, 28005 (2009).
- [44] S. Still, J. P. Crutchfield, and C. Ellison, Chaos **20**, 037111 (2010).
- [45] R. Shaw, *The dripping faucet as a model chaotic system* (Aerial Press, 1984).
- [46] J. P. Crutchfield and K. Young, Phys. Rev. Lett. **63**, 105 (1989).
- [47] W. Bialek, I. Nemenman, and N. Tishby, Neural Comput. **13**, 2409 (2001).
- [48] J. P. Crutchfield and D. P. Feldman, Chaos **13**, 25 (2003).
- [49] G. E. Crooks, J. Stat. Phys. **90**, 1481 (1998).
- [50] B. Gaveau and L. S. Schulman, Phys. Lett. A **229**, 347 (1997).
- [51] G. E. Crooks, Phys. Rev. E **75**, 041119 (2007).
- [52] D. A. Sivak and G. E. Crooks, Phys. Rev. Lett. (2012), in press, arXiv:0912.2468.
- [53] A. Imparato, L. Peliti, G. Pesce, G. Rusciano, and A. Sasso, Phys. Rev. E **76**, 050101 (2007).
- [54] C. E. Shannon, Bell Syst. Tech. J. **27**, 379 (1948).
- [55] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley-Interscience, 2006), 2nd ed.
- [56] R. Landauer, IBM J. Res. Develop. **5**, 183 (1961).
- [57] S. Still and D. Precup, to appear in Theory in Biosciences, GSO-2009 Special Issue. (2011), available at <http://www2.hawaii.edu/~sstill/StillPrecup2011.pdf>.