Synchronizing to the Environment: Information Theoretic Constraints on Agent Learning

James P. Crutchfield
Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501
Electronic Address: chaos@santafe.edu

David P. Feldman
College of the Atlantic, 105 Eden St., Bar Harbor, ME 04609
and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501
Electronic Address: dpf@santafe.edu

(July 6, 2001)

Using an information-theoretic framework, we examine how an intelligent agent, given an accurate model of its environment, synchronizes to the environment—i.e., comes to know in which state the environment is. We show that the total uncertainty experienced by the agent during the process is closely related to the transient information, a new quantity that captures the manner in which the environment's entropy growth curve converges to its asymptotic form. We also discuss how an agent's estimates of its environment's structural properties are related to its estimate of the environment entropy rate. If structural properties are ignored, the missed regularities are converted to apparent randomness. Conversely, using representations that assume too much memory results in false predictability.

PACS: 02.50.Ey 05.45.-a 05.45.Tp 89.75.Kd Santa Fe Institute Working Paper 01-03-020

I. INTRODUCTION: AGENTS AND ENVIRONMENTS

The question of how an intelligent agent learns about its environment arises in one fashion or another in many disciplines—such as, economics [1,2], social psychology [3,4], collective cognition [5,6], distributed computing [7,8], automata theory [9], and reinforcement learning [10]. For our purposes here, *intelligent agent* simply refers to an observer that actively builds internal models of its environment using available sensory stimuli and takes action based on these models. This terminology follows the basic framework laid out in the field of reinforcement learning [10]. Here, however, we use an information-theoretic approach to examine how an agent learns about its environment. In so doing, we consider two scenarios—learning about randomness and structure in the environment and synchronizing to the environment's hidden states.

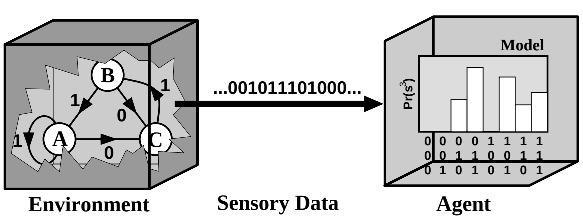

FIG. 1. The Learning Channel: The internal states $\{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$ of the system are reflected, only indirectly, in the observed measurement of 1s and 0s. An agent works with this impoverished data to build a model of the underlying system. After Ref. [11].

A. The Learning Channel

The first of these scenarios, an adaptation of Shannon's communication channel [12], is illustrated in Fig. 1. We assume that there is an *environment* (source or process) that produces a *sensory data stream* (message)—a string of symbols drawn from a finite alphabet ($\mathcal{A}$). The task for the *agent* (receiver or observer) is to estimate the probability distribution of sequences and, thereby, estimate how random the environment is. In this scenario, we assume that the agent does not know the environment's structure or internal dynamics; the range of the environment's states and their transition structure are hidden from the agent. Since the agent does not have direct access to the environment's internal, hidden states, we picture instead that the agent simply collects blocks of measurements from the data stream and stores the block probabilities in a histogram. This histogram functions for the agent as an internal model of the environment. In this scheme, the agent can estimate, to arbitrary accuracy, the probability of measurement sequences by observing for arbitrary lengths of time.

In the particular case illustrated in Fig. 1, the environment is a three-state deterministic finite automaton. However, the agent does not see the internal states $\{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$. Instead, it has access only to the measurement symbols $\mathcal{A} = \{0, 1\}$ generated on state-to-state transitions by the hidden automaton. The environment depicted in Fig. 1 belongs to the class of stochastic processes known as *hidden Markov models*. The transitions from internal state to internal state are Markovian, in that the probability of a given transition depends only upon which state the process is currently in. However, these internal states are not seen by the agent—hence the name *hidden* Markov model [13,14] is often used to describe this type of environment. (In general, however, we do not require that the environment be hidden Markovian.)

B. Synchronization

In the second scenario we assume that the agent already has a correct, finite model of the environment in hand before encountering the environment and making measurements. The issue, then, is how difficult is it for the agent, using sequences of environmental observations, to determine in which hidden state the environment is.

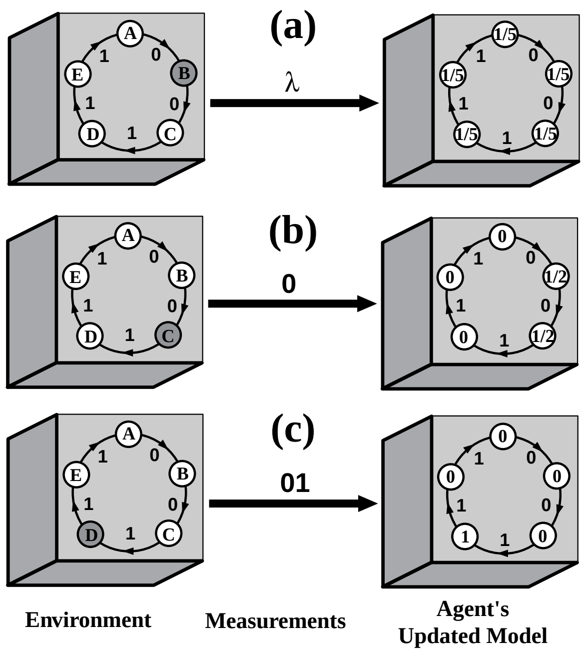

FIG. 2. Synchronizing to the Environment: The agent has an accurate (5-state) model of the environment. In (a), the agent has yet to make an observation and so can assume in this case that all environmental states are equally likely. This is denoted by the fractions inscribed in the agent's model states. In (b), one observation has been made, a 0, and the agent adjusts its guess of the state probabilities accordingly. Having seen a 0, the environment can only be in one of two hidden states. In (c), sufficient measurement information has been gathered so that the agent is synchronized. Having observed the sequence 01, the agent is certain in which state (**D**) the environment is.

An example of this scenario is illustrated in Fig. 2. Imagine that you have planned a vacation to a remote, deserted island where the weather is always three days of sun followed by two days of rain. (Sunny days are encoded in the figure by a 1; rainy days by a 0.) Thus, you, the intelligent agent, have an accurate model of the environment. Having not yet arrived on the island, the least biased inference is that each day in the weather pattern is equally likely and that the chance of rain on your arrival is 40%. (See Fig. 2(a).) If, after landing on the island, you observe that it is raining, you realize that only two days out of the five-day cycle are compatible with that information and so infer a new distribution of the likelihood of the environment being in only two states. Your estimate now is that tomorrow there is a 50% chance of rain. (See Fig. 2(b).) The next day, as it turns out, no matter whether it is sunny or rainy, you will be certain in which state the environment is. (See Fig. 2(c) for the inference that follows from the day being sunny.) We call this condition of model-state certainty being *synchronized*, which will be defined more carefully below. One benefit of being synchronized is that, unlike the previous days, from today forward, you can accurately (and exactly) predict the weather.

In this and related scenarios, several questions naturally arise about the relationship between the agent and its environment. For one, we might be interested in knowing how many observations, on average, must be made before the agent is synchronized. We might also wonder *how* uncertain the environment appears to the agent while it is synchronizing to it. This issue is of relevance, for example, if the agent is compelled to act before synchronization; perhaps it must make decisions—what to wear, what vacation activities to plan—even though it is not fully synchronized. We analyze this kind of uncertainty in Sec. IV.

Questions concerning synchronization have also received considerable attention in other domains. For example, Refs. [15–18] look at schemes that allow an agent to quickly determine the phase of a long periodic sequence. Such schemes are central to a range of communications and engineering applications. More recently, LeBaron suggested that, in the context of an autonomous-agent stock market, understanding the manner in which the agents synchronize to their background environment is essential to understanding market dynamics [19].

II. INFERRING RANDOMNESS

We will revisit the two scenarios below. Before doing so, however, we review several information theoretic measures of unpredictability, structure, and synchronization.

Let $\Pr(s^L)$ denote the probability distribution over blocks $s^L = s_0, s_1, \ldots, s_{L-1}$ of L consecutive environment observations, $s_i \in \mathcal{A}$. Then the *total Shannon entropy* of these L consecutive measurements is defined to be:

$$H(L) \equiv - \sum_{s^L \in \mathcal{A}^L} \Pr(s^L) \log_2 \Pr(s^L) \; , \qquad (1)$$

where $L > 0$. The sum runs over all possible blocks of L consecutive symbols. The units of $H(L)$ are *bits*. The entropy $H(L)$ measures the uncertainty associated with sequences of length L. For a more detailed discussion of the Shannon entropy and related information theoretic quantities, see, e.g., Ref. [20].

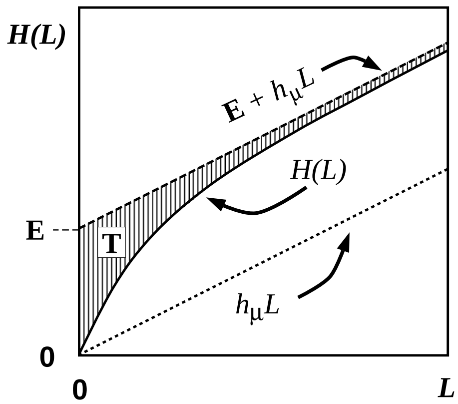

FIG. 3. Total Shannon entropy growth for a finite-memory information source: a schematic plot of $H(L)$ versus L. $H(L)$ increases monotonically and asymptotes to the line $\mathbf{E} + h_\mu L$, where $\mathbf{E}$ is the excess entropy and h_μ is the source entropy rate. The shaded area is the transient information $\mathbf{T}$. After Ref. [21].

The *environment entropy rate* h_μ is the rate of increase with respect to L of the total Shannon entropy in the large-L limit:

$$h_\mu \equiv \lim_{L\to\infty} \frac{H(L)}{L} \ , \tag{2}$$

The units of h_μ are *bits/symbol*, and μ denotes the measure over infinite sequences that induces the L-block joint distribution $\Pr(s^L)$. Alternatively, one can define a finite-L approximation to h_μ,

$$h_\mu(L) \ = \ H(L) - H(L-1) \ , \tag{3}$$

$$\qquad = \ H[S_L | S_{L-1} S_{L-2} \ldots S_0] \ , \tag{4}$$

where $H[X|Y]$ is the entropy of the random variable X conditioned on the random variable Y:

$$H[X|Y] \ = \ \sum_{x,y} \Pr(x,y) \log_2 \Pr(\mathrm{x|y}) \ . \tag{5}$$

One can then show [20] that $h_\mu = \lim_{L\to\infty} h_\mu(L)$.

Thus, the entropy rate h_μ is the uncertainty of a single measurement, given that statistics over infinitely long blocks of measurements have been taken into account. In this sense the entropy rate quantifies the irreducible randomness in the sequences observed by the agent—the randomness that persists even after the agent accounts for statistics over longer and longer blocks of observations.

III. INFERRING MEMORY

Having looked at length-L sequences, an agent can estimate the true environment randomness h_μ by calculating $h_\mu(L)$, defined in Eq. (3). With enough sensory data it can get good approximations to h_μ by using long sequences. But what if the agent has insufficient resources to allow this? To answer this we must determine how the estimates $h_\mu(L)$ converge to h_μ. One measure of convergence is provided by the *excess entropy* $\mathbf{E}$:

$$\mathbf{E} \ \equiv \ \sum_{L=1}^{\infty} [h_\mu(L) - h_\mu] \ . \tag{6}$$

The units of $\mathbf{E}$ are *bits*. The excess entropy is not a new quantity; it was first introduced almost two decades ago and sometimes goes by the names "stored information," "predictive information," and "effective measure complexity" [22–26]. For recent reviews see [21,27,26].

The excess entropy $\mathbf{E}$ can also be given a direct geometric interpretation. As is well known—see, e.g., Refs. [23,24,26,28], the excess entropy is the subextensive part of $H(L)$: that is,

$$\mathbf{E} = \lim_{L\to\infty} [H(L) - h_\mu L] \ . \tag{7}$$

When the limit exists, this implies the following asymptotic form for entropy growth:

$$H(L) \ \sim \ \mathbf{E} + h_\mu L \, , \ \text{as} \ L \to \infty \ . \tag{8}$$

Thus, we see that $\mathbf{E}$ is the $L = 0$ intercept of the linear function Eq. (8) to which $H(L)$ asymptotes. This observation is shown graphically in Fig. 3.

Fig. 3 illustrates the relationships between the growth of the total entropy $H(L)$, its rate of growth h_μ, and excess entropy $\mathbf{E}$. It is clear geometrically from the figure that $\mathbf{E}$ is one quantity that measures the convergence of $h_\mu(L)$ and so it plays a role in how an agent comes to know how random its environment is. But what property does $\mathbf{E}$ quantify? The length-L approximation $h_\mu(L)$ overestimates the entropy rate h_μ at finite L by an amount $h_\mu(L) - h_\mu$. This difference measures how much more random single measurements appear using the finite L-block statistics than the statistics of infinite sequences. In other words, this excess randomness tells us how much additional information must be gained from the environment in order to reveal the actual per-symbol uncertainty h_μ. Thus, we can think of the difference $h_\mu(L) - h_\mu$ as the redundancy per symbol in length-L sequences: that portion of information-carrying capacity in the L-blocks which is not actually random, but is due instead to correlations. The excess entropy $\mathbf{E}$, then, is the total amount of this redundancy and, as such, a measure of one type of memory intrinsic to an environment.

Another way to understand excess entropy is through its expression as a type of mutual information. One can show [23,24] that the excess entropy is the mutual information between the past and the future:

$$\mathbf{E} = \lim_{L\to\infty} I[s_0 s_1 \cdots s_{2L-1}; s_{2L} s_{2L+1} s_{2L-1}] , \qquad (9)$$

when the limit exists. The mutual information [20] is the reduction in the entropy of one variable due to knowledge of another; $I[X;Y] \equiv H[X] - H[X|Y]$. The variable Y carries information about X to the extent to which, on average, knowledge of Y reduces the uncertainty about X.

Eq. (9) says that $\mathbf{E}$ measures the extent to which observations of the past provide information about the future environment behavior. This information can be used to predict the environment's future behavior, but typically more than $\mathbf{E}$ bits of information are required for optimal prediction. For a discussion of the subtleties associated with interpreting $\mathbf{E}$ and also the limitations of using $\mathbf{E}$, see Ref. [27]. Due to these limitations, we interpret $\mathbf{E}$ as the amount of *apparent* memory of the environment and do not use the descriptive phrases quoted above.

Eq. (9) also shows that $\mathbf{E}$ can be interpreted as the *cost of amnesia*: If an agent suddenly loses track of its environment, so that it cannot be predicted at an error level determined by the entropy rate h_μ, then the environment suddenly appears more random by a total of $\mathbf{E}$ bits.

IV. MEASURING SYNCHRONIZATION

A. Transient Information

With this review of measuring randomness and the environment's apparent memory, we are now in a position to address the questions raised in Sec. I B: If an agent has a correct model of the environment in hand, how uncertain is it while it makes its initial observations and synchronizes to the environment? We begin by first defining synchronization more precisely.

For finite-memory ($\mathbf{E} < \infty$) environments, $H(L)$ scales as $\mathbf{E} + h_\mu L$ for large L, Eq. (8). When this scaling form is attained, we say that the agent is *synchronized* to the environment. In other words, when

$$\mathbf{T}(L) \equiv \mathbf{E} + h_\mu L - H(L) = 0 , \qquad (10)$$

we say the agent is synchronized at length-L sequences. The quantity $\mathbf{T}(L)$ provides a measure of the agent's departure from synchronization. Note that $\mathbf{T}(L) \geq 0$.

We now define the the *transient information* $\mathbf{T}$:

$$\mathbf{T} \equiv \sum_{L=0}^{\infty} \mathbf{T}(L) = \sum_{L=0}^{\infty} \left[\mathbf{E} + h_\mu L - H(L)\right] . \qquad (11)$$

Note that the units of $\mathbf{T}$ are *bits* $\times$ *symbols*. In contrast with $\mathbf{E}$, the transient information is a new quantity, recently introduced by us in Ref. [21].

The environment's transient information, being a sum of the agent's finite-L departures from synchronization, captures how difficult it is for an agent to synchronize to that environment. We refer to $\mathbf{T}$ as *transient* since during synchronization the agent's prediction probabilities change, stabilizing only after it has collected a sufficient number of observations.

B. Synchronizing to Markovian Environments

To ground our interpretation, we can establish a direct relation between the transient information $\mathbf{T}$ and the amount of information required for synchronization to order-R Markovian environments. An environment is order-R Markovian if its states are directly observable and the probability of the next state depends only upon the values of the previous R states. The environment depicted in Fig. 1 is not Markovian, since the states are not directly observable. In the context of dynamical systems, a restriction to Markov processes might seem quite limiting. However, in the field of reinforcement learning, such restrictions are fairly common [10]; as they are in statistical mechanics [29].

Assume that the agent has a correct model $\mathcal{M} = \{\mathcal{V}, T\}$ of the environment, where $\mathcal{V}$ is a set of states and T is the rule governing transitions between states. The task for the agent is to make observations and determine the state $v \in \mathcal{V}$ of the environment. This is exactly the process depicted in Figs. 2(a)-(c). Once the agent knows with certainty the current state, it is *synchronized* to the environment, and the average per-symbol uncertainty is exactly h_μ.

The agent's knowledge of $\mathcal{V}$ is given by a distribution over the states $v \in \mathcal{V}$. Let $\Pr(v|s^L, \mathcal{M})$ denote the distribution over $\mathcal{V}$, given that the particular sequence s^L has been observed and the agent has internal model $\mathcal{M}$. The entropy of this distribution measures the agent's average uncertainty in inferring $v \in \mathcal{V}$. Averaging this uncertainty over the possible length-L sequences, we obtain the *average agent-environment uncertainty*:

$$\mathcal{H}(L) \equiv -\sum_{s^L} \Pr(s^L) \sum_{v\in\mathcal{V}} \Pr(v|s^L, \mathcal{M}) \log_2 \Pr(v|s^L, \mathcal{M}) . \qquad (12)$$

The quantity $\mathcal{H}(L)$ can be used as a criterion for synchronization. The agent is synchronized to the environment when $\mathcal{H}(L) = 0$—that is, when the agent is completely certain about the state $v \in \mathcal{V}$ of the mechanism generating the sequence. When the condition in Eq. (10) is met, $\mathcal{H}(L) = 0$, and the uncertainty associated with the prediction of the next state is exactly h_μ.

However, while the agent is still unsynchronized, $\mathcal{H}(L) > 0$. We refer to the total average uncertainty experienced by an agent during the synchronization process as the *synchronization information* $\mathbf{S}$:

$$\mathbf{S} \equiv \sum_{L=0}^{\infty} \mathcal{H}(L) . \qquad (13)$$

The synchronization information measures the total agent-environment uncertainty $\mathcal{H}(L)$ experienced by an agent during synchronization. If the agent is in a position where it must take immediate action, it does not have the option of waiting for synchronization. In this circumstance, the synchronization information $\mathbf{S}$ provides an information-theoretic (average-case) measure of the error incurred by the agent during the synchronization process.

C. Synchronization and Transient Informations

For an order-R Markovian environment, one can establish a direct relationship between the synchronization information $\mathbf{S}$ and the transient information $\mathbf{T}$:

$$\mathbf{S} = \mathbf{T} + \frac{1}{2} R(R+1) h_\mu \ . \tag{14}$$

A detailed proof of this result was given by us in Ref. [21]. Here, we sketch the key pieces in the argument.

First, note that for $L \geq R$, $H(L) = \mathbf{E} + h_\mu L$, and thus $\mathbf{T}(L) = 0$ and $\mathcal{H}(L) = 0$, for all $L \geq R$. Second, observe that $\mathcal{H}(L) = H(R) - H(L)$ for $L < R$. This result follows by recalling that, since the process is order-R Markovian, the states of the agent's model are in a one-to-one correspondence with the directly observable states of the Markov process. At $L = 0$, no measurements have been made, and the agent's uncertainty over its states is exactly $H(R)$, the entropy of an R-block. At $L > 0$, the observer has gained $H(L)$ bits of information about the current state, and so the agent's state-uncertainty is reduced from $H(R)$ to $H(R) - H(L)$. Once one establishes that $\mathcal{H}(L) = H(R) - H(L)$, Eq. (14) follows from relatively simple manipulations of finite sums.

Before moving on, there are several additional comments we should make to clarify the transient information and its relevance to intelligent agents.

First, the transient information $\mathbf{T}$—together with the entropy rate h_μ and the order R of the Markov process—measures how difficult it is to synchronize to an environment. If a system has a large $\mathbf{T}$, then, on average, an agent will be highly uncertain about the internal state of the environment *while* synchronizing to it. Thus, $\mathbf{T}$ measures an important structural feature of the environment: how difficult it is for an agent to synchronize to it.

The second comment concerns a distinction between the transient information $\mathbf{T}$ and the excess entropy $\mathbf{E}$. Note that Eq. (14) implies that the $\mathbf{E}$ *does not* play a direct role in synchronization. (Although it is the case the $\mathbf{E}$ forms a lower bound for $\mathbf{T}$ [21].) Moreover, the excess entropies for all periodic processes of a given period are the same; a period-p process has $\mathbf{E} = \log_2 p$. However, the transient information is *not* the same for all period-p sequences. In fact, the transient information of period-p sequences varies considerably [21,30]. Since $h_\mu = 0$ for a periodic sequence, we have $\mathbf{T} = \mathbf{S}$, giving, in this case, a direct interpretation of $\mathbf{T}$ as the difficulty of synchronization. Thus, there is a range of different synchronization behaviors within the set of periodic sequences of a given period; these structural distinctions are not captured by the excess entropy $\mathbf{E}$.

Finally, it turns out that the transient information $\mathbf{T}$ is not directly proportional to the average number of measurements needed to synchronize. For example, there are three sets of period-5 sequences with distinct $H(L)$ behaviors. Direct calculation shows that, of these three sets, the set including the sequence $(10101)^\infty$ requires the most observations, on average, to synchronize. It does not have the largest $\mathbf{T}$ among the period-5 sequences, however. That honor is reserved for $(10000)^\infty$ [21]. Thus, the transient information does not directly measure the *time* it takes to synchronize. Instead, it measures the total *uncertainty* experienced by the agent during the process of synchronizing.

V. UNTANGLING SOURCES OF RANDOMNESS AND STRUCTURE

Using h_μ, $\mathbf{E}$, and $\mathbf{T}$, one can distinguish between environments structured in qualitatively and quantitatively different ways. See, e.g., Refs. [21,26,31] and references therein. But how might an agent, gathering statistics over larger and larger blocks of observations, estimate h_μ and $\mathbf{E}$? Not surprisingly, errors in the estimation of these quantities are linked; this can be seen graphically in Fig. 3. A misestimate of h_μ affects estimates of $\mathbf{E}$ and $\mathbf{T}$ and vice versa. We consider three special cases of the inter-relationship of these quantities and draw out the consequences for how an agent comes to know its environment.

A. Disorder as the Price of Ignorance

First, let's recall the consequences for an agent attempting to estimate the environment's randomness h_μ via the approximation $h_\mu(L)$. Assuming the agent has a finite memory, stopping the estimate at finite L is a necessity. This results in an entropy rate $h_\mu(L)$ which is almost always larger (and never smaller) than the actual rate h_μ. That is, the environment appears *more random* if the agent ignores correlations between observations separated by more than L steps.

For example, suppose the environment is periodic: the environment consists of the repeating length-16 sequence 1010111011101110. The environment entropy rate is thus 0. However, if the agent only keeps track of statistics over length-4 sequences, then the agent will estimate $h_\mu(4) \approx 0.303$ bits per symbol. The environment appears random to the agent, even though $h_\mu = 0$.

We are assuming here that, while the agent is only able to account for statistics over sequences of length L, it is nevertheless able to estimate these probabilities to arbitrary accuracy. When this is not the case, there are schemes that an agent can employ to improve on its estimation of block probabilities; see, for example, Ref. [32].

B. Instantaneous Synchronization and Predictability

Second, let's consider a scenario in which the agent happens to know the exact amount of apparent environmental memory $\mathbf{E}$. When this is the case, what happens to the agent's estimates of how random the environment is? In particular, what happens if the agent assumes it is synchronized to the environment at some finite L?

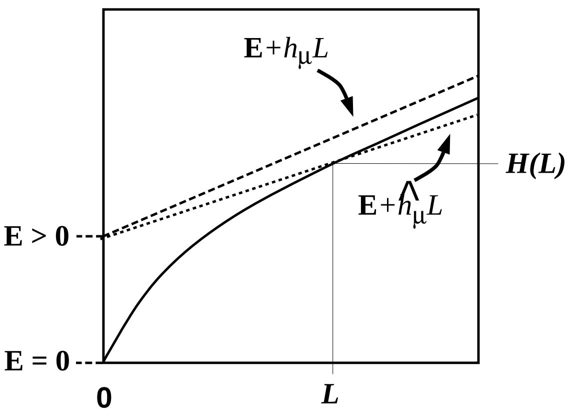

FIG. 4. Assumed synchronization converted to false predictability: Schematic illustration of how an agent, assuming it is synchronized, makes an underestimate $\widehat{h_\mu}$ (slope of dotted line) for an environment with excess entropy $\mathbf{E} > 0$ and entropy rate h_μ (slope of dashed line).

In so doing, the agent is assuming $H(L) = \mathbf{E} + h_\mu L$ at that L. The geometric construction for this scenario is given in Fig. 4. In effect the environment is erroneously considered to be a completely observable Markovian process in which $H(L)$ has converged to its asymptotic form exactly at some finite L [21,31]. If the agent then uses its value for $\mathbf{E}$, one arrives at the estimator $\widehat{h_\mu}$ where

$$\widehat{h_\mu} \equiv \frac{H(L) - \mathbf{E}}{L} \neq h_\mu \ . \qquad (15)$$

The line $\mathbf{E} + \widehat{h_\mu} L$ appears fixed at $\mathbf{E}$ when that intercept should be lower at the given L. The result, easily gleaned from Fig. 4, is that the entropy rate h_μ is underestimated as $\widehat{h_\mu}$. In other words, the agent will believe the environment to be more predictable than it actually is.

C. Assumed Synchronization Implies Reduced Apparent Memory

We now consider a different, less straightforward situation. Suppose that, because of some prior knowledge, the agent knows the exact environment entropy rate h_μ. For example, the agent could know before making any observations that the environment was periodic and, hence, had an h_μ of zero.

Given this situation, what happens if the agent assumes it is synchronized, when it is not? Figure 5 illustrates this situation. In this case, the agent infers an excess entropy $\widehat{\mathbf{E}}$ that is less than the true environment excess entropy $\mathbf{E}$.

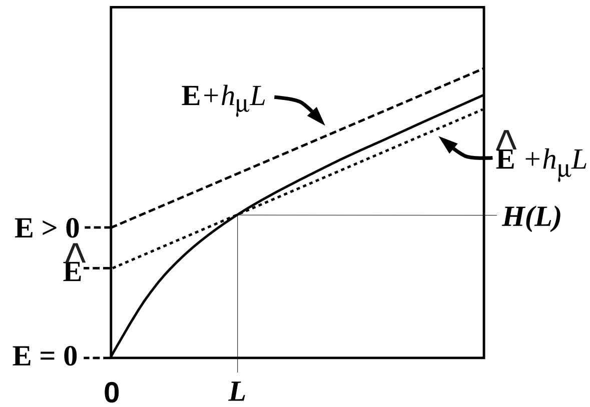

FIG. 5. Assumed synchronization leads to less apparent memory: Schematic illustration of how assuming synchronization to an environment, in this case implicitly assuming $H(L) = \mathbf{E} + h_\mu L$, leads to an underestimate $\widehat{\mathbf{E}}$ of the actual memory $\mathbf{E} > 0$.

If, at a given L, the agent approximates the entropy-rate estimate $h_\mu(L) = H(L) - H(L-1)$ by the true entropy h_μ, then the offset between the asymptote and $H(L)$ is simply $\mathbf{E} + h_\mu L - H(L)$. From Fig. 5 we see that we have a reduced apparent memory $\widehat{\mathbf{E}} \leq \mathbf{E}$ of

$$\widehat{\mathbf{E}} = H(L) - h_\mu L \ . \qquad (16)$$

If, instead of using the exact environment entropy rate h_μ, as we just did, the agent uses the estimate $h_\mu(L)$, the agent will infer an excess entropy that is even smaller than $\widehat{\mathbf{E}}$. To see this, note that $h_\mu(L) \geq h_\mu$ and replace h_μ by $h_\mu(L)$ in Eq. (16). Thus, assuming synchronization, in the sense that $h_\mu(L) = h_\mu$, leads one to underestimate the apparent memory $\mathbf{E}$. And so, the environment appears less structurally complex than it is.

In this section we considered just three among a number of possible tradeoffs encountered during an agent's inferring its environment's structure and memory. We discuss these tradeoffs more thoroughly, although in a different context, in Ref. [21]. The main lessons that emerge from this analysis are that unseen memory is converted to unpredictability, assumed synchronization leads to underestimating environmental structure, and ignoring environmental structure leads to underestimating its randomness.

VI. CONCLUSION

We have looked at a range of issues concerning how an agent synchronizes to its environment. In so doing, we reviewed several information theoretic properties of the environment: the entropy rate h_μ, the excess entropy $\mathbf{E}$, and the transient information $\mathbf{T}$. The main result reported here is contained in Eq. (14), which states that the total uncertainty experienced while an agent synchronizes to a Markovian environment is directly related to $\mathbf{T}$. Thus, the $\mathbf{T}$ captures that feature of an environment which makes it difficult to synchronize to.

In an effort to understand the roles of the various structural quantities and to show the consequences of ignoring them in estimating environmental randomness, structure, and synchronization, we considered various trade-offs between finite-L estimates of the excess entropy $\mathbf{E}$ and the entropy rate h_μ. In particular, we argued that if an agent does not take one or another into account (by, say, assuming it is synchronized, when it is not), the agent will systematically *over-* or *underestimate* an environment's entropy rate h_μ. As a result, even if an agent focuses only on quantifying the randomness of its environment, it must nevertheless have some method that accounts for the environment's structural features.

ACKNOWLEDGMENTS

The authors thank Kristian Lindgren and Cosma Shalizi for comments on the manuscript. This work was supported at the Santa Fe Institute under the Computation, Dynamics, and Inference Program via SFI's core grants from the National Science and MacArthur Foundations. Direct support was provided from DARPA contract F30602-00-2-0583. DPF thanks the Department of Physics and Astronomy at the University of Maine, Orono, for their hospitality during the summers of 2000 and 2001, when part of this work was completed.

[1] H. A. Simon. *Sciences of the Artifcial*. MIT Press, Cambridge, Massachusetts, 1969.

[2] H. Peyton Young. *Individual strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton, New Jersey, 1998.

[3] R. E. Nisbett and L. Ross. *Human Inference: Strategies and Shortcomings of Social Judgment*. Prentice-Hall, Englewood Cliffs, New Jersey, 1980.

[4] Z. Kunda. *Social Cognition: Making Sense of People*. MIT Press, Cambridge, Massachusetts, 1999.

[5] L. Vygotsky. *Thought and Language*. MIT Press, Cambridge, Massachusetts, revised edition, 1986. First published as *Myshlenie irech'*, Moscow, 1934.

[6] E. Hutchins. *Cognition in the Wild*. MIT Press, Cambridge, Massachusetts, 1996.

[7] J. M. Vidal. *Computational Agents That Learn About Agents: Algorithms for Their Design and a Predictive Theory of Their Behavior*. PhD thesis, University of Michigan, 1998.

[8] J. M. Vidal and E. H. Durfee. Predicting the expected behavior of agents that learn about agents: the CLRI framework. E-print, arxiv.org, cs.MA/0001008, 2000.

[9] E. F. Moore. Gedanken experiments on sequential automata. In C. E. Shannon and J. McCarthy, editors, *Automata Studies*, pages 129–153. Princeton Univeristy Press, Princeton, New Jersey, 1956.

[10] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[11] J. P. Crutchfield. Semantics and thermodynamics. In M. Casdagli and S. Eubank, editors, *Nonlinear Modeling and Forecasting*, volume XII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 317–359, Reading, Massachusetts, 1992. Addison-Wesley.

[12] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423, 1948. as reprinted in "The Mathematical Theory of Communication", C. E. Shannon and W. Weaver, University of Illinois Press, Champaign-Urbana (1963).

[13] D. Blackwell and L. Koopmans. On the identifiability problem for functions of Markov chains. *Ann. Math. Statist.*, 28:1011–1015, 1957.

[14] R. J. Elliot. *Hidden Markov Models: Estimation and Control*. Springer-Verlag, 1995.

[15] R. C. Titsworth. Optimal ranging codes. *IEEE Trans. Space Elec. & Telem.*,, SET-10:19–30, 1964.

[16] L. I. Bluestein. Interleaving of pseudorandom sequences for synchronization. *IEEE Trans. Aerospace & Elec. Sys.*, AES-4:551–556, 1968.

[17] J. J. Stiller. Rapid acquisition sequences. *IEEE Trans. Info Th.*, IT-14:221–225, 1968.

[18] J. L. Massey. Noisy sequence acquisition with minimum computation. Lecture Notes, Ulm Winter School, 2000.

[19] B. LeBaron. Evolution and time horizons in an agent-based stock market. *Macroeconomic Dynamics*, 5:225–254, 2001.

[20] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.

[21] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. Technical Report 01-02-012, arXiv.org/abs/cond-mat/0102181, Santa Fe Institute, 2001.

[22] J. P. Crutchfield and N. H. Packard. Symbolic dynamics of noisy chaos. *Physica D*, 7:201–223, 1983.

[23] R. Shaw. *The Dripping Faucet as a Model Chaotic System*. Aerial Press, Santa Cruz, California, 1984.

[24] P. Grassberger. Toward a quantitative theory of self-generated complexity. *Intl. J. Theo. Phys.*, 25(9):907–938, 1986.

[25] K. Lindgren and M. G. Norhdal. Complexity measures and cellular automata. *Complex Systems*, 2(4):409–440, 1988.

[26] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. physics/0007070v2, 2000.

[27] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 2001. in press.

[28] W. Li. On the relationship between complexity and entropy for Markov chains and regular languages. *Complex Systems*, 5(4):381–399, 1991.

[29] O. Penrose. *Foundations of statistical mechanics; a deductive treatment*. Pergamon Press, Oxford, 1970.

[30] D. P. Feldman and J. P. Crutchfield. Synchronizing to a periodic signal: The transient information of periodic sequences. manuscript in preparation.

[31] W. Ebeling. Prediction and entropy of nonlinear dynamical systems and symbolic sequences with LRO. *Physica D*, 109:42–52, 1997.

[32] T. Schürmann and P. Grassberger. Entropy estimation of symbol sequences. *Chaos*, 6:414–427, 1996.