

UC Davis

UC Davis Previously Published Works

Title

Prediction and Power in Molecular Sensors: Uncertainty and Dissipation When Conditionally Markovian Channels Are Driven by Semi-Markov Environments

Permalink

<https://escholarship.org/uc/item/5dj1n7jm>

Authors

Marzen, Sarah E
Crutchfield, James P

Publication Date

2017-07-12

Peer reviewed



Prediction and Dissipation in Nonequilibrium Molecular Sensors: Conditionally Markovian Channels Driven by Memoryful Environments

Sarah E. Marzen^{1,2} · James P. Crutchfield^{1,3}

Received: 5 December 2018 / Accepted: 31 December 2019
© Society for Mathematical Biology 2020

Abstract

Biological sensors must often predict their input while operating under metabolic constraints. However, determining whether or not a particular sensor is evolved or designed to be accurate and efficient is challenging. This arises partly from the functional constraints being at cross purposes and partly since quantifying the prediction performance of even *in silico* sensors can require prohibitively long simulations, especially when highly complex environments drive sensors out of equilibrium. To circumvent these difficulties, we develop new expressions for the prediction accuracy and thermodynamic costs of the broad class of conditionally Markovian sensors subject to complex, correlated (unifilar hidden semi-Markov) environmental inputs in nonequilibrium steady state. Predictive metrics include the instantaneous memory and the total predictable information (the mutual information between present sensor state and input future), while dissipation metrics include power extracted from the environment and the nonpredictive information rate. Success in deriving these formulae relies on identifying the environment's causal states, the input's minimal sufficient statistics for prediction. Using these formulae, we study large random channels and the simplest nontrivial biological sensor model—that of a Hill molecule, characterized by the number of ligands that bind simultaneously—the sensor's cooperativity. We find that the seemingly impoverished Hill molecule can capture an order of magnitude more predictable information than large random channels.

✉ James P. Crutchfield
chaos@ucdavis.edu

Sarah E. Marzen
smarzen@cmc.edu

- ¹ W. M. Keck Science Department, Claremont McKenna, Scripps and Pitzer Colleges, Claremont, CA 91711, USA
- ² Physics of Living Systems, Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
- ³ Complexity Sciences Center and Department of Physics, University of California at Davis, One Shields Avenue, Davis, CA 95616, USA

Keywords Predictive information rate · Information processing · Nonequilibrium steady state · Thermodynamics

1 Introduction

To perform functional tasks, synthetic nanoscale machines and their macromolecular cousins simultaneously manipulate energy, information, and matter. They are *information engines*—systems that operate by synergistically balancing the energetics of their physical substrate against required information generation, storage, loss, and transformation to support a given functionality. Classically, information engines were conceived as either potential computers (Bennett 1982)—that is, physical systems that can compute anything given the right program—or as Maxwellian-like demons that use information as a resource to convert disordered energy to useful work (Maxwell 1888; Szilard 1929; Mandal and Jarzynski 2012; Deffner and Jarzynski 2013; Boyd et al. 2017, 2016; Chapman and Miyake 2015; McGrath et al. 2017; Horowitz et al. 2013; Toyabe et al. 2010). Recently, investigations into functional computation¹ embedded in physical systems led to studies of the thermodynamics of various kinds of information processing (Parrondo et al. 2015), including the thermodynamic costs of information creation (Aghamohammadi and Crutchfield 2017), noise suppression (Hinczewski and Thirumalai 2014), error correction and synchronization (Boyd et al. 2017), prediction (Still et al. 2012; Barato et al. 2014; Horowitz and Esposito 2014), homeostasis (Boyd et al. 2016), learning (Goldt and Seifert 2017), structure (Boyd et al. 2017), and intelligent control (Boyd and Crutchfield 2016). Due to its broad importance to the survival of biological organisms, here we focus on a specific functional computation in information engines: How do sensors predict their environment?

Evolved and designed sensory systems can be tasked with at least two, potentially competing, objectives: accurately predicting inputs (Rao and Ballard 1999; Palmer et al. 2015) and contending with metabolic constraints (Chklovskii and Koulakov 2004; Hasenstaub et al. 2010).² Accurate and energy-efficient predictive feature extraction can be used to reap increased rewards from the environment, no matter the particular “reward function” (Sutton and Barto 1998; Littman et al. 2001; Brodus 2011; Little and Sommer 2014).

Optimizing such sensors requires a nontrivial matching of sensory statistics and sensor structure (Becker et al. 2015; Boyd et al. 2016). Undaunted, much effort has been invested to find energetically efficient and maximally predictive sensors, often

¹ Here, when analyzing sensory information processing in biological systems, we take care to distinguish intrinsic, functional, and useful computation (Crutchfield and Young 1989; Crutchfield 1994; Crutchfield and Mitchell 1995). *Intrinsic computation* refers to how a physical system stores and transforms its historical information. We take *functional computation* as information processing in a physical device that promotes the performance of a larger, encompassing system. We take *useful computation* as information processing in a physical device used to achieve an external user’s goal. The first is well suited to analyzing structure in physical processes and determining if they are candidate substrates for any kind of information processing. The second is well suited for discussing biological sensors, while the third is well suited for discussing the benefits of contemporary digital computers.

² We only consider online or real-time computations, so that the oft-considered energy–speed–accuracy trade-off (Lan et al. 2012; Lahiri et al. 2016) reduces to an energy–accuracy trade-off.

simplifying the challenges by ignoring action policies—how the sensed information is used. Some seek sensor models that maximize a combination of prediction power and (energetic) efficiency, e.g., as in Barato et al. (2014) and Becker et al. (2015). Others validate learning rules based on whether or not they maximize the aforementioned objective function (Creutzig and Sprekeler 2008; Creutzig et al. 2009). Finally, others compare real biological sensors to *in silico* null models, e.g., as in Palmer et al. (2015).

All these efforts require estimating prediction and dissipation metrics of given sensors. Doing so can be surprisingly difficult. Consider the total predictable information captured by a sensor (Bialek et al. 2001; Palmer et al. 2015), which we take to be the mutual information between the present sensor state and the future of the sensory input. There are uncountably infinite possible futures. How, then, are we to estimate the total predictable information from simulations given that we are *always* in the undersampled limit?

One approach is to simulate and employ sophisticated techniques for estimating entropy in the undersampled limit, e.g., as in Nemenman et al. (2002). However, even estimating the joint probability distribution of environment and sensor states from simulations can require long simulations, as we find in one example below. Alternatively, one can rewrite prediction and dissipation metrics in terms of generators of the environment and sensor, similar to how the predictive information bottleneck can be recast in terms of forward- and reverse-time causal states (Marzen and Crutchfield 2016). The following pursues this second line of inquiry by solving a Chapman–Kolmogorov equation that yields a partial differential equation for the joint probability distribution of forward-time causal states and channel state. We find new closed-form expressions for total predictable information, power consumption, instantaneous memory, and nonpredictive information rate in nonequilibrium steady-state sensors. The challenge to this analytical approach, of course, is that it requires accurate models of the environment and sensor dynamics.

We address a very general class of environments and sensors, here conceptualized as input-dependent dynamical systems: conditionally Markovian channels subject to a realization of a stationary unifilar hidden semi-Markov process (Marzen and Crutchfield 2017a) with an energy function associated with each environment and channel state combination. This setting is sufficiently general that it allows nonequilibrium sensing. The latter is crucial for accurate sensing (Mehta and Schwab 2012; Govern and ten Wolde 2014; Lang et al. 2014; Bo et al. 2015; Mancini et al. 2016), in that detailed balance may not hold and the Boltzmann distribution need not be equivalent to the steady-state distribution over sensor states for a frozen environment. Unlike previous treatments, we do not explicitly demarcate a separate memory or energy source accessed by the sensor (Bo et al. 2015; Ito and Sagawa 2015; Sartori et al. 2014), but include components as part of the sensor. Our framework allows for nondetailed-balance dynamics, implying the presence of multiple reservoirs (Van den Broeck and Esposito 2015). A lack of detailed balance in sensor dynamics is a reminder of the presence of multiple reservoirs. Finally, the environment is assumed to be unaffected by the sensor. That is, either there is no feedback or the environment is infinitely large. After all these assumptions, we calculate the heat dissipated by the sensor into the reservoir of interest.




Fig. 1 Hill molecule: (Top) Ion channel in a membrane in the open (left) and closed (right) states in which ions can and cannot travel through. (Below) Ligand concentrations during the open-closed cycle of the molecule

For illustration, we study both large random channels (randomly wired channels with a large number of states) and an “optimal” Hill molecule, which is a simple model of a ligand-gated channel; see Fig. 1. For these examples, we assume that the ligand concentration is a realization of a semi-Markov process—a generalization over previous efforts that assumed Markovian (Barato et al. 2014; Brittain et al. 2017) or Gaussian (Becker et al. 2015) processes. This generalization is necessary when, for instance, the Hill molecule represents a nicotinic acetylcholine receptor on a synapse and ligands are acetylcholine molecules since, as a practical matter, neuronal dynamics are often non-Markovian and non-Gaussian (Izhikevich 2007).

We find that (1) increases in cooperativity (sharpness of response) of the Hill molecule lead to increases in both prediction power and heat dissipation rate, (2) a large fraction of the heat dissipation rate comes from inefficient prediction, and (3) simple gradient-based adaptation rules lead to hysteresis. Furthermore, we find that the total predictable information captured by this seemingly impoverished Hill molecule exceeds the total predictable information captured by large random channels by an order of magnitude, despite the latter’s ability to capture potentially useful information about the past of the stimulus. This latter result would be impossible to obtain via interactions with a continuous-time Markovian environment.

2 Background

Central to our analysis is an appreciation of causal states (minimal sufficient statistics of prediction or retrodiction), unifilar hidden semi-Markov processes, and conditionally Markovian channels. We review these concepts here, simultaneously introducing relevant notation.

2.1 Environment

Input symbols x take on any value in the observation alphabet \mathcal{A} . We code the *past* of the input time series as $\overleftarrow{x} = \dots(x_{-2}, \tau_{-2}), (x_{-1}, \tau_{-1}), (x_0, \tau_+)$ and the input’s

future as $\vec{x} = (x_0, \tau_-), (x_1, \tau_1), (x_2, \tau_2), \dots$, where τ_i is the total dwell time for symbol x_i . To ensure a unique coding, we stipulate that $x_i \neq x_{i+1}$. Note that symbol x_0 is seen for a total dwell time of $\tau_+ + \tau_- = \tau_0$; that is, the present splits the dwell time τ_0 in two.

As is typical, \overleftarrow{X} is the random variable corresponding to semi-infinite input pasts and \overrightarrow{X} the random variable corresponding to semi-infinite input futures. We now briefly review the definition of causal states, as described in Shalizi and Crutchfield (2001) and measure-theoretically defined in Löhner (2010, 2012). Forward-time causal states S^+ , the minimal sufficient statistics for prediction, are defined via the following equivalence relation: two semi-infinite pasts, \overleftarrow{x} and \overleftarrow{x}' , are considered “predictively” equivalent if:

$$\overleftarrow{x} \sim_{\epsilon^+} \overleftarrow{x}' \Leftrightarrow \Pr(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}) = \Pr(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}').$$

The relation partitions the set of semi-infinite pasts into clusters of pasts. Each cluster is a *forward-time causal state* σ^+ , a realization of the random variable S^+ . *Reverse-time causal states* S^- , the minimal sufficient statistics for retrodiction, are defined similarly. Two semi-infinite futures, \overrightarrow{x} and \overrightarrow{x}' , are considered “retrodictively” equivalent if:

$$\overrightarrow{x} \sim_{\epsilon^-} \overrightarrow{x}' \Leftrightarrow \Pr(\overleftarrow{X} | \overrightarrow{X} = \overrightarrow{x}) = \Pr(\overleftarrow{X} | \overrightarrow{X} = \overrightarrow{x}').$$

This equivalence relation partitions the set of semi-infinite futures into clusters, each cluster being a *reverse-time causal state* σ^- , a realization of the random variable S^- .

Forward- and reverse-time causal states are useful in the ensuing calculations due to the following Markov chains. First, forward-time causal states are a deterministic function of the input past ($\sigma^+ = \epsilon^+(\overleftarrow{x})$), and reverse-time causal states are a deterministic function of the input future ($\sigma^- = \epsilon^-(\overrightarrow{x})$). Hence, we have the Markov chains $S^+ \rightarrow \overleftarrow{X} \rightarrow \overrightarrow{X}$ and $S^- \rightarrow \overrightarrow{X} \rightarrow \overleftarrow{X}$; so that, for instance:

$$p(\overrightarrow{x}, \sigma^+ | \overleftarrow{x}) = p(\overrightarrow{x} | \overleftarrow{x})p(\sigma^+ | \overleftarrow{x}),$$

where we introduce a simplified notation for the probability distributions; e.g., $p(\sigma^+ | \overleftarrow{x}) = \Pr(S^+ = \sigma^+ | \overleftarrow{X} = \overleftarrow{x})$. Note that $p(\sigma^+ | \overleftarrow{x})$ is singly supported—that is, a Kronecker delta $\delta_{\sigma^+, \epsilon^+(\overleftarrow{x})}$. However, causal states are minimal sufficient statistics of the past relative to the future and vice versa. And so, $\overleftarrow{X} \rightarrow S^+ \rightarrow \overrightarrow{X}$ and $\overrightarrow{X} \rightarrow S^- \rightarrow \overleftarrow{X}$ are also valid Markov chains; so that, for instance:

$$p(\overrightarrow{x}, \overleftarrow{x} | \sigma^+) = p(\overrightarrow{x} | \sigma^+)p(\overleftarrow{x} | \sigma^+).$$

Invoking these Markov chains is called *causal shielding*.

For the discussion that follows, we must define hidden Markov models and unifiarity. We first do so in the case of discrete-time, discrete-event processes. A hidden Markov model is equipped with hidden states and labeled transition probabilities that give the probability for emitting a particular symbol and transitioning to a new




Fig. 2 Two presentations of a unifilar hidden semi-Markov model. At bottom, a generative model for a discrete-alphabet, continuous-time stochastic process. Dwell times τ are drawn upon transitions between states, and the corresponding symbol is shown for that amount of time. At top, the corresponding “conveyor belt” representation of the process generated by the model beneath. Conveyor belts represent the time since last symbol based on the height along the conveyor belt traveled; each conveyor belt has a symbol. (Used with permission from Marzen and Crutchfield (2017a).)

particular hidden state. (These are edge-emitting hidden Markov models, which are completely general and equivalent upon transformation to the likely more familiar state-emitting hidden Markov models.) In a unifilar hidden Markov model, the state to which you transition is determined by the state that you are in and the symbol that you emit. In other words, if g is the hidden state at time t , then $p(g_{t+1}|x_{t+1}, g_t)$ is singly supported. Any nonunifilar hidden Markov model can be converted into a unifilar hidden Markov model with equivalent statistics. However, typically, such a unifilar hidden Markov model will have an uncountable infinity of states (Marzen and Crutchfield 2017b). The process generated by a finite unifilar hidden Markov model has a finite number of causal states (Travers and Crutchfield 2014), and the generalization of this statement to the continuous-time, discrete-event domain causes us to focus on unifilar hidden semi-Markov models.

Now, we turn our attention to continuous-time, discrete-event processes. We focus on unifilar hidden semi-Markov input, as in Marzen and Crutchfield (2017a). Causal states of unifilar hidden semi-Markov processes have three distinct components, in a way that can be visualized as state- and symbol-dependent conveyor belts that transition into one another. See Fig. 2. Each conveyor belt is labeled by a forward-time

hidden state g . The forward-time hidden state is analogous to hidden states in unifilar hidden Markov models, except that forward-time hidden states now emit both symbols and dwell times. Dwell times are drawn from $\phi_g(\tau)$; the corresponding emitted symbol x is chosen with probability $p(x|g)$; and $\epsilon^+(g, x)$ is the next hidden state given that the current hidden state is g and the current emitted symbol is x . In other words, x is emitted upon transition from g to the next hidden state g' . The last of these properties is called *unifilarity*.

Having set up this machinery, forward-time causal states are labeled by $\sigma^+ = (g_+, x_+, \tau_+)$. That is, the forward-time hidden state g_+ , current emitted symbol x_+ , and time since last symbol τ_+ (or equivalently, the time since x_+ was emitted) together comprise the forward-time causal states for unifilar hidden semi-Markov processes. Similarly, the reverse-time causal state $\sigma^- = (g_-, x_-, \tau_-)$ for a unifilar hidden semi-Markov process is a combination of reverse-time hidden state g_- , current emitted symbol x_- , and time to next symbol τ_- .

In what follows, we limit ourselves to *finite* unifilar hidden semi-Markov processes—that is, unifilar hidden semi-Markov processes with a finite number of hidden states. While Marzen and Crutchfield (2017a)'s hidden semi-Markov processes are completely general, they are restricted in that the unifilar hidden semi-Markov process corresponding to a particular *nonunifilar* hidden semi-Markov process might have unordered uncountably infinite states. That said, we believe that any continuous-time, discrete-event process can be approximated arbitrarily well by a countable unifilar hidden semi-Markov process. (Alas, a proof of this is still lacking.) As such, we do not believe that restricting ourselves to finite unifilar hidden semi-Markov processes is limiting.

In our examples in Sect. 5, we focus on semi-Markov input, such that the aforementioned hidden states g are simply the current emitted symbol x . This greatly constrains the forward- and reverse-time causal states, to the point that $p(\sigma^+, \sigma^-)$ is easily attainable. (See App. C or below.) The forward-time causal states are now described by the pair (x_+, τ_+) , where x_+ is the input symbol infinitesimally prior to the present and τ_+ is the time since last symbol—i.e., since x_{-1} . The reverse-time causal states are similarly described by the pair (x_-, τ_-) , where x_- is the input symbol infinitesimally after the present and τ_- is the time to next symbol—i.e., until x_1 . Let T_{\pm} be the random variable describing time since (to) last (next) symbol. The dwell time of symbol x has probability density function $\phi_x(\tau)$, and the probability of observing symbol x after x' is $q(x|x')$. By virtue of how we have chosen to encode our input: $q(x|x) = 0$.

Finally, the development to come requires finding the joint density $\rho(\sigma^+, \sigma^-)$ of forward- and reverse-time causal states. Let:

$$\Phi_{x_+}(\tau_+) = \int_{\tau_+}^{\infty} \phi_{x_+}(t)dt \text{ and}$$

$$\mu_{x_+}^{-1} = \int_0^{\infty} t\phi_{x_+}(t)dt.$$

As detailed in Marzen and Crutchfield (2017a), for the conditional density, we can say:

$$\begin{aligned} \rho(\sigma^-|\sigma^+) &= \rho(g_-, x_-, \tau_-|g_+, x_+, \tau_+) \\ &= \delta_{x_+, x_-} \frac{\phi_{g_+}(\tau_+ + \tau_-)}{\Phi_{g_+}(\tau_+)} p(g_+|g_-, x_-). \end{aligned}$$

For semi-Markov input, this simplifies the density: $\rho(\sigma^+, \sigma^-) = \rho((x_+, \tau_+), (x_-, \tau_-))$. As described in Marzen and Crutchfield (2017a), we have:

$$\rho(x_+, \tau_+) = \mu_{x_+} \Phi_{x_+}(\tau_+) p(x_+). \tag{1}$$

Let $p(x_+)$ be the probability of observing symbol x_+ . This probability is given by:

$$p(x_+) = (\text{diag}(1/\mu_x) \text{eig}_1(q))_{x_+},$$

where $\text{diag}(1/\mu_x)$ is a diagonal matrix and $\text{eig}_1(q)$ is the eigenvector of $q(\cdot)$ associated with eigenvalue 1. The conditional density of reverse-time causal states given forward-time causal states is then:

$$\begin{aligned} \rho(\sigma^-|\sigma^+) &= \rho((x_-, \tau_-)|(x_+, \tau_+)) \\ &= \frac{\phi_{x_+}(\tau_+ + \tau_-)}{\Phi_{x_+}(\tau_+)} \delta_{x_+, x_-}. \end{aligned} \tag{2}$$

Together, Eqs. (1) and (2) give the joint density $\rho(\sigma^+, \sigma^-) = \rho(\sigma^+) \rho(\sigma^-|\sigma^+)$.

2.2 Sensory Channel

We assume the channel is conditionally Markovian. As such, its dynamics are fully specified by input state-dependent kinetic rates. More precisely, the channel state y , with corresponding random variable Y , can take on any value in \mathcal{Y} , and the rate at which channel state y transfers to channel state y' when the input has value x is given by $k_{y \rightarrow y'}(x)$. Probability conservation dictates that:

$$k_{y \rightarrow y}(x) := - \sum_{y'} k_{y \rightarrow y'}(x).$$

Then, the probability $p(y, t)$ of being in channel state y at time t evolves as:

$$\dot{p}(y, t) = \sum_{y'} k_{y' \rightarrow y}(x(t)) p(y', t),$$

where $x(t)$ is the input symbol at time t .

To simplify notation and ease computation, we write dynamical evolution rules in matrix-vector form. Let $\mathbf{p}(y, t)$ be the column vector of probabilities that the channel is in a particular state y at time t , and let $M(x)$ be a matrix of rates: $M_{y', y}(x) = k_{y \rightarrow y'}(x)$.




Fig. 3 Sensor information diagram (James et al. 2011) giving the relationship between prediction metrics and nonpredictive information rate. Instantaneous memory is $I_{\text{mem}} = I[X_0; Y_0] = b + d + f + h$ and total predictable information $I_{\text{fut}} = \lim_{\Delta t \rightarrow 0} I[Y_0; X_0] = \lim_{\Delta t \rightarrow 0} (b + c + d + e + f + g + h)$, while the lower bound on power consumption, the nonpredictive information rate is $\dot{I}_{\text{np}} = \lim_{\Delta t \rightarrow 0} (b + h - c - e) / \Delta t$. Recall $a = H[Y_0|X_0, X_{\Delta t}, \vec{X}_{2\Delta t}]$, $b = I[Y_0; X_0|X_{\Delta t}, \vec{X}_{2\Delta t}]$, $f = I[Y_0; X_0; X_{\Delta t}; \vec{X}_{2\Delta t}]$, and so on

Then, we have:

$$\dot{\mathbf{p}}(y, t) = M(x(t)) \mathbf{p}(y, t). \tag{3}$$

with this, it is clear that $\mathbf{p}(y, t)$ can oscillate or decay to a steady state. The Perron–Frobenius theorem guarantees that:

$$p_{\text{eq}}(x) := \text{eig}_0(M(x)),$$

the probability distribution over channel states, when ligand concentration is set to x , is unique. This need not be a Boltzmann distribution due to the presence of unspecified reservoirs.

3 Prediction and Dissipation Metrics

We employ several metrics to quantify the sensor’s prediction performance and its energetic efficiency. *Instantaneous memory* I_{mem} (Still et al. 2012) and *total predictable information* I_{fut} (Still et al. 2010; Creutzig et al. 2009) characterize prediction power, while the *nonpredictive information rate* \dot{I}_{np} (Still et al. 2012; Horowitz and Esposito 2014; Barato et al. 2014) and temperature-normalized *power consumption* βP monitor dissipation. In nonequilibrium steady state, if detailed balance holds, heat dissipation rate and power consumption are equivalent. Figure 3 uses an information diagram (Yeung 1991; James et al. 2011) to illustrate the relationships between the various information-theoretic quantities in terms of the elementary information atoms—entropies, conditional entropies, and mutual information—out of which they are constructed.

In defining the sensor, we grouped into a single sensor what might separately be called sensor, memory, and internal energy source(s). This allows us to consider active sensors with memory. However, reading the sensor by other downstream parts of an organism might entail a different grouping. In any case, since we use information-theoretic quantities to define prediction metrics, the data processing inequality employed below guarantees that we establish a (not necessarily tight) upper bound on the instantaneous memory and total predictable information when reading a sensor.

Our selection of metrics is inspired by previous efforts to characterize prediction and dissipation in sensors. Instantaneous memory and nonpredictive information were first defined by Still et al. (2012). Horowitz and Esposito (2014) also focused on nonpredictive information rate and power consumption, but did not calculate total predictable information nor a more standard prediction-related metric. Barato et al. (2014) used the ratio of nonpredictive information rate to entropy production to characterize learning, though nonpredictive information rate is not necessarily an intuitive metric for learning. Moreover, this ratio can be greater than unity when nonpredictive information rate is negative. Finally, Becker et al. (2015) focused on metrics for prediction, including a natural continuous-time extension of instantaneous predictive information, but not on metrics for dissipation.

3.1 Prediction Metrics

Instantaneous memory and total predictable information are far from being the only metrics useful for characterizing a sensor's prediction capability. One could instead study instantaneous predictable information $I[Y_t; X_{t+\Delta t}]$ which requires choosing a relevant Δt (Still et al. 2012; Palmer et al. 2015). Or, following the echo-state network literature, one might instead calculate “memory capacity” (Jaeger 2001) or “prediction capacity” (Marzen 2017). Or, if a reward function is known (Sutton and Barto 1998), then the relevant metric would depend on the reward function and the organism's action policy.

Interestingly, information-theoretic prediction metrics might have more relevance to biology than, say, the so-called prediction capacity. In a discrete-time setting, extending Kelly's classic bet-hedging analysis shows that an increase in expected log-growth rate via an increase in sensory information is equal to the instantaneous predictable information $I[Y_0; X_{\Delta t}]$ (Marzen and Crutchfield 2018).

The *total predictable information* I_{fut} , defined as:

$$I_{\text{fut}} := I[Y; \vec{X}],$$

is an upper bound on this increase in expected log-growth rate. (Note that we dropped time subscripts t due to stationarity.) This is the mutual information between present channel state and the input's future. It is the amount of information that is predictable about the input future from the present channel state. We merely assert that on ontogenetic timescales total predictable information constitutes a reasonable reward function (Still 2009; Tishby and Polani 2011).

Even in a discrete-time setting, calculating the total predictable information appears intractable, as there is an uncountable infinity of possible input futures. [An exception to this is a Gaussian process input to linear dynamical systems; e.g., see Creutzig et al. (2009).] To overcome this, we employ the causal-shielding relations specified in Sect. 2. Due to the feedforward nature of the channel-input setup—that is, the channel’s state does not affect the input—and the Markov chains given earlier—e.g., see Crutchfield et al. (2009)—we have the Markov chain $Y \rightarrow \mathcal{S}^+ \rightarrow \mathcal{S}^- \rightarrow \overrightarrow{X}$ and the Markov chain $Y \rightarrow \mathcal{S}^+ \rightarrow \overrightarrow{X} \rightarrow \mathcal{S}^-$. Two applications of the data processing inequality yield:

$$\begin{aligned} I_{\text{fut}} &= I[Y; \mathcal{S}^-] \\ &= I[Y; X_-, \mathcal{T}_-, \mathcal{G}_-]. \end{aligned}$$

In other words, when we calculate how much information we share with the future, in fact, we are calculating how much information we share with the aspects of the future that we can predict—the minimal sufficient statistics of retrodiction. (That is, \mathcal{S}^- , and somewhat confusingly not \mathcal{S}^+ , is a sufficient statistic for the future trajectory.) I_{fut} then decomposes into:

$$I_{\text{fut}} = I[Y; X_-] + I[Y; \mathcal{T}_-, \mathcal{G}_- | X_-].$$

Note that X_- is equal to X_+ almost surely.

The term $I[Y; X_-]$ is called the *instantaneous memory* I_{mem} (Still et al. 2012), since it is the amount of information available from the channel state about the just-seen input symbol. Rewriting we have:

$$I_{\text{fut}} = I_{\text{mem}} + I[Y; \mathcal{T}_-, \mathcal{G}_- | X_-]. \tag{4}$$

Thus, the total predictable information is the sum of instantaneous memory and information that is truly about the future, which here is the combined time to next-symbol and reverse-time hidden state. For the special case of semi-Markov input, the reverse-time hidden state is equivalent to the present observed symbol and so:

$$I_{\text{fut}} = I_{\text{mem}} + I[Y; \mathcal{T}_- | X_-].$$

The difference between the total predictable information and instantaneous memory for semi-Markov input is the information that the present sensor state captures about the time to next observed symbol. Generally, the time to next observed symbol can be better predicted using the time since last symbol. However, for a Markov input, $I_{\text{fut}} = I_{\text{mem}}$ because the time to next symbol is independent of the time since last symbol—a property that follows from the exponential nature of the dwell time distribution.

3.2 Dissipation Metrics

Next, we quantify the power extracted from the environment by the sensor system. Note that this does not include the power required to maintain the sensor in nonequilibrium

steady state, even at fixed environment, due to the violation of detailed balance. In other words, we do not account for the power needed for internal maintenance of the sensor. Assuming access to a temperature-normalized “energy function” $\beta E(x, y)$, the temperature-normalized power βP is given by:

$$\beta P = \lim_{\Delta t \rightarrow 0} \frac{\langle \beta E(x_{t+\Delta t}, y_t) \rangle - \langle \beta E(x_t, y_t) \rangle}{\Delta t}. \quad (5)$$

In effect, we group any internal energy sources into the sensor state, so that we calculate the work done by the environment on the entire sensor and its internal energy sources. Given an energy function, we do not assume that the distribution over sensor states reached, if the environment is fixed, is identical to the Boltzmann distribution.

If determining an energy function is not possible, we can calculate a lower bound using a continuous-time adaptation of the inequality in Still et al. (2012):

$$\dot{I}_{\text{np}} := \lim_{\Delta t \rightarrow 0} \frac{I[Y_t; X_t] - I[Y_t; X_{t+\Delta t}]}{\Delta t} \leq \beta P, \quad (6)$$

with an alternate equivalent definition in Horowitz and Esposito (2014). Note that Eq. (6) only holds in nonequilibrium steady state. See Appendix A.

\dot{I}_{np} is called the *nonpredictive information rate* since it loosely corresponds to how much of the instantaneous memory is useless for predicting the next input. Barato et al. (2014) viewed $\dot{I}_{\text{np}}/\beta P$ as a learning efficiency, though; see Brittain et al. (2017). We take the view that \dot{I}_{np} is a potentially useful lower bound on temperature-normalized power consumption and use I_{mem} and I_{fut} instead to characterize learning. Any differences between the formulae shown here and in Still et al. (2012 Eq. (2)) are superficial; we merely adapted the derivation for continuous-time processes. Unfortunately, the nonpredictive information rate is not necessarily a tight lower bound on temperature-normalized power. When the environment is Markovian, there is another lower bound on the heat dissipation rate that is proportional to the instantaneous memory (Das et al. 2017), but we focus mainly on non-Markovian environments here.

There are other information measures related to dissipation that we do not address here. Most popular of these is the transfer entropy, which is interpreted to quantify the causal influence of one system on another (Ito and Sagawa 2013; Hartich et al. 2014; Ito and Sagawa 2015); however, see James et al. (2016) for arguments that such information measures can be related to causal structure in nonintuitive ways. We also do not calculate the predictive capacity (Spinney et al. 2018), not to be confused with that in Marzen (2017), nor the causal irreversibility.

4 Results

Given a known environment and sensor, two approaches to evaluate the aforementioned prediction and dissipation metrics present themselves:

- First, simulate the environment and sensor and approximate prediction and dissipation metrics based on observed frequencies; or,
- Second, find (new) closed-form expressions for prediction and dissipation metrics in terms of the environment generators ($\phi_x(\tau), \epsilon^+(g, x), p(x|g)$) and sensor $M(x)$.

The first can lead to prohibitively long simulations for accurate estimates. The following pursues the second. This requires solving a Chapman–Kolmogorov equation that turns into partial differential equations with coupled boundary conditions. The explicit derivations are lengthy and therefore are relegated to the appendices. The appropriate equations there are referenced here, where relevant.

That said, let us briefly expand upon the first approach to highlight its several potential difficulties. To calculate prediction and dissipation metrics other than the total predictable information I_{fut} , one can estimate $p(x_t, y_t)$ and $p(x_{t+\Delta t}, y_t)$ for some small Δt from simulation and then use standard formulae ($I[U; V] = \sum_{u,v} p(u, v) \log(p(u, v)/p(u)p(v))$) and Eqs. (7) and (10) to calculate I_{mem} , \dot{I}_{np} , and βP . More sophisticated algorithms for calculating prediction metrics would directly approximate $I[X_t, Y_t]$ (Nemenman et al. 2002) instead of “plugging-in” estimates of the corresponding probability distribution. As total simulation time T increases, estimates of the various probability distributions from empirical frequencies become more and more accurate and our corresponding estimates of prediction and dissipation metrics also become more accurate. As discussed in Sect. 5, this might require surprisingly long times T to obtain accurate estimates.

Using the first approach to calculate the total predictable information I_{fut} from simulations can be more difficult. Suppose we do not utilize the causal state machinery and ignore Eq. (4). Then, we work with a continuous-time system and semi-infinite futures and are always in a grossly undersampled limit. Alternatively, suppose that we do employ causal-state machinery and calculate I_{fut} using Eq. (4). Then, progress can be made, as there are mutual information estimators to calculate mutual information of mixed random variables (Victor 2002). Rather than test the efficacy of mutual information estimators, we will not calculate I_{fut} from simulations.

Notably, from the closed-form expressions of Sect. 4 later on, we find that I_{fut} is close to I_{mem} for random channels and the simple Hill molecule. This result, however, will likely not hold for the interesting case of highly predictive sensors in complex environments (Bialek et al. 2001).

4.1 Calculating Prediction Metrics

I_{fut} and I_{mem} can be analytically calculated, though once $\rho(\sigma^+, y) = \Pr(S^+ = \sigma^+, Y = y)$ is in hand. This follows, in turn, by manipulating a Chapman–Kolmogorov equation, shown in Appendix B. Set any ordering on the pairs (g, x) ; e.g., the ordering $(g_1, x_1), (g_1, x_2), \dots, (g_{|\mathcal{G}|}, x_{|\mathcal{A}|})$. An expression for $p(y|\sigma^+) = p(y|g, x, \tau)$ is given by a combination of Eqs. (B5) and (B9):

$$p(y|g, x, \tau) = \left(e^{M(x)\tau} \text{eig}_1(\mathbf{C})_{(g,x)} / \mu_g p(g) \right)_y,$$

where \mathbf{C} is a block matrix with entries:

$$\mathbf{C}_{(g,x),(g',x')} = \delta_{g,\epsilon^+(g',x')} p(x'|g') \int_0^\infty \phi_{g'}(t) e^{M(x')t} dt,$$

and $\text{eig}_1(\mathbf{C})$ is a block vector. Normalization forces $\mathbf{1}^\top \text{eig}_1(\mathbf{C})_{(g,x)} = \mu_g p(g)$.

We then find the joint probability density as $\rho(y, \sigma^+) = p(y|\sigma^+) \rho(\sigma^+)$, which enables computation of all prediction metrics. Instantaneous memory is given by $I_{\text{mem}} = I[X; Y]$, whereas $I_{\text{fut}} = I[Y; \mathcal{S}^-]$. All the relevant probabilities—namely $p(x, y)$ and $\rho(\sigma^-, y)$ —are obtained from the previously derived density $\rho(\sigma^+, y)$. For instance, to calculate $\rho(\sigma^-, y)$, we employ the Markov chain $Y \rightarrow \mathcal{S}^+ \rightarrow \mathcal{S}^-$ to find:

$$\rho(\sigma^-, y) = \sum_{\sigma^+} \rho(\sigma^-|\sigma^+) p(y|\sigma^+) \rho(\sigma^+).$$

Equation (2) gives $\rho(\sigma^-|\sigma^+)$. And, to calculate $p(x, y)$, we recall that $\sigma^- = (g_-, x, \tau_-)$, so we only need marginalize the joint density $\rho((g_-, x, \tau_-), y)$.

4.2 Calculating Dissipation Metrics

Additionally, dissipation metrics can be calculated once:

$$\frac{\delta p(x, y)}{\delta t} = \lim_{\Delta t \rightarrow 0} \left(\Pr(X_{t+\Delta t} = x, Y_t = y) - \Pr(X_t = x, Y_t = y) \right) / \Delta t$$

is obtained. (Note that this is not the total time derivative dp/dt .) An expression for $\delta p/\delta t$ in terms of input and channel properties is given in Eq. (B11):

$$\begin{aligned} \frac{\delta p}{\delta t} &= \sum_{g', x' \neq x} \int d\tau' p(x|\epsilon^+(g', x')) p(x'|g') \phi_{g'}(\tau') \\ &\quad \times \left(e^{M(x')\tau'} \text{eig}_1(\mathbf{C})_{(g', x')} \right)_y \\ &\quad - \sum_{g'} \int d\tau' p(x|g') \phi_{g'}(\tau') \left(e^{M(x)\tau'} \text{eig}_1(\mathbf{C})_{(g', x)} \right)_y, \end{aligned}$$

where normalization again requires $\mathbf{1}^\top \text{eig}_1(\mathbf{C})_{(g,x)} = \mu_x p(g)$. Then, from Eq. 6, we find that:

$$\begin{aligned}
 \dot{I}_{\text{np}} &= \lim_{\Delta t \rightarrow 0} \frac{H[Y_t, X_{t+\Delta t}] - H[Y_t, X_t]}{\Delta t} \\
 &= \lim_{\Delta t \rightarrow 0} \left(\sum_{x,y} \left(p(x, y) + \frac{\delta p}{\delta t} \Delta t \right) \log \frac{1}{p(x, y) + \frac{\delta p}{\delta t} \Delta t} - \sum_{x,y} p(x, y) \log \frac{1}{p(x, y)} \right) / \Delta t \\
 &= - \sum_{x,y} \frac{\delta p(x, y)}{\delta t} \log p(x, y). \tag{7}
 \end{aligned}$$

If there is an energy function $E(x, y)$, then it is straightforward to show that power dissipated as heat into the reservoir of interest can be obtained from:

$$\beta P = \sum_{x,y} \frac{\delta p(x, y)}{\delta t} E(x, y). \tag{8}$$

For passive sensors, despite lacking direct access to an energy function, we find βP by calculating the steady-state distribution over channel states with fixed input:

$$\begin{aligned}
 \text{Peq}(y|x) &= \text{eig}_0(M(x)) \\
 &= \frac{e^{-\beta E(x,y)}}{Z_\beta(x)},
 \end{aligned}$$

where the partition function is $Z_\beta(x) := \sum_y e^{-\beta E(x,y)}$. Hence:

$$\beta E(x, y) = \log \frac{1}{\text{peq}(y|x)} - \log Z_\beta(x). \tag{9}$$

Recalling Eq. (5) and invoking stationarity—that $\Pr(X_t = x) = \Pr(X_{t+\Delta t} = x)$ —yields:

$$\begin{aligned}
 \beta P &= \lim_{\Delta t \rightarrow 0} \left(\left\langle \log \frac{1}{\text{peq}(x, y)} \right\rangle_{\Pr(X_{t+\Delta t}=x, Y_t=y)} - \left\langle \log \frac{1}{\text{peq}(x, y)} \right\rangle_{\Pr(X_t=x, Y_t=y)} \right) / \Delta t \\
 &= \sum_{x,y} \frac{\delta p(x, y)}{\delta t} \log \frac{1}{\text{peq}(x, y)}. \tag{10}
 \end{aligned}$$

The distributions $\Pr(X_{t+\Delta t} = x, Y_t = y)$ and $\Pr(X_t = x, Y_t = y)$ can be obtained from $M(x)$. In other words, for passive sensors, we can calculate βP directly from the kinetic rates $k_{y \rightarrow y'}(x)$ and input generator $(\phi_g(\tau), \epsilon^+(g, x), p(x|g))$ alone.

5 Examples

Using this theory, we now study four sensor–environment examples. The first two—a two-state sensor channel in a binary environment—serves to both check our formulae and quantify the gains in computational efficiency that result from using these formulae. The third studies the performance of large random sensor channels at lossy

prediction. The fourth example is a Hill molecule sensor in a fluctuating environment. In this, we wish to find Hill molecules that “optimally” balance the need to predict sensory input with constraints on heat dissipation. Thereafter, we compare to prior results on optimized biochemical sensors.

All code was implemented in Python. In this section, we focus on semi-Markov environments, for which g is commensurate with x . This important special case is also treated separately in the Supplemental Information.

5.1 Two-State Sensors in Markovian Environments

When the environment is Markovian, we can check our formula for $p(x, y)$ in two ways. First, we can analytically calculate (using similar ideas) a seemingly different expression for $p(x, y)$. Though an analytical match between the formula given here and the formulae shown below is not obvious, there is an exact numerical match for all tested randomly chosen sensor channels. Next, we simulate the system using the temporal Gillespie algorithm (Vestergaard and Génois 2015), and this results in a nearly exact numerical match with simulations.

The Markovian environment that we consider has two observed symbols with dwell time densities of $\phi_A(t) = 4e^{-4t}$ and $\phi_B(t) = 5e^{-5t}$. This implies that the kinetic rate at which one moves from A to B is $k_{A \rightarrow B} = 4$ and that at which one moves from B to A is $k_{B \rightarrow A} = 5$.

For a Markovian environment, only the present environmental symbol is needed to predict future environmental states. In the language above, the present environmental symbol is the causal state. As a result, to calculate $p(x, y)$, we need only find the stationary distribution of a master equation with kinetic rates given by:

$$k_{(x,y) \rightarrow (x',y')} = \begin{cases} 0 & x \neq x', y \neq y' \\ k_{y \rightarrow y'}^{(x)} & x = x', y \neq y' \\ k_{x \rightarrow x'} & x \neq x', y = y' \\ -k_{y \rightarrow y}^{(x)} - k_{x \rightarrow x} & x = x', y = y'. \end{cases} \tag{11}$$

We use these kinetic rates to analytically characterize the evolution of the system’s joint probability distribution $p(x, y)$ as:

$$\frac{dp(x, y)}{dt} = \sum_{x', y'} k_{(x',y') \rightarrow (x,y)} p(x', y'). \tag{12}$$

That is, the steady state $p(x, y)$ is an eigenvector of eigenvalue 0 of the kinetic rate matrix $k_{(x',y') \rightarrow (x,y)}$. Using this technique, we get an exact match with Eqs. (B5) and (B9) for all randomly chosen two-state channels with the Markovian environment listed above. Ideally, we would find an analytic match in the case of Markovian environments between the two analytic methods. However, we could not find an obvious mapping. The exact numerical match is good evidence that it exists, however.

Finally, we compare theory to simulation using the temporal Gillespie algorithm (Vestergaard and Génois 2015). The standard Gillespie algorithm (Gillespie 1976) assumes constant environments in which the waiting-time density to the next transition between sensor states is a simple exponential. When the environment fluctuates, the waiting-time density changes from a simple exponential to the exponential of an integral of the time-varying kinetic rate out of the current sensor state. In the present case, this time-varying kinetic rate is piecewise constant. We first generate exponentially distributed random numbers that correspond to the dwell times in the environment and, then, simulate transitions between sensor states using this temporal Gillespie algorithm. The convergence to agreement with analytics is shown in Fig. 4. Error is less than 10^{-3} .

5.2 Two-State Sensors Driven by Semi-Markov Environments

The advantages of the analytic approach championed here are not clear when considering a Markovian environment. Here, the more typical analytic approach is faster and more accurate as it requires no integrals.

To better illustrate the advantages, we now consider a random two-state sensor channel driven by a semi-Markov environment in which the dwell-time density for symbol A is $\phi_A(t) = 16te^{-4t}$ and the dwell-time density for symbol B is $\phi_B(t) = 25te^{-5t}$. Figure 5 shows comparison of the aforementioned temporal Gillespie algorithm to the formulae developed above. In particular, it shows that as time increases, the simulation becomes more accurate, as expected.

The simulation is quite efficient, partly since the environment is easily simulated as the negative log of the product of two uniform random variables and partly since the temporal Gillespie algorithm is fast. We speculate that the advantages of analytic methods are better illustrated by an environment that is less easy to simulate.

5.3 Random Multistate Sensors Driven by Semi-Markov Environments

Now consider lossy prediction—in which not all predictive information is captured—by large (multistate) random sensor channels in a semi-Markov environment. The fact that random channels achieve coding-theoretic limits (Cover and Thomas 1991) strongly suggests that random sensors must be explored. On the one hand, in principle they potentially can store information about an environment's past that is useful for prediction. On the other, they cannot store all the information about the past necessary to predict the future as well as possible (Marzen and Crutchfield 2017a). As such, large random sensors are *lossy* predictors.

We ask two questions. First, how well do random sensors perform relative to the optimal possible performance? One can calculate bounds on lossy prediction (Marzen and Crutchfield 2016) using a combination of causal-state machinery and the generalized Blahut–Arimoto algorithm. Second, are larger random sensors better or worse at lossy prediction? Potentially, larger sensor channels have more capacity to store information about the environment's past (Marzen 2017), but whether or not they deploy this capacity well is still undetermined.




Fig. 4 Two-state sensor in Markov environment—theory versus simulation: A random two-state sensor is chosen, such that $k_{B \rightarrow A}^{(A)} = 0.449$, $k_{A \rightarrow B}^{(A)} = 0.192$, $k_{B \rightarrow A}^{(B)} = 0.646$, and $k_{A \rightarrow B}^{(B)} = 0.397$. (Top) Comparing total variational distance $\sum_{x,y} |p(x,y) - p_{\text{sim}}(x,y)|$ between analytic estimate $p(x,y)$ from Appendix B and simulation estimate $p_{\text{sim}}(x,y)$ as a function of simulation time. (Bottom) The corresponding difference $\hat{I}_{\text{mem}} - I_{\text{mem}}$ between estimates of the instantaneous memory also decreases as simulation time increases, with $I_{\text{mem}} = 2.5 \times 10^{-5}$ nats. As time increases, error between the true $p(x,y)$ and $p_{\text{sim}}(x,y)$ decreases

We subjected random sensors of varying number of states $N \in \{3, 6, 10, 20, 30, 40, 50, 60, 100, 300, 1000\}$ to a semi-Markov environment in which the dwell-time density of symbol A is $\phi_A(\tau) = 16\tau e^{-4\tau}$ and the dwell-time density of symbol B is $\phi_B(\tau) = 25\tau e^{-5\tau}$. The kinetic rate $k_{y \rightarrow y'}^{(x)}$ between one sensor state y and another y' , conditioned on observing a particular environmental symbol x , is drawn uniformly from the unit interval.

Figure 6 (Top) shows that random sensors are not good lossy predictors in that their total predictable information $I[Y; \overline{X}]$ falls three orders of magnitude below the maximal achievable total predictable information for their coding rate $I[Y; \mathcal{S}^+]$. In using this bound on $I[Y; \overline{X}]$, we assumed that one recodes the dynamical system's




Fig. 5 Two-state sensor in a semi-Markov environment—theory versus simulation: Random two-state sensor as in Fig. 4. (Top) Total variational distance $\sum_{x,y} |p(x, y) - p_{\text{sim}}(x, y)|$ between true $p(x, y)$ and simulation estimate $p_{\text{sim}}(x, y)$ as a function of simulation time. As time increases, error between the true $p(x, y)$ and $p_{\text{sim}}(x, y)$ decreases. (Bottom) The corresponding difference $\hat{I}_{\text{mem}} - I_{\text{mem}}$ between estimates of the instantaneous memory also decreases as simulation time increases. Same kinetic rates as in Fig. 4. The true value is $I_{\text{mem}} = 2.9 \times 10^{-5}$ nats

state so as to eliminate extraneous information about the past using the procedure outlined in the proof of Theorem 1 of Marzen and Crutchfield (2016). This achieves a coding rate of $I[Y; S^+]$ rather than one of $I[Y; \vec{X}]$.

Note that we can calculate $I[Y; \vec{X}]$ exactly because $I[Y; \vec{X}] = I[Y; S^-]$ (Marzen and Crutchfield 2016), as the shared information between our sensor state and the future environment is exactly the shared information between our sensor state and the predictable parts of the future environment—the minimal sufficient statistics of retrodiction.

Perhaps more interestingly, Fig. 6 (Top) also shows that as sensor size grows, large random sensors become worse lossy predictors, in that they store information about

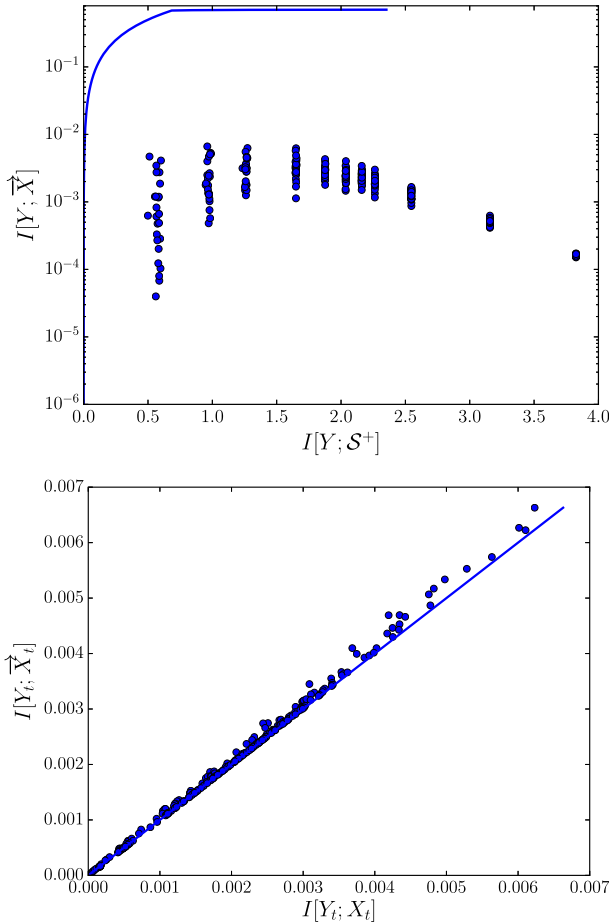


Fig. 6 (Top) Random sensors are not good lossy predictors of semi-Markov stimuli: Blue solid curve is the theoretical bound on total predictable information $I[Y; \vec{X}]$ for a given constraint on mutual information $I[Y; \mathcal{S}^+]$ between past and representation, calculated using the algorithm described in Marzen and Crutchfield (2016). Blue dots represent realizations of prediction performance and compression for randomly generated channels with states sets of size $N \in \{3, 6, 10, 20, 30, 40, 50, 60, 100, 300, 1000\}$ from left to right, with 25 random channels per size. All pieces of mutual information are reported in nats. (Bottom) Most prediction power comes from correctly storing information about the present environmental symbol: Each dot indicates the combination of instantaneous memory $I[Y_t; X_t]$ and total predictable information $I[Y_t; \vec{X}]$ of a random channel. The solid line corresponds to their equality. At maximum, $I[Y_t; X_t]$ is 0.0.687 nats and $I[Y_t; \vec{X}]$ is 0.7065 nats. All mutual information are reported in nats

the environment’s forward-time causal states that is not the “right” information to remember for predicting environment futures. This result is only potentially true in a non-Markovian environment, as in a continuous-time Markovian environment every bit stored about the forward-time causal states is useful for understanding environmental futures. Typically, both prediction and memory capacity of random sensors increase with increasing channel size (Marzen 2017). And so, the lack of efficacy of random sensors is somewhat surprising.

In addition, random sensors at fixed size tend to have similar achievable coding rates $I[Y; \mathcal{S}^+]$ and total predictable information $I[Y; \vec{X}]$, despite differences between random instantiation. Likely, this results from central limit theorem-like behavior. However, proving as much in even simpler settings can be difficult (Marzen and DeDeo 2016; Marzen 2018).

In a semi-Markov environment, the total predictable information I_{fut} is the shared information between our sensor's representation and both the present symbol and the time to next symbol. The dynamical system (sensor) obtains information about the time to next symbol from the time since last symbol. In particular, the time since last symbol can be used to better predict the time to next symbol, as the time between successive symbols is nonexponential (Marzen and Crutchfield 2017a). Figure 6 (Bottom) shows that random channels essentially are sensitive only to information about the current environmental symbol, which is predictive of immediately following environmental symbols.

If large random dynamical systems had been “good enough”, then evolutionary search to find predictive sensors would not be needed. Naturally, such a result flies in the face of current thinking (Palmer et al. 2015). Though many more examples need be considered, the present results for this class demonstrate that one cannot reliably design good lossy predictors simply by generating large randomly wired sensors. This hints at the challenge faced in the evolution of biological sensors.

Note that these results hold no matter the energy function. To better understand their thermodynamic efficiency, we now assume that they are passive sensors, so that the steady-state probability distribution in a fixed environment is equivalent to the Boltzmann distribution. We then calculate both nonpredictive information rate \dot{I}_{np} and power dissipated as heat into the reservoir of interest βP . While we average the dissipation metric over many channels of the same size, we track the standard deviations in Fig. 7. Mean \dot{I}_{np} varies nonmonotonically with channel size, peaking at around $N = 30$, while βP (equivalent to heat dissipation rate into the reservoir of interest in nonequilibrium steady state) decreases monotonically with channel size.

Therefore, larger channels have smaller entropy production rates, as shown by Fig. 6 (Top), even though they are emphatically not at equilibrium. The sensory capacity (Hartich et al. 2016)—the ratio of the nonpredictive information rate or learning rate to entropy production rate—increases monotonically with channel size. This qualitatively agrees with Barato et al. (2014) which claimed that adding memory to a sensor increases its sensory capacity. Similar to Brittain et al. (2017), then, we find that maximizing sensory capacity and maximizing prediction metrics are at odds. Maximizing sensory capacity yields large random channels, while maximizing prediction metrics from Fig. 6 yields small random channels.

5.4 Hill Molecule Interacting with a Semi-Markov Environment

The Hill molecule—a common fixture in theoretical biochemistry—is the simplest mechanistic model of binding cooperativity (Marzen et al. 2013). Recall Fig. 1. A Hill molecule can be in one of two states, open or closed. When open, n ligand molecules are bound; when closed, no ligand molecules are bound. Hence, the Hill molecule




Fig. 7 Dissipation metrics for random channels vary with channel size: Mean nonpredictive information rate \dot{I}_{np} and mean power consumption βP (equivalent to heat dissipation rate in nonequilibrium steady state) averaged over 25 random channels at each of several sizes: $N \in \{3, 6, 10, 20, 30, 40, 50, 60, 100, 300, 1000\}$. Error bars give standard deviations across these 25 realized random channels. Higher sensory capacity—larger ratio of nonpredictive information rate to power consumption $\dot{I}_{np}/\beta P$ —is attained with larger random channels, even though this yields the lower prediction metrics shown in Fig. 6 (Top)

state carries information about the number of bound ligand molecules. In other words, Hill molecules sense their environment’s ligand concentration. Though they seem impoverished in their simplicity, the following concludes that they are sensitive to an order-of-magnitude more total predictable information in their environment than the large random sensors just studied.

Let us outline a simple dynamical model of the Hill molecule’s operation. The transition rate from closed C to open O given a ligand concentration x is:

$$k_{C \rightarrow O} = k_O x^n, \tag{13}$$

while the rate from open O to closed C is:

$$k_{O \rightarrow C} = k_C. \tag{14}$$

Given fixed ligand concentration, the steady-state distribution is familiar:

$$\begin{aligned} \text{Pr}_{\text{eq}}(Y = O | X = x) &= \frac{k_O x^n}{k_C + k_O x^n} \\ &= \frac{x^n}{(k_C/k_O) + x^n}. \end{aligned}$$

Thus, the rates are determined by the number n of ligands that bind simultaneously—this is called the *cooperativity*. Although the mechanistic model makes sense only when

n is a nonnegative integer, this model is often used when n is any nonnegative real number; increases in n can still be thought of as increases in cooperativity. Equations (13) and (14) constitute a complete characterization of its channel properties.

Note that these dynamics obey detailed balance, so no power is required to sustain operation out of equilibrium in a fixed environment. This is a passive sensor for which the power extracted from the environment is identical to the heat dissipation rate, according to the first law of thermodynamics. It may seem *a priori* that sensor designers should not care how much energy a sensor must draw from its environment. However, the power extracted from the environment is dissipated as heat, and increased heat dissipation can require additional biological machinery; e.g., sweating when internal temperature rises. With this in mind, it would seem that biology should favor sensors with smaller power consumption.

Increasing the cooperativity n increases the steepness of the molecule's sigmoidal "binding curve"—the probability of being "on" as a function of concentration. In other words, the sensor becomes more switch-like and less a proportionately responding transducer of the input. If the concentration is greater than $(k_C/k_O)^{1/n}$, the switch is essentially "on" if n is high. A more switch-like sensor is useful if the optimal phenotype depends only upon the condition "ligand concentration greater than X". A less switch-like, smoother-responding sensor helps if the optimal phenotype depends on ligand concentration in a more graded manner.

The concentration scale is set by $(k_C/k_O)^{1/n}$, while the timescale is set by $1/k_C$; as such, we set both to $k_O = k_C = 1$ without loss of generality. We imagine that the ligand concentration alternates between high and low values: x_l and x_h , respectively. When there is less ligand ($x \approx x_l$), the Hill molecule reverts to and stays in the closed state. When there is ligand ($x \approx x_h > x_l$), it reverts to and stays in the open state. With no particular application in mind, we again imagine that the dwell-time densities take the form $\phi_x(\tau) = \lambda^2(x)\tau e^{-\lambda(x)\tau}$ with $\lambda(x_l) = 5$ and $\lambda(x_h) = 4$.

We now deploy the earlier formulae to study the prediction capabilities and dissipation tendencies of a Hill molecule subject to semi-Markov input. Previous studies of biological sensors found that increases in cooperativity accompanied increases in channel capacity (Tkačik et al. 2009; Walczak et al. 2010; Martins and Swain 2011) and in sensor accuracy (Govern and ten Wolde 2014) in equilibrium systems. Others studied the thermodynamics of prediction of cooperative biological sensors (Barato et al. 2014; Becker et al. 2015), but did not use the more general class of semi-Markov input which, we argued, is more typical of real environmental stimuli. They also did not calculate the total predictable information.

Absent a reward function, we assert that energetic rewards are proportional to I_{fut} (Still 2009; Tishby and Polani 2011). Altogether, this argues that $\alpha I_{\text{fut}} - \beta P$ is a reasonable objective function for sensor design, similar to that of Mancini et al. (2016). However, unlike there, we hold the environment fixed and do not consider the scenario in which the environment tries to adversely impact the total predictable information captured by the sensor. The proportionality constant α is a conversion factor between information and energy units. It is ultimately set by the type of environment in which the sensor finds itself. It determines the energetic reward for prediction. All said, a sensor that maximizes this objective is predictive of its input and does not dissipate much heat.

An example— $x_l = 0.5$, $x_h = 2.0$, $k_O = k_C = 1.0$, and $n = 2$ —illustrates that roughly 99% of I_{fut} is devoted to capturing instantaneous memory I_{mem} and roughly 25% of βP is devoted to \dot{I}_{np} . That is, the inefficiency in choosing what information to store about the present input contributes substantively to energetic inefficiency. These results hold qualitatively even when the dwell-time densities are lognormal; i.e., heavier-tailed.

Figure 8 shows that increased cooperativity—that is, increases in n —lead to increases in prediction performance, qualitatively in line with Martins and Swain (2011) and Mancini et al. (2016). Additionally, the larger the cooperativity, the higher the fraction $I_{\text{mem}} / I_{\text{fut}}$. Larger cooperativity, however, leads to roughly linear increases in the power consumption and nonpredictive information rate, whereas increases in prediction power take a more sigmoidal shape. The correlation between prediction power and power consumption are qualitatively in line with the results of Mehta and Schwab (2012).

This suggests a preference for intermediate values of cooperativity (e.g., $n \approx 5$) over larger values of cooperativity (e.g., $n \geq 10$). This is qualitatively similar to the results of Tkačik et al. (2009) and Walczak et al. (2010) in that physical constraints force optimal information transmission at intermediate levels of cooperativity. Notice that instantaneous memory and total predictable information sit at around 0.05 bits for Hill factors $n \approx 4$. This is an order of magnitude higher than the instantaneous memory and total predictable information of the large random sensors above.

The sensor's objective function $\alpha I_{\text{fut}} - \beta P$ is optimized by the cooperativity $\hat{n} = \arg \max_n (\alpha I_{\text{fut}} - \beta P)$. As both I_{fut} and βP increase monotonically with n , there is generically only one such cooperativity \hat{n} . At lower α , however, there are two local maxima of the function of n given by $\alpha I_{\text{fut}} - \beta P$, as shown in Fig. 9 (Top). Let us pursue the consequences of this regime dependence.

There are rules for how sensor biochemical parameters adapt to the present environment. If adaptation rules for cooperativity of a Hill molecule increase the total energy budget by gradient descent then, for a range of α , we expect cooperativity to be in either of the two local maxima of $\alpha I_{\text{fut}} - \beta P$ just noted. Let us assume a separation of timescales—namely, that cooperativity adapts much more slowly than the longest environmental timescale and that the sensor does not adapt quickly enough to always find the global maximum in cooperativity space.

Figure 9 (Bottom) shows the optimal cooperativity \hat{n} as a function of conversion factor α . As expected from the presence of two local maxima at lower α , there is a discontinuity [*supercritical bifurcation* (Strogatz 1994)] in the function of α given by $\hat{n} = \arg \max_n (\alpha I_{\text{fut}} - \beta P)$. Thus, initially, if the energetic reward α for prediction increases, the cooperativity discontinuously jumps to a higher \hat{n} at a critical α_h . From there, if one decreases α , optimal cooperativity slowly decreases, but stays high well below α_h , suddenly decreasing to zero cooperativity at the lower value α_l . Thus, there is a substantial hysteresis loop built into the optimal trade-off between energy and sensitivity.

What could be the functional benefit of this hysteresis? Recall that in switching circuits hysteresis is essential to adding stability to a switch's response. Hysteresis stops "race" conditions in which the switch oscillates wildly just as the threshold is passed, amplifying any noise in the control and internal dynamics. In the Hill




Fig. 8 (Top) Hill molecule prediction-related metrics: Instantaneous memory I_{mem} and total prediction power I_{fut} , as functions of n in nats. (Bottom) Hill molecule temperature-normalized dissipation-related metrics: Nonpredictive information rate \dot{I}_{np} and temperature-normalized power βP , both in nats per unit time. Ligand concentrations take one of two values, $x_l = 0.5$ and $x_h = 2.0$, in units of $(k_C/k_O)^{1/n}$. Dwell-time densities are parametrized as $\phi_x(t) = \lambda(x)^2 t e^{-\lambda(x)t}$ with $\lambda(x_l) = 5$ and $\lambda(x_h) = 4$ in units of $1/k_C$. Hill molecule parameters are set to $k_O = k_C = 1$, with varying cooperativity n

molecule, hysteresis is helpful if a memory of past environmental conditions (α) provides insight into future conditions (future α). For example, the environment might shift α suddenly to being low, but there is a replenishment mechanism for the available energy that will soon increase α again. Thus, we see that robustness to environmental noise emerges naturally if a Hill molecule sensor adapts to and anticipates changing external conditions.




Fig. 9 (Top) “Lagrangian” $\alpha I_{\text{fut}} + \beta P$ for several α —that represent the energy rewards of I_{fut} —as a function of cooperativity n . Notice that a particular α singles out a particular optimal \hat{n} . (Bottom) Optimal cooperativity \hat{n} as a function of conversion factor α . Circles (solid blue) are the values of the global maximum \hat{n} and circles (solid green) are local maxima. Arrows indicate the directionality in the hysteresis loop: \hat{n} jumps up at the upper value α_h , if starting at low α , and jumps down at the lower value α_l , if starting at high α

Finally, we note that if both x_l and x_h are larger than 1 or less than 1, the Hill coefficient n that maximizes predictable information is not infinite, but finite, and dependent on the exact values of x_l and x_h . If n is too large, the Hill molecule acts as a switch that opens when x is larger than 1; then, x_l and x_h will have a nearly identical response, and the sensor will be unable to distinguish the two. However, if n is too small, the concentration will have a reduced effect on the switching rates. A middling Hill coefficient leads to maximal predictive power, as shown in Fig. 10, though the maximal predictive power is still an order of magnitude smaller than the predictive power of the sensor when its natural concentration is between the low and




Fig. 10 (Top) Hill molecule prediction-related metrics: Instantaneous memory I_{mem} and total prediction power I_{fut} , as functions of n in nats. (Bottom) Hill molecule temperature-normalized dissipation-related metrics: nonpredictive information rate \dot{I}_{np} and temperature-normalized power βP , both in nats per unit time. Ligand concentrations take one of two values, $x_l = 1.5$ and $x_h = 2.0$, in units of $(k_C/k_O)^{1/n}$. Dwell-time densities are parametrized as $\phi_x(t) = \lambda(x)^2 t e^{-\lambda(x)t}$ with $\lambda(x_l) = 5$ and $\lambda(x_h) = 4$ in units of $1/k_C$. Hill molecule parameters are set to $k_O = k_C = 1$, with varying cooperativity n

high ligand concentrations. Similar results hold for the case that both x_l and x_h are below 1. Hysteresis remains regardless of x_l and x_h .

6 Conclusion

We provided closed-form expressions for instantaneous memory, total predictable information, nonpredictive information rate, and heat dissipated into the reservoir

of interest for a conditionally Markovian channel subject to the broad class of very complex environments described by unifilar hidden semi-Markov input. The information-theoretic quantities among that list are best motivated by Marzen and Crutchfield (2018) and Still et al. (2012).

We showed that determining these closed-form expressions requires knowledge of the environment's causal states. In some cases, one has generated the environment, and so a reasonable model for its dynamics is known. In other cases, the environment in question is well-studied and a reasonable model for its dynamics is known. Causal states can be straightforwardly derived from these models. In other cases, however, reasonable dynamical models might need to be inferred from data. This is quite a difficult problem on which much work has been done, e.g., Refs. Rabiner (1989), Pfau et al. (2011) and Strelioff and Crutchfield (2014).

It is not necessary to use the closed-form expressions in Sect. 4 to calculate prediction and dissipation metrics. Instead, one may numerically estimate them as follows:

1. Generate a realization of the environment, and
2. Simulate transitions between states using the temporal Gillespie algorithm (Vestergaard and Génois 2015).

This approach has its challenges, however. Generating realizations of non-Markovian environments, as considered here, is straightforward but could be quite compute intensive, if one desires accurate simulations. Simulating transitions between Hill molecule states given an environmental realization is straightforward, too.

This highlights the benefits of Sect. 4's closed-form expressions. Their accuracy, however, is limited by that of the numerical algorithms for calculating integrals and finding eigenvectors. Available routines exhibit several weaknesses in accurately evaluating the requisite integrals when, for instance, $\phi_x(\tau)$ is heavy-tailed. For the dwell-time densities considered here, there were no such issues.

All in all, the ease with which the sensor metrics were numerically estimated masks the difficulty of deriving the underlying closed-form expressions in the first place. Appendices B and C give those derivations. That said, leveraging this one-time calculation effort, we provided universal estimators of prediction and dissipation metrics for conditionally Markovian channels. The universality claimed here arises from the fact that unifilar hidden semi-Markov processes are very general memoryful processes that capture highly complex environmental behaviors.

One practical consequence going forward is that analyzing sensor prediction and dissipation no longer requires simulating arbitrarily long trajectories. Instead, one can now validate or invalidate predictive learning rules and sensor designs using the universal estimators. This will also greatly accelerate searching through parameter space for "optimal" (predictive and energy efficient) sensors. In addition, given that the theories of random dynamical systems and of input-dependent dynamical systems are still under development (Arnold 2013), we believe the formulae presented here will eventually lead in those domains to a precise generalization of timescale matching for nonlinear systems (Becker et al. 2015). Finally, we hope to extend these new calculational techniques to nonstationary environments and thereafter to analyze sensory adaptation in more familiar scenarios, as in Sartori et al. (2014) and Ito and Sagawa (2015).

Acknowledgements We thank N. Ay, A. Bell, W. Bialek, S. Dedeo, C. Hillar, I. Nemenman, and S. Still for useful conversations and the Santa Fe Institute and the Telluride Science Research Center for their hospitality during visits, where JPC is an External Faculty and a Board Member, respectively. This material is based upon work supported by, or in part by, the John Templeton Foundation grant 52095, the Foundational Questions Institute grant FQXi-RFP-1609, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract W911NF-13-1-0390. S.E.M. was funded by an MIT Physics of Living Systems Fellowship and AFOSR Grant FA9550-19-1-0411.

Appendix A Revisiting the Thermodynamics of Prediction

For completeness, we review the derivation of Eq. (6). Let x_t represent the input at time t , y_t represent the sensor state at time t , and $E(x, y)$ the system’s energy function for the reservoir of interest. We assume constant temperature. The system’s temperature-normalized nonequilibrium free energy F_{neq} is given by:

$$\beta F_{\text{neq}}[p(y|x)] = \beta \langle E(x, y) \rangle_{p(y|x)} - H[Y|X = x]. \tag{A1}$$

Even if this is not a valid expression for nonequilibrium free energy, the validity of Still et al. (2012)’s derivation only rests on this expression being a Lyapunov function for the dynamics. Intuitively, this corresponds to an assumption that the system reduces its nonequilibrium free energy when the sensor thermalizes to its attached thermal bath. [Accordingly, the β in the above expression refers to the temperature of the sensor when the environment is fixed, indirectly circumventing the difficulty with defining a nonequilibrium temperature (Casas-Vázquez and Jou 2003).] If so, then:

$$\beta F_{\text{neq}}[p(y_t|x_{t+\Delta t})] \geq \beta F_{\text{neq}}[p(y_{t+\Delta t}|x_{t+\Delta t})],$$

giving:

$$\begin{aligned} 0 &\leq \beta F_{\text{neq}}[p(y_t|x_{t+\Delta t})] - \beta F_{\text{neq}}[p(y_{t+\Delta t}|x_{t+\Delta t})] \\ &\leq (\beta \langle E(x_{t+\Delta t}, y_t) \rangle_{p(y_t|x_{t+\Delta t})} - H[Y_t|X_{t+\Delta t} = x_{t+\Delta t}]) - (\beta \langle E(x_{t+\Delta t}, y_{t+\Delta t}) \rangle_{p(y_{t+\Delta t}|x_{t+\Delta t})} - H[Y_{t+\Delta t}|X_{t+\Delta t} = x_{t+\Delta t}]) \\ &\leq \beta \lim_{\Delta t \rightarrow 0} \left(\frac{\langle E(x_{t+\Delta t}, y_t) \rangle_{p(y_t|x_{t+\Delta t})}}{\Delta t} - \frac{\langle E(x_{t+\Delta t}, y_{t+\Delta t}) \rangle_{p(y_{t+\Delta t}|x_{t+\Delta t})}}{\Delta t} \right) \\ &\quad + \lim_{\Delta t \rightarrow 0} \left(\frac{H[Y_{t+\Delta t}|X_{t+\Delta t} = x_{t+\Delta t}]}{\Delta t} - \frac{H[Y_t|X_{t+\Delta t} = x_{t+\Delta t}]}{\Delta t} \right). \end{aligned}$$

Finally, we average over possible environmental realizations, equivalent in nonequilibrium steady states (NESSs) to averages over time, to find:

$$\begin{aligned} 0 &\leq \beta \lim_{\Delta t \rightarrow 0} \frac{\langle E(x_{t+\Delta t}, y_t) \rangle - \langle E(x_{t+\Delta t}, y_{t+\Delta t}) \rangle}{\Delta t} \\ &\quad + \lim_{\Delta t \rightarrow 0} \frac{H[Y_{t+\Delta t}|X_{t+\Delta t}] - H[Y_t|X_{t+\Delta t}]}{\Delta t}. \end{aligned}$$

The former term could be called $-\dot{Q}$ —the negative of the sensor’s heat dissipation rate into the reservoir of interest—and the latter \dot{I}_{lost} —the rate of lost information. And so:

$$0 \leq -\beta \dot{Q} + \dot{I}_{\text{lost}}. \tag{A2}$$

This is valid even outside of NESSs. In a NESS, however, we can invoke stationarity, concluding that:

$$\dot{Q} = -P, \tag{A3}$$

where P is the part of the power extracted by the sensor from the environment that is dissipated as heat into the reservoir of interest. Furthermore, \dot{I}_{lost} reduces to the negative of the nonpredictive information rate, since $H[Y_{t+\Delta t}|X_{t+\Delta t}] = H[Y_t|X_t]$, giving:

$$\dot{I}_{\text{lost}} = \lim_{\Delta t \rightarrow 0} \frac{H[Y_t|X_t] - H[Y_{t+\Delta t}|X_{t+\Delta t}]}{\Delta t}.$$

Calling on a standard information theory identity (Cover and Thomas 2006)— $H[U|V] = H[U] - I[U; V]$ —leads to:

$$\dot{I}_{\text{lost}} = \lim_{\Delta t \rightarrow 0} \frac{I[Y_t; X_{t+\Delta t}] - I[Y_t; X_t]}{\Delta t}.$$

We recognize this as the continuous-time version of the nonpredictive information rate \dot{I}_{np} ; also called the learning rate. Hence, the nonpredictive information rate is the increase in unpredictability of sensor state Y_t given a slightly delayed environmental state:

$$\dot{I}_{\text{np}} = \lim_{\Delta t \rightarrow 0} \frac{H[Y_t|X_{t+\Delta t}] - H[Y_t|X_t]}{\Delta t}. \tag{A4}$$

In a NESS, then:

$$\dot{I}_{\text{lost}} = -\dot{I}_{\text{np}}. \tag{A5}$$

Outside of NESSs, these terms are augmented by the time derivative $dH[Y_t|X_t]/dt$ of the conditional entropy. This leads to the addition of an Landauer-erasure information (Landauer 1961) when integrated.

One of Still et al. (2012)’s main results follows directly from Eqs. (A2), (A3), and (A5):

$$\dot{I}_{\text{np}} = \lim_{\Delta t \rightarrow 0} \frac{I[Y_t; X_t] - I[Y_t; X_{t+\Delta t}]}{\Delta t} \leq \beta P. \tag{A6}$$

In contrast to Still et al. (2012)’s implication, this is true only in a NESS and Eq. (A2) should be used otherwise.

Differences in presentation between the derivation here and that of Still et al. (2012) come from the difference between discrete- and continuous-time formulations. To make this clear, we present a continuous-time formulation of the same result, following Horowitz and Esposito (2014). We start from $\beta F_{\text{neq}}[p(y_{t'}|x_t)]$ being a Lyapunov function in t' :

$$\begin{aligned} 0 &\geq \beta \frac{\partial F_{\text{neq}}[p(y_{t'}|x_t)]}{\partial t'} \\ &= \beta \frac{\partial}{\partial t'} \langle E(x_t, y_{t'}) \rangle_{p(y_{t'}|x_t)} \Big|_{t'=t} - \frac{\partial}{\partial t'} H[Y_{t'}|X_t = x_t] \Big|_{t'=t} \\ &= \beta \left\langle \frac{\partial E(x_t, y_t)}{\partial y_t} \dot{y}_t \right\rangle_{p(y_{t'}|x_t)} - \frac{\partial}{\partial t'} H[Y_{t'}|X_t = x_t] \Big|_{t'=t}. \end{aligned}$$

Next, as before, we average over protocols (or, equivalently in NESS, over time) to find:

$$0 \geq \beta \left\langle \frac{\partial E(x_t, y_t)}{\partial y_t} \dot{y}_t \right\rangle_{p(x_t, y_t)} - \frac{\partial}{\partial t'} H[Y_{t'}|X_t] \Big|_{t'=t}$$

We then recognize $\beta \left\langle \frac{\partial E(x_t, y_t)}{\partial y_t} \dot{y}_t \right\rangle_{p(x_t, y_t)}$ as the temperature-normalized rate of heat dissipation $\beta \dot{Q}$, so that:

$$\beta \dot{Q} \leq \frac{\partial}{\partial t'} H[Y_{t'}|X_t] \Big|_{t'=t}.$$

The quantity on the right-hand side is simply the rate \dot{I}_{lost} of information loss, defined earlier. In NESS, $d\langle E \rangle/dt$ and $dH[Y_t|X_t]/dt$ vanish. As a result, $\beta \dot{Q} + \beta P = 0$ and:

$$\frac{\partial}{\partial t'} H[Y_{t'}|X_t] \Big|_{t'=t} = -\frac{\partial}{\partial t'} H[Y_t|X_{t'}] \Big|_{t'=t},$$

giving:

$$\beta P \geq \frac{\partial}{\partial t'} H[Y_t|X_{t'}] \Big|_{t'=t}. \tag{A7}$$

We recognize this as the continuous-time formulation of Eq. (A4). Again invoking stationarity, $dH[X_t]/dt$ vanishes and so:

$$\beta P \geq -\frac{\partial}{\partial t'} H[X_{t'}; Y_t] \Big|_{t'=t}, \tag{A8}$$

the continuous-time formulation of Eq. (A6). We have, in Eqs. (A4), (A6), (A7), and (A8), four equivalent definitions for the nonpredictive information rate in the NESS limit.

Appendix B Closed-form Expressions for Unifilar Hidden Semi-Markov Environments

To find $\rho(\sigma^+, y)$, we start with the following:

$$\begin{aligned} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y) &= \sum_{g',x',\tau',y'} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g', x', \tau'), Y_t = y') \Pr(\mathcal{S}_t^+ = (g', x', \tau'), Y_t = y'). \end{aligned} \tag{B1}$$

We decompose the transition probability using the lack of feedback as:

$$\begin{aligned} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g', x', \tau'), Y_t = y') &= \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau) | \mathcal{S}_t^+ = (g', x', \tau')) \Pr(Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g', x', \tau'), Y_t = y'). \end{aligned}$$

From the setup, we have:

$$\Pr(Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g', x', \tau'), Y_t = y') = \begin{cases} k_{y' \rightarrow y}(x') \Delta t & y \neq y' \\ 1 - k_{y' \rightarrow y'}(x') \Delta t & y = y' \end{cases}$$

with corrections of $O(\Delta t^2)$.

Now split this into two cases. As long as $\tau > \Delta t$, so that $x = x'$, we have:

$$\Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau) | \mathcal{S}_t^+ = (g', x', \tau')) = \frac{\Phi_g(\tau)}{\Phi_g(\tau')} \delta(\tau - (\tau' + \Delta t)) \delta_{x,x'} \delta_{g,g'}.$$

Then, Eq. (B1) reduces to:

$$\begin{aligned} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y) &= \sum_{y'} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau) | \mathcal{S}_t^+ = (g, x, \tau - \Delta t)) \\ &\quad \times \Pr(Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \\ &\quad \times \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \\ &= \sum_{y' \neq y} \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau) | \mathcal{S}_t^+ = (g, x, \tau - \Delta t)) \\ &\quad \times \Pr(Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \\ &\quad \times \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \\ &\quad + \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau) | \mathcal{S}_t^+ = (g, x, \tau - \Delta t)) \\ &\quad \times \Pr(Y_{t+\Delta t} = y | \mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y) \end{aligned}$$

$$\begin{aligned}
 & \times \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y) \\
 = & \sum_{y' \neq y} \frac{\Phi_g(\tau)}{\Phi_g(\tau - \Delta t)} k_{y' \rightarrow y}(x) \\
 & \times \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \Delta t \\
 & + \frac{\Phi_g(\tau)}{\Phi_g(\tau - \Delta t)} (1 - k_{y \rightarrow y}(x) \Delta t) \\
 & \times \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y), \tag{B2}
 \end{aligned}$$

plus terms of $O(\Delta t^2)$. We Taylor expand $\Phi_g(\tau + \Delta t) = \Phi_g(\tau) - \phi_g(\tau) \Delta t$ to find:

$$\frac{\Phi_g(\tau)}{\Phi_g(\tau - \Delta t)} = 1 - \frac{\phi_g(\tau)}{\Phi_g(\tau)} \Delta t,$$

plus terms of $O(\Delta t^2)$. And, similarly, assuming differentiability, we write:

$$\begin{aligned}
 & \Pr(\mathcal{S}_t^+ = (g, x, \tau - \Delta t), Y_t = y') \\
 = & \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y') - \frac{d}{d\tau} \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y') \Delta t,
 \end{aligned}$$

plus terms of $O(\Delta t^2)$. Substitution into Eq. (B2) then gives:

$$\begin{aligned}
 & \Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y) \\
 = & \left(\sum_{y' \neq y} k_{y' \rightarrow y}(x) \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y') \right) \Delta t + \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y) \\
 & - \frac{d \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y)}{d\tau} \Delta t - \frac{\phi_g(\tau)}{\Phi_g(\tau)} \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y) \Delta t \\
 & - k_{y \rightarrow y}(x) \Pr(\mathcal{S}_t^+ = (g, x, \tau), Y_t = y) \Delta t,
 \end{aligned}$$

plus terms of $O(\Delta t^2)$. For notational ease, we denote:

$$\rho((g, x, \tau), y) := \Pr(\mathcal{S}_t^+ = (x, \tau), Y_t = y),$$

which is equal to $\Pr(\mathcal{S}_{t+\Delta t}^+ = (g, x, \tau), Y_{t+\Delta t} = y)$ since we assumed the system is in a NESS. Then we have:

$$\begin{aligned}
 \rho((g, x, \tau), y) = & \left(\sum_{y' \neq y} k_{y' \rightarrow y}(x) \rho((g, x, \tau), y') \right) \Delta t + \rho((g, x, \tau), y) - \frac{d\rho((g, x, \tau), y)}{d\tau} \Delta t \\
 & - \frac{\phi_g(\tau)}{\Phi_g(\tau)} \rho((g, x, \tau), y) \Delta t - k_{y \rightarrow y}(x) \rho((g, x, \tau), y) \Delta t
 \end{aligned}$$

plus corrections of $O(\Delta t^2)$. We are left equating the coefficient of the $O(\Delta t)$ term to 0:

$$\frac{d\rho((g, x, \tau), y)}{d\tau} = \sum_{y' \neq y} k_{y' \rightarrow y}(x) \rho((g, x, \tau), y') - \frac{\phi_g(\tau)}{\Phi_g(\tau)} \rho((g, x, \tau), y) - k_{y \rightarrow y}(x) \rho((g, x, \tau), y). \tag{B3}$$

Our task is simplified if we separate:

$$\rho((g, x, \tau), y) = p(y|g, x, \tau) \rho(g, x, \tau)$$

and if we recall that:

$$\rho(g, x, \tau) = \mu_g \Phi_g(\tau) p(g) p(x|g).$$

These give:

$$\frac{d\rho(g, x, \tau)}{d\tau} = -\mu_g \phi_g(\tau) p(g) p(x|g). \tag{B4}$$

Plugging Eq. (B4) into Eq. (B3) yields:

$$\begin{aligned} \frac{dp(y|x, \tau)}{d\tau} \rho(g, x, \tau) - \mu_g \phi_g(\tau) p(g) p(x|g) p(y|g, x, \tau) \\ = \sum_{y' \neq y} k_{y' \rightarrow y}(x) \rho(g, x, \tau) p(y'|g, x, \tau) \\ - \frac{\phi_g(\tau)}{\Phi_g(\tau)} \rho(g, x, \tau) p(y|g, x, \tau) \\ - k_{y \rightarrow y}(x) \rho(g, x, \tau) p(y|g, x, \tau), \end{aligned}$$

where we note that:

$$\mu_g \phi_g(\tau) p(g) p(x|g) p(y|g, x, \tau) = \frac{\phi_g(\tau)}{\Phi_g(\tau)} \rho(g, x, \tau) p(y|g, x, \tau).$$

Hence, we are left with:

$$\frac{dp(y|g, x, \tau)}{d\tau} = \sum_{y' \neq y} k_{y' \rightarrow y}(x) p(y'|g, x, \tau) - k_{y \rightarrow y}(x) p(y|g, x, \tau).$$

We can summarize this ordinary differential equation in matrix-vector notation as follows. Let $\mathbf{v}(g, x, \tau)$ be the vector:

$$\mathbf{v}(g, x, \tau) := \begin{pmatrix} p(y_1|g, x, \tau) \\ \vdots \\ p(y_{|Y|}|g, x, \tau). \end{pmatrix}$$

We have:

$$\frac{d\mathbf{v}}{d\tau} = M(x)\mathbf{v},$$

with solution:

$$\mathbf{v}(g, x, \tau) = e^{M(x)\tau}\mathbf{v}(g, x, 0). \tag{B5}$$

The structure of $M(x)$ guarantees that probability is conserved, as long as $\mathbf{1}^\top \mathbf{v}(g, x, 0) = 1$ for all $x \in \mathcal{A}$.

Our next task is to find expressions for $\mathbf{v}(g, x, 0)$. We do this by considering Eq. (B1) in the limit that $\tau < \Delta t$. More straightforwardly, we consider the equation:

$$\rho((g, x, 0), y) = \sum_{g',x'} \int_0^\infty d\tau \frac{\phi_{g'}(\tau)}{\Phi_{g'}(\tau)} \delta_{g,\epsilon+(g',x')} p(x|g) \rho((g', x', \tau), y), \tag{B6}$$

which is based on the following logic. For probability to flow into $\rho((g, x, 0), y)$ from $\rho((g', x', \tau), y')$, we need the dwell time for symbol x' to be exactly τ and for $y' = y$. (The latter comes from the unlikelihood of switching both channel state and input symbol at the same time.) Again decomposing:

$$\begin{aligned} \rho((g', x', \tau), y) &= p(y|g', x', \tau) \rho(g', x', \tau) \\ &= \mu_{g'} \Phi_{g'}(\tau) p(g') p(x'|g') p(y|g', x', \tau) \end{aligned} \tag{B7}$$

and, thus, as a special case:

$$\rho((g, x, 0), y) = p(y|g, x, 0) p(g) p(x|g) \mu_g. \tag{B8}$$

Plugging both Eqs. (B7) and (B8) into Eq. (B6), we find:

$$\begin{aligned} \mu_g p(g) p(x|g) p(y|g, x, 0) &= \sum_{g',x'} \int_0^\infty \mu_{g'} p(g') p(x'|g') \phi_{g'}(\tau) \delta_{g,\epsilon+(g',x')} p(x|g) p(y|g', x', \tau) d\tau \\ \mu_g p(g) p(y|g, x, 0) &= \sum_{g',x'} \int_0^\infty \mu_{g'} p(g') p(x'|g') \phi_{g'}(\tau) \delta_{g,\epsilon+(g',x')} p(y|g', x', \tau) d\tau. \end{aligned}$$

Using Eq. (B5), we see that $p(y|g', x', \tau) = \left(e^{M(x')\tau} \mathbf{v}(g', x', 0) \right)_y$ and $p(y|g, x, 0) = (\mathbf{v}(g, x, 0))_y$. So, we have:

$$\mu_g p(g) \mathbf{v}(g, x, 0) = \sum_{g',x'} \mu_{g'} \delta_{g,\epsilon+(g',x')} p(g') p(x'|g') \left(\int_0^\infty \phi_{g'}(\tau) e^{M(x')\tau} d\tau \right) \mathbf{v}(g', x', 0).$$

If we form the composite vector:

$$\begin{aligned} \mathbf{U} &= \begin{pmatrix} \mathbf{u}(g_1, x_1) \\ \mathbf{u}(g_1, x_2) \\ \vdots \\ \mathbf{u}(g_{|\mathcal{G}|}, x_{|\mathcal{A}|}) \end{pmatrix} \\ &= \begin{pmatrix} \mu_{g_1} p(g_1) \mathbf{v}(g_1, x_1, 0) \\ \vdots \\ \mu_{g_{|\mathcal{G}|}} p(g_{|\mathcal{G}|}) \mathbf{v}(g_{|\mathcal{G}|}, x_{|\mathcal{A}|}, 0) \end{pmatrix} \end{aligned}$$

and the matrix (written in block form) as:

$$\mathbf{C} := \begin{pmatrix} C_{(g_1, x_1) \rightarrow (g_1, x_1)} & C_{(g_1, x_2) \rightarrow (g_1, x_1)} & \cdots \\ C_{(g_1, x_1) \rightarrow (g_1, x_2)} & C_{(g_1, x_2) \rightarrow (g_1, x_2)} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}$$

with:

$$C_{(g', x') \rightarrow (g, x)} = \delta_{g, \epsilon^+(g', x')} p(x' | g') \int_0^\infty \phi_{g'}(t) e^{M(x')t} dt,$$

we then have:

$$\mathbf{U} = \text{eig}_1(\mathbf{C}). \tag{B9}$$

Finally, we must normalize $\mathbf{u}(x)$ appropriately. We do this by recalling that $\mathbf{1}^\top \mathbf{v}(g, x, 0) = 1$, since $\mathbf{v}(g, x, 0)$ is a vector of probabilities. Then we have:

$$\mathbf{u}(g, x) \rightarrow \frac{\mathbf{u}(g, x)}{\mathbf{1}^\top \mathbf{u}(g, x)} \mu_g p(g).$$

for each g, x .

To calculate prediction metrics—i.e., I_{mem} and I_{fut} —we need $p(x, y)$ and $p(y, \sigma^-)$. The former is a marginalization of $p(\sigma^+, y)$ that we just calculated. The second can be calculated via:

$$p(\sigma^-, y) = \sum_{\sigma^+} p(\sigma^- | \sigma^+) p(y, \sigma^+),$$

where:

$$\begin{aligned} p(\sigma^- | \sigma^+) &= p((g_-, x_-, \tau_-) | (g_+, x_+, \tau_+)) \\ &= \delta_{x_+, x_-} p(g_- | g_+, x_+) \mu_{g_+} \phi_{g_+}(\tau_+ + \tau_-). \end{aligned}$$

Hence, we turn our attention to calculating dissipation metrics, for which we only need:

$$\frac{\delta p}{\delta t} = \lim_{\Delta t \rightarrow 0} \frac{\Pr(X_{t+\Delta t} = x, Y_t = y) - \Pr(X_t = x, Y_t = y)}{\Delta t}.$$

Moreover, we can use the Markov chain $Y_t \rightarrow \mathcal{S}_t^+ \rightarrow X_{t+\Delta t}$ to compute it:

$$\Pr(X_{t+\Delta t} = x, Y_t = y) = \sum_{\sigma^+} \Pr(X_{t+\Delta t} = x | \mathcal{S}_t^+ = \sigma^+) \Pr(Y_t = y, \mathcal{S}_t^+ = \sigma^+).$$

We have:

$$\begin{aligned} \Pr(X_{t+\Delta t} = x | \mathcal{S}_t^+ = \sigma^+) &= \Pr(X_{t+\Delta t} = x | \mathcal{S}_t^+ = (g', x', \tau')) \\ &= \begin{cases} \frac{\Phi_{g'}(\tau'+\Delta t)}{\Phi_{g'}(\tau')} & x = x' \\ \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} p(x | \epsilon^+(g', x')) \Delta t & x \neq x'. \end{cases} \end{aligned}$$

This, combined with $p(\sigma^+, y)$, gives:

$$\begin{aligned} \Pr(X_{t+\Delta t} = x, Y_t = y) &= \sum_{g', x' \neq x} \int d\tau' \rho((g', x', \tau'), y) \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} \Delta t p(x | \epsilon^+(g', x')) \\ &\quad + \sum_{g'} \int d\tau' \frac{\Phi_{g'}(\tau'+\Delta t)}{\Phi_{g'}(\tau')} \rho((g', x', \tau'), y) \\ &= \Pr(X_t = x, Y_t = y) \\ &\quad + \Delta t \left(\sum_{g', x' \neq x} \int d\tau' p(x | \epsilon^+(g', x')) \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} \rho((g', x', \tau'), y) \right. \\ &\quad \left. - \sum_{g'} \int d\tau' \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} \rho((g', x, \tau'), y) \right), \end{aligned}$$

correct to $O(\Delta t)$. Recalling that:

$$\begin{aligned} \rho((g', x', \tau'), y) &= \rho(g', x', \tau') p(y | g', x', \tau') \\ &= p(x' | g') \Phi_{g'}(\tau') \left(e^{M(x')\tau'} \mathbf{u}(g', x') \right)_y, \end{aligned}$$

gives:

$$\begin{aligned} \frac{\delta p}{\delta t} &= \lim_{\Delta t \rightarrow 0} \frac{\Pr(X_{t+\Delta t} = x, Y_t = y) - \Pr(X_t = x, Y_t = y)}{\Delta t} \\ &= \sum_{g', x' \neq x} \int d\tau' p(x | \epsilon^+(g', x')) \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} \rho((g', x', \tau'), y) - \sum_{g'} \int d\tau' \frac{\phi_{g'}(\tau')}{\Phi_{g'}(\tau')} \rho((g', x, \tau'), y) \end{aligned} \tag{B10}$$

$$\begin{aligned}
 &= \sum_{g', x' \neq x} \int d\tau' p(x|\epsilon^+(g', x')) p(x'|g') \phi_{g'}(\tau') \left(e^{M(x')\tau'} \mathbf{u}(g', x') \right)_y \\
 &\quad - \sum_{g'} \int d\tau' p(x|g') \phi_{g'}(\tau') \left(e^{M(x)\tau'} \mathbf{u}(g', x) \right)_y.
 \end{aligned}
 \tag{B11}$$

From this, Eqs. (7) and (10) can be used to calculate $\dot{\mathbf{I}}_{np}$ and βP .

Appendix C Specialization to Semi-Markov Input

Up to this point, we wrote expressions for the general case of unifilar hidden semi-Markov environment inputs to the sensor. We now specialize to the semi-Markov input case: the environment’s states are directly observed, not hidden. Not surprisingly, a great simplification ensues: hidden states g are the current emitted symbols x . Recall that, in an abuse of notation, $q(x|x')$ is now the probability of observing symbol x after seeing symbol x' .

Hence, forward-time causal states are given by the pair (x, τ) . The analog of Eq. (B5) is:

$$\mathbf{p}(y|x, \tau) = e^{M(x)\tau} \mathbf{p}(y|x, 0),$$

and we define vectors:

$$\mathbf{u}(x) := \mu_x p(x) \mathbf{p}(y|x, 0).$$

The large vector:

$$\mathbf{U} := \begin{pmatrix} \mathbf{u}(x_1) \\ \vdots \\ \mathbf{u}(x_{|\mathcal{A}|}) \end{pmatrix}$$

is the eigenvector $\text{eig}_1(\mathbf{C})$ of eigenvalue 1 of the matrix:

$$\mathbf{C} = \begin{pmatrix} 0 & q(x_1|x_2) \int_0^\infty \phi_{x_2}(\tau) e^{M(x_2)\tau} d\tau & \dots \\ q(x_2|x_1) \int_0^\infty \phi_{x_1}(\tau) e^{M(x_1)\tau} d\tau & 0 & \dots \\ \vdots & \vdots & \ddots \end{pmatrix},$$

where normalization requires $\mathbf{1}^\top \mathbf{u}(x) = \mu_x p(x)$.

We continue by finding $p(y)$, since from this we obtain $H[Y]$. We do this via straightforward marginalization:

$$p(y) = \sum_{\sigma^+} \rho(\sigma^+, y) = \sum_{\sigma^+} p(y|\sigma^+) \rho(\sigma^+)$$

$$\begin{aligned}
 &= \sum_x \int_0^\infty p(y|x, \tau) \rho(x, \tau) \, d\tau \\
 &= \sum_x \int_0^\infty \left(e^{M(x)\tau} \mathbf{v}(x, 0) \right)_y \mu_x p(x) \Phi_x(\tau) \, d\tau \\
 &= \sum_x \left(\left(\int_0^\infty e^{M(x)\tau} \Phi_x(\tau) \, d\tau \right) \mathbf{u}(x) \right)_y .
 \end{aligned}$$

This implies that:

$$\mathbf{p}(y) = \sum_x \left(\int_0^\infty e^{M(x)\tau} \Phi_x(\tau) \, d\tau \right) \mathbf{u}(x) .$$

From earlier, recall that $\mathbf{u}(x) := \mu_x p(x) \mathbf{p}(y|x, 0)$.

Next, we aim to find $p(x, y)$, again via marginalization:

$$\begin{aligned}
 p(x, y) &= \int_0^\infty \rho((x, \tau), y) \, d\tau \\
 &= \int_0^\infty \mu_x p(x) \Phi_x(\tau) p(y|x, \tau) \, d\tau \\
 &= \int_0^\infty \mu_x p(x) \Phi_x(\tau) \left(e^{M(x)\tau} \mathbf{v}(x, 0) \right)_y \, d\tau \\
 &= \left(\left(\int_0^\infty e^{M(x)\tau} \Phi_x(\tau) \, d\tau \right) \mathbf{u}(x) \right)_y . \tag{C1}
 \end{aligned}$$

From the joint distribution $p(x, y)$, we easily numerically obtain $I[X; Y]$, since $|\mathcal{A}| < \infty$ and $|\mathcal{Y}| < \infty$.

For notational ease, we introduced \mathcal{T}_i in this section as the random variable for the time since last symbol, whose realization is τ . Finally, we require $p(y|\sigma^-)$ to calculate $H[Y|\mathcal{S}^-]$, which we can then combine with $H[Y]$ to get an estimate for I_{fut} . We utilize the Markov chain $Y \rightarrow \mathcal{S}^+ \rightarrow \mathcal{S}^-$, as stated earlier, and so have:

$$\begin{aligned}
 p(y|\sigma^-) &= \sum_{\sigma^+} \rho(y, \sigma^+|\sigma^-) \\
 &= \sum_{\sigma^+} p(y|\sigma^+, \sigma^-) \rho(\sigma^+|\sigma^-) \\
 &= \sum_{\sigma^+} p(y|\sigma^+) \rho(\sigma^+|\sigma^-) .
 \end{aligned}$$

Eq. (B5) gives us $p(y|\sigma^+)$ as:

$$\begin{aligned}
 p(y|\sigma^+) &= p(y|x_+, \tau_+) \\
 &= \left(e^{M(x_+)\tau_+} \mathbf{v}(x_+, 0) \right)_y
 \end{aligned}$$

and Eq. (2) gives us $\rho(\sigma^+|\sigma^-)$ after some manipulation:

$$\begin{aligned} \rho(\sigma^+|\sigma^-) &= \rho((x_+, \tau_+)|(x_-, \tau_-)) \\ &= \delta_{x_+,x_-} \frac{\phi_{x_-}(\tau_+ + \tau_-)}{\Phi_{x_-}(\tau_-)}. \end{aligned}$$

Combining the two equations gives:

$$\begin{aligned} p(y|x_-, \tau_-) &= \sum_{x_+} \int_0^\infty \delta_{x_+,x_-} \frac{\phi_{x_-}(\tau_+ + \tau_-)}{\Phi_{x_-}(\tau_-)} \left(e^{M(x_+)\tau_+} \mathbf{v}(x_+, 0) \right)_y d\tau_+ \\ &= \frac{1}{\Phi_{x_-}(\tau_-)} \left(\left(\int_0^\infty \phi_{x_-}(\tau_+ + \tau_-) e^{M(x_-)\tau_+} d\tau_+ \right) \mathbf{v}(x_-, 0) \right)_y. \end{aligned}$$

From this conditional distribution, we compute $H[Y|S^- = \sigma^-]$, and so $H[Y|S^-] = \langle H[Y|S^- = \sigma^-] \rangle_{\rho(\sigma^-)}$. In more detail, define:

$$D_x(\tau) := \int_0^\infty \phi_x(\tau + s) e^{M(x)s} ds,$$

and we have:

$$\mathbf{p}(y|x_-, \tau_-) = D_{x_-}(\tau_-) \mathbf{u}(x_-) / \mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-).$$

This conditional distribution gives:

$$\begin{aligned} H[Y|X_- = x_-, T_- = \tau_-] &= - \sum_y p(y|x_-, \tau_-) \log p(y|x_-, \tau_-) \\ &= -1^\top \left(\frac{D_{x_-}(\tau_-) \mathbf{u}(x_-)}{\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)} \log \left(\frac{D_{x_-}(\tau_-) \mathbf{u}(x_-)}{\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)} \right) \right) \\ &= - \frac{1}{\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)} \left(1^\top ((D_{x_-}(\tau_-) \mathbf{u}(x_-)) \log(D_{x_-}(\tau_-) \mathbf{u}(x_-))) \right. \\ &\quad \left. - 1^\top (D_{x_-}(\tau_-) \mathbf{u}(x_-)) \log(\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)) \right). \end{aligned}$$

We recognize the factor $\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)$ as $\rho(x_-, \tau_-)$ and so we find that:

$$\begin{aligned} H[Y|X_-, T_-] &= \sum_{x_-} \int_0^\infty \rho(x_-, \tau_-) H[Y|X_- = x_-, T_- = \tau_-] d\tau_- \\ &= - \int_0^\infty \left(\sum_{x_-} 1^\top ((D_{x_-}(\tau_-) \mathbf{u}(x_-)) \log(D_{x_-}(\tau_-) \mathbf{u}(x_-))) \right) d\tau_- \\ &\quad + \int_0^\infty \left(\sum_{x_-} 1^\top D_{x_-}(\tau_-) \mathbf{u}(x_-) \log(\mu_{x_-} p(x_-) \Phi_{x_-}(\tau_-)) \right) d\tau_-. \end{aligned}$$

This, combined with earlier formula for $H[Y]$, gives I_{fut} .

Finally, we wish to find an expression for the nonpredictive information rate \dot{I}_{np} . We review the somewhat compact derivation of $\delta p/\delta t$ in the more general case, specialized for semi-Markov input. This requires finding an expression for $\Pr(Y_t = y, X_{t+\Delta t} = x)$ as an expansion in Δt . We start as usual:

$$\Pr(Y_t = y, X_{t+\Delta t} = x) = \sum_{x'} \int_0^\infty \Pr(Y_t = y, X_{t+\Delta t} = x, X_t = x', \mathcal{T}_t = \tau) d\tau$$

and utilize the Markov chain $Y_t \rightarrow \mathcal{S}_t^+ \rightarrow X_{t+\Delta t}$, giving:

$$\Pr(Y_t = y, X_{t+\Delta t} = x) = \sum_{x'} \int_0^\infty \Pr(Y_t = y|X_t = x', \mathcal{T}_t = \tau) \Pr(X_{t+\Delta t} = x|X_t = x', \mathcal{T}_t = \tau) \rho(x', \tau) d\tau. \tag{C2}$$

We have $\Pr(Y_t = y|X_t = x, \mathcal{T}_t = \tau)$ from Eq. (B5). So, we turn our attention to finding $\Pr(X_{t+\Delta t} = x|X_t = x', \mathcal{T}_t = \tau)$. Some thought reveals that:

$$\Pr(X_{t+\Delta t} = x|X_t = x', \mathcal{T}_t = \tau) = \begin{cases} \Delta t q(x|x') \phi_{x'}(\tau) / \Phi_{x'}(\tau) & x \neq x' \\ \Phi_{x'}(\tau + \Delta t) / \Phi_{x'}(\tau) & x = x' \end{cases}, \tag{C3}$$

plus corrections of $O(\Delta t^2)$. We substitute Eq. (C3) into Eq. (C2) to get:

$$\Pr(Y_t = y, X_{t+\Delta t} = x) = \left(\sum_{x' \neq x} \int_0^\infty \Pr(Y_t = y|X_t = x', \mathcal{T}_t = \tau) q(x|x') \frac{\phi_{x'}(\tau)}{\Phi_{x'}(\tau)} \rho(x', \tau) d\tau \right) \Delta t + \int_0^\infty \Pr(Y_t = y|X_t = x, \mathcal{T}_t = \tau) \frac{\Phi_x(\tau + \Delta t)}{\Phi_x(\tau)} \rho(x, \tau) d\tau,$$

plus corrections of $O(\Delta t^2)$. Recalling:

$$\frac{\Phi_x(\tau + \Delta t)}{\Phi_x(\tau)} = 1 - \frac{\phi_x(\tau)}{\Phi_x(\tau)} \Delta t,$$

plus corrections of $O(\Delta t^2)$, we simplify further:

$$\Pr(Y_t = y, X_{t+\Delta t} = x) = \left(\sum_{x' \neq x} \int_0^\infty \Pr(Y_t = y|X_t = x', \mathcal{T}_t = \tau) q(x|x') \frac{\phi_{x'}(\tau)}{\Phi_{x'}(\tau)} \rho(x', \tau) d\tau \right) \Delta t + \int_0^\infty \Pr(Y_t = y|X_t = x, \mathcal{T}_t = \tau) \rho(x, \tau) d\tau - \left(\int_0^\infty \Pr(Y_t = y|X_t = x, \mathcal{T}_t = \tau) \frac{\phi_x(\tau)}{\Phi_x(\tau)} \rho(x, \tau) d\tau \right) \Delta t,$$

plus $O(\Delta t^2)$ corrections. We notice that:

$$\int_0^\infty \Pr(Y_t = y | X_t = x, \mathcal{T}_t = \tau) \rho(x, \tau) d\tau = \Pr(Y_t = y, X_t = x),$$

so that:

$$\begin{aligned} & \lim_{\Delta t \rightarrow 0} \frac{\Pr(Y_t = y, X_{t+\Delta t} = x) - \Pr(Y_t = y, X_t = x)}{\Delta t} \\ &= \sum_{x' \neq x} \int_0^\infty \Pr(Y_t = y | X_t = x', \mathcal{T}_t = \tau) q(x|x') \frac{\phi_{x'}(\tau)}{\Phi_{x'}(\tau)} \rho(x', \tau) d\tau \\ & \quad - \int_0^\infty \Pr(Y_t = y | X_t = x, \mathcal{T}_t = \tau) \frac{\phi_x(\tau)}{\Phi_x(\tau)} \rho(x, \tau) d\tau. \end{aligned}$$

Substituting Eqs. (B5) and (1) into the above expressions yields:

$$\begin{aligned} & \sum_{x' \neq x} \int_0^\infty \Pr(Y_t = y | X_t = x', \mathcal{T}_t = \tau) q(x|x') \frac{\phi_{x'}(\tau)}{\Phi_{x'}(\tau)} \rho(x', \tau) d\tau \\ &= \sum_{x'} q(x|x') \left(\left(\int_0^\infty \phi_{x'}(\tau) e^{M(x')\tau} d\tau \right) \mathbf{u}(x') \right)_y \end{aligned}$$

and:

$$\int_0^\infty \Pr(Y_t = y | X_t = x, \mathcal{T}_t = \tau) \frac{\phi_x(\tau)}{\Phi_x(\tau)} \rho(x, \tau) d\tau = \left(\left(\int_0^\infty \phi_x(\tau) e^{M(x)\tau} d\tau \right) \mathbf{u}(x) \right)_y,$$

so that we have:

$$\begin{aligned} & \lim_{\Delta t \rightarrow 0} \frac{\Pr(Y_t = y, X_{t+\Delta t} = x) - \Pr(Y_t = y, X_t = x)}{\Delta t} \\ &= \left(\sum_{x'} q(x|x') \left(\int_0^\infty \phi_{x'}(\tau) e^{M(x')\tau} d\tau \right) \mathbf{u}(x') - \left(\int_0^\infty \phi_x(\tau) e^{M(x)\tau} d\tau \right) \mathbf{u}(x) \right)_y. \end{aligned}$$

For notational ease, denote the left-hand side as $\delta p(x, y) / \delta t$. The nonpredictive information rate is given by:

$$\begin{aligned} \dot{I}_{np} &= \lim_{\Delta t \rightarrow 0} \frac{I[X_t; Y_t] - I[X_{t+\Delta t}; Y_t]}{\Delta t} \\ &= \lim_{\Delta t \rightarrow 0} \frac{(\mathbb{H}[X_t] + \mathbb{H}[Y_t] - \mathbb{H}[X_t, Y_t]) - (\mathbb{H}[X_{t+\Delta t}] + \mathbb{H}[Y_t] - \mathbb{H}[X_{t+\Delta t}, Y_t])}{\Delta t} \\ &= \lim_{\Delta t \rightarrow 0} \frac{\mathbb{H}[X_{t+\Delta t}, Y_t] - \mathbb{H}[X_t, Y_t]}{\Delta t}, \end{aligned}$$

where we utilize stationarity to assert $H[X_t] = H[X_{t+\Delta t}]$. Then, correct to $O(\Delta t)$, we have:

$$\begin{aligned} H[X_{t+\Delta t}, Y_t] &= - \sum_{x,y} \left(p(x, y) + \frac{\delta p(x, y)}{\delta t} \Delta t \right) \log \left(p(x, y) + \frac{\delta p(x, y)}{\delta t} \Delta t \right) \\ &= - \sum_{x,y} p(x, y) \log p(x, y) - \sum_{x,y} p(x, y) \frac{\delta p(x, y)/\delta t}{p(x, y)} \Delta t \\ &\quad - \sum_{x,y} \frac{\delta p(x, y)}{\delta t} \log p(x, y) \Delta t \\ &= H[X_t; Y_t] - \sum_{x,y} \frac{\delta p(x, y)}{\delta t} \log p(x, y) \Delta t, \end{aligned}$$

which implies:

$$\dot{I}_{np} = \sum_{x,y} \frac{\delta p(x, y)}{\delta t} \log p(x, y),$$

with:

$$\frac{\delta p(x, y)}{\delta t} = \left(\sum_{x'} q(x|x') \left(\int_0^\infty \Phi_{x'}(\tau) e^{M(x')\tau} d\tau \right) \mathbf{u}(x') - \left(\int_0^\infty \phi_x(\tau) e^{M(x)\tau} d\tau \right) \mathbf{u}(x) \right)_y \quad (C4)$$

and $p(x, y)$ given in Eq. (C1).

References

- Aghamohammadi C, Crutchfield JP (2017) Thermodynamics of random number generation. *Phys Rev E* 95(6):062139
- Arnold L (2013) *Random dynamical systems*. Springer, New York
- Barato AC, Hartich D, Seifert U (2014) Efficiency of cellular information processing. *New J Phys* 16(10):103024
- Becker NB, Mugler A, ten Wolde PR (2015) Optimal prediction by cellular signaling networks. *Phys Rev Lett* 115(25):258103
- Bennett CH (1982) The thermodynamics of computation: a review. *Int J Theor Phys* 21(12):905–940
- Bialek W, Nemenman I, Tishby N (2001) Predictability, complexity, and learning. *Neural Comput* 13:2409–2463
- Bo S, Del Giudice M, Celani A (2015) Thermodynamic limits to information harvesting by sensory systems. *J Stat Mech Theory Exp* 2015(1):P01014
- Boyd AB, Crutchfield JP (2016) Maxwell demon dynamics: Deterministic chaos, the Szilard map, and the intelligence of thermodynamic systems. *Phys Rev Lett* 116:190601
- Boyd AB, Mandal D, Crutchfield JP (2016) Leveraging environmental correlations: the thermodynamics of requisite variety. *J Stat Phys* 167(6):1555–1585
- Boyd AB, Mandal D, Crutchfield JP (2017) Correlation-powered information engines and the thermodynamics of self-correction. *Phys Rev E* 95(1):012152
- Boyd AB, Mandal D, Riechers PM, Crutchfield JP (2017) Transient dissipation and structural costs of physical information transduction. *Phys Rev Lett* 118:220602
- Brittain RA, Jones NS, Ouldrige TE (2017) What we learn from the learning rate. *J Stat Mech* 2017:063502

- Brodu N (2011) Reconstruction of ϵ -machines in predictive frameworks and decisional states. *Adv Complex Syst* 14(05):761–794
- Casas-Vázquez J, Jou D (2003) Temperature in non-equilibrium states: a review of open problems and current proposals. *Rep Prog Phys* 66(11):1937
- Chapman A, Miyake A (2015) How an autonomous quantum Maxwell demon can harness correlated information. *Phys Rev E* 92(6):062125
- Chklovskii DB, Koulakov AA (2004) Maps in the brain: what can we learn from them? *Annu Rev Neurosci* 27:369–392
- Cover TM, Thomas JA (1991) *Elements of Information Theory*. Wiley-Interscience, New York
- Cover TM, Thomas JA (2006) *Elements of information theory*, 2nd edn. Wiley-Interscience, New York
- Creutzig F, Sprekeler H (2008) Predictive coding and the slowness principle: an information-theoretic approach. *Neural Comput* 20(4):1026–1041
- Creutzig F, Globerson A, Tishby N (2009) Past-future information bottleneck in dynamical systems. *Phys Rev E* 79(4):041925
- Crutchfield JP (1994) The calculi of emergence: computation, dynamics, and induction. *Phys D* 75:11–54
- Crutchfield JP, Mitchell M (1995) The evolution of emergent computation. *Proc Natl Acad Sci* 92:10742–10746
- Crutchfield JP, Young K (1989) Inferring statistical complexity. *Phys Rev Lett* 63:105–108
- Crutchfield JP, Ellison CJ, Mahoney JR (2009) Time's barbed arrow: irreversibility, crypticity, and stored information. *Phys Rev Lett* 103(9):094101
- Das SG, Rao M, Iyengar G (2017) Universal lower bound on the free-energy cost of molecular measurements. *Phys Rev E* 95(6):062410
- Deffner S, Jarzynski C (2013) Information processing and the second law of thermodynamics: an inclusive Hamiltonian approach. *Phys Rev X* 3(4):041003
- Gillespie DT (1976) A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comput Phys* 22(4):403–434
- Goldt S, Seifert U (2017) Stochastic thermodynamics of learning. *Phys Rev Lett* 118(1):010601
- Govern CC, ten Wolde PR (2014) Energy dissipation and noise correlations in biochemical sensing. *Phys Rev Lett* 113(25):258102
- Hartich D, Barato AC, Seifert U (2014) Stochastic thermodynamics of bipartite systems: transfer entropy inequalities and a Maxwell's demon interpretation. *J Stat Mech Theory Exp* 2014(2):P02016
- Hartich D, Barato AC, Seifert U (2016) Sensory capacity: an information theoretical measure of the performance of a sensor. *Phys Rev E* 93(2):022116
- Hasenstaub A, Otte S, Callaway E, Sejnowski TJ (2010) Metabolic cost as a unifying principle governing neuronal biophysics. *Proc Natl Acad Sci USA* 107(27):12329–12334
- Hinczewski M, Thirumalai D (2014) Cellular signaling networks function as generalized Wiener-Kolmogorov filters to suppress noise. *Phys Rev X* 4(4):041017
- Horowitz JM, Esposito M (2014) Thermodynamics with continuous information flow. *Phys Rev X* 4:031015
- Horowitz JM, Sagawa T, Parrondo JMR (2013) Imitating chemical motors with optimal information motors. *Phys Rev Lett* 111(1):010602
- Ito S, Sagawa T (2013) Information thermodynamics on causal networks. *Phys Rev Lett* 111(18):180603
- Ito S, Sagawa T (2015) Maxwell's demon in biochemical signal transduction with feedback loop. *Nat Commun* 6:7498
- Izhikevich EM (2007) *Dynamical systems in neuroscience*. MIT press, Cambridge
- Jaeger H (2001) Short Term Memory in Echo State Networks, vol 5. GMD-Forschungszentrum Informationstechnik
- James RG, Ellison CJ, Crutchfield JP (2011) Anatomy of a bit: information in a time series observation. *CHAOS* 21(3):037109
- James RG, Barnett N, Crutchfield JP (2016) Information flows? A critique of transfer entropies. *Phys Rev Lett* 116(23):238701
- Lahiri S, Sohl-Dickstein J, Ganguli S (2016) A universal tradeoff between power, precision and speed in physical communication. [arXiv:1603.07758](https://arxiv.org/abs/1603.07758)
- Lan G, Sartori P, Neumann S, Sourjik V, Tu Y (2012) The energy-speed-accuracy trade-off in sensory adaptation. *Nat Phys* 8(5):422–428
- Landauer R (1961) Irreversibility and heat generation in the computing process. *IBM J Res Dev* 5(3):183–191

- Lang AH, Fisher CK, Mora T, Mehta P (2014) Thermodynamics of statistical inference by cells. *Phys Rev Lett* 113(14):148103
- Little DY, Sommer FT (2014) Learning and exploration in action-perception loops. Closing the loop around neural systems, p 295
- Littman ML, Sutton RS, Singh SP (2001) Predictive representations of state. In: *NIPS*, vol 14, pp 1555–1561
- Löhr W (2010) Models of discrete-time stochastic processes and associated complexity measures. Ph.D. thesis, University of Leipzig
- Löhr W (2012) Predictive models and generative complexity. *J Syst Sci Complex* 25:30–45
- Mancini F, Marsili M, Walczak AM (2016) Trade-offs in delayed information transmission in biochemical networks. *J Stat Phys* 162(5):1088–1129
- Mandal D, Jarzynski C (2012) Work and information processing in a solvable model of Maxwell's demon. *Proc Natl Acad Sci USA* 109(29):11641–11645
- Martins BMC, Swain PS (2011) Trade-offs and constraints in allosteric sensing. *PLoS Comput Biol* 7(11):e1002261
- Marzen SE (2017) Difference between memory and prediction in linear recurrent networks. *Phys Rev E* 96(3):032308
- Marzen S (2018) Infinitely large, randomly wired sensors cannot predict their input unless they are close to deterministic. *PLoS One*
- Marzen SE, Crutchfield JP (2016) Predictive rate-distortion for infinite-order Markov processes. *J Stat Phys* 163(6):1312–1338
- Marzen SE, Crutchfield JP (2017a) Structure and randomness of continuous-time discrete-event processes. *J Stat Phys* 169(2):303–315
- Marzen SE, Crutchfield JP (2017b) Nearly maximally predictive features and their dimensions. *Phys Rev E* 95(5):051301(R)
- Marzen SE, Crutchfield JP (2018) Optimized bacteria are environmental prediction engines. *Phys Rev E* 98:012408
- Marzen SE, DeDeo S (2016) Weak universality in sensory tradeoffs. *Phys Rev E* 94(6):060101
- Marzen S, Garcia HG, Phillips R (2013) Statistical mechanics of Monod–Wyman–Changeux (MWC) models. *J Mol Biol* 425(9):1433–1460
- Maxwell JC (1888) *Theory of heat*, 9th edn. Longmans Green and Co, London
- McGrath T, Jones NS, ten Wolde PR, Ouldridge TE (2017) Biochemical machines for the interconversion of mutual information and work. *Phys Rev Lett* 118(2):028101
- Mehta P, Schwab DJ (2012) Energetic costs of cellular computation. *Proc Natl Acad Sci USA* 109(44):17978–17982
- Nemenman I, Shafee F, Bialek W (2002) Entropy and inference, revisited. In: *Advances in neural information processing systems*, pp 471–478
- Palmer SE, Marre O, Berry MJ, Bialek W (2015) Predictive information in a sensory population. *Proc Natl Acad Sci USA* 112(22):6908–6913
- Parrondo JMR, Horowitz JM, Sagawa T (2015) Thermodynamics of information. *Nat Phys* 11(2):131–139
- Pfau D, Bartlett N, Wood F (2011) Probabilistic deterministic infinite automata. In: *Advances in neural information processing systems*, MIT Press, pp 1930–1938
- Rabiner LR (1989) A tutorial on hidden Markov models and selected applications. *IEEE Proc* 77:257
- Rao RPN, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2(1):79–87
- Sartori P, Granger L, Lee Fan CF, Horowitz JM (2014) Thermodynamic costs of information processing in sensory adaptation. *PLoS Comput Biol* 10(12):e1003974
- Shalizi CR, Crutchfield JP (2001) Computational mechanics: pattern and prediction, structure and simplicity. *J Stat Phys* 104:817–879
- Spinney RE, Lizier JT, Prokopenko M (2018) Entropy balance and information processing in bipartite and nonbipartite composite systems. *Phys Rev E* 98(3):032141
- Still S (2009) Information-theoretic approach to interactive learning. *EuroPhys Lett* 85:28005
- Still S, Crutchfield JP, Ellison CJ (2010) Optimal causal inference: estimating stored information and approximating causal architecture. *CHAOS* 20(3):037111
- Still S, Sivak DA, Bell AJ, Crooks GE (2012) Thermodynamics of prediction. *Phys Rev Lett* 109:120604
- Strelhoff CC, Crutchfield JP (2014) Bayesian structural inference for hidden processes. *Phys Rev E* 89:042119

- Strogatz SH (1994) *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Addison-Wesley, Reading
- Sutton RS, Barto AG (1998) *Reinforcement learning: an introduction*. MIT Press, Cambridge
- Szilard L (1929) On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. *Z Phys* 53:840–856
- Tishby N, Polani D (2011) *Information theory of decisions and actions. Perception-action cycle*. Springer, New York, pp 601–636
- Tkačik G, Walczak AM, Bialek W (2009) Optimizing information flow in small genetic networks. *Phys Rev E* 80(3):031920
- Toyabe S, Sagawa T, Ueda M, Muneyuki E, Sano M (2010) Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality. *Nat Phys* 6(12):988
- Travers N, Crutchfield JP (2014) Equivalence of history and generator ϵ -machines. *Phys Rev E*, page in press, SFI Working Paper 11-11-051; [arXiv:1111.4500](https://arxiv.org/abs/1111.4500) [math.PR]
- Van den Broeck C, Esposito M (2015) Ensemble and trajectory thermodynamics: a brief introduction. *Phys A Stat Mech Appl* 418:6–16
- Vestergaard CL, Génois M (2015) Temporal Gillespie algorithm: fast simulation of contagion processes on time-varying networks. *PLoS Comput Biol* 11(10):e1004579
- Victor JD (2002) Binless strategies for estimation of information from neural data. *Phys Rev E* 66(5):051903
- Walczak AM, Tkačik G, Bialek W (2010) Optimizing information flow in small genetic networks II. Feed-forward interactions. *Phys Rev E* 81(4):041905
- Yeung RW (1991) A new outlook on Shannon's information measures. *IEEE Trans Info Theory* 37(3):466–474

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.