

Information accessibility and cryptic processes

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2009 J. Phys. A: Math. Theor. 42 362002

(<http://iopscience.iop.org/1751-8121/42/36/362002>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.155

The article was downloaded on 03/06/2010 at 08:06

Please note that [terms and conditions apply](#).

FAST TRACK COMMUNICATION

Information accessibility and cryptic processesJohn R Mahoney¹, Christopher J Ellison¹ and James P Crutchfield^{1,2}¹ Complexity Sciences Center and Physics Department, University of California at Davis,
One Shields Avenue, Davis, CA 95616, USA² Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USAE-mail: jrmahoney@ucdavis.edu, cellison@cse.ucdavis.edu and chaos@cse.ucdavis.edu

Received 8 June 2009, in final form 31 July 2009

Published 19 August 2009

Online at stacks.iop.org/JPhysA/42/362002**Abstract**

We give a systematic expansion of the crypticity—a recently introduced measure of the inaccessibility of a stationary process's internal state information. This leads to a hierarchy of k -cryptic processes and allows us to identify finite-state processes that have infinite cryptic order—the internal state information is present across arbitrarily long, observed sequences. The crypticity expansion is exact in both the finite- and infinite-order cases. It turns out that k -crypticity is complementary to the Markovian finite-order property that describes state information in processes. One application of these results is an efficient expansion of the excess entropy—the mutual information between a process's infinite past and infinite future—that is finite and exact for finite-order cryptic processes.

PACS numbers: 02.50.-r, 89.70.+c, 05.45.Tp, 02.50.Ey

1. Introduction

The data of phenomena come to us through observation. A large fraction of the theoretical activity of model building, though, focuses on internal mechanism. How are observation and modeling related? A first step is to frame the problem in terms of hidden processes—internal mechanisms probed via instruments that, in particular, need not accurately report a process's internal state. A practical second step is to measure the difference between internal structure and the information in observations.

We recently established that the amount of observed information a process communicates from the past to the future—the *excess entropy*—is the mutual information between its forward- and reverse-time minimal causal representations [1, 2]. This closed-form expression gives a concrete connection between the observed information and a process's internal structure.

Excess entropy and related mutual information quantities are widely used diagnostics for complex systems. They have been applied to detect the presence of organization in

dynamical systems [3–6], in spin systems [7–9], in neurobiological systems [10, 11] and even in language [12, 13], to mention only a very few uses. Thus, understanding how much internal state structure is reflected in the excess entropy is critical to whether or not these and other studies of complex systems can draw structural inferences about the internal mechanisms that produce observed behavior.

Unfortunately, there is a fundamental problem. The excess entropy is *not* the internal state information the process stores—rather, the latter is the process’s *statistical complexity* [1, 2]. On the positive side, there is a diagnostic. The difference between, if you will, experiment and theory (between observed information and internal structure) is controlled by the difference between a process’s excess entropy and its statistical complexity. This difference is called the *crypticity*—how much internal state information is inaccessible [1, 2]. Here we introduce a classification of processes using a systematic expansion of crypticity.

The starting point is *computational mechanics*’s minimal causal representation of a stochastic process \mathcal{P} —the ϵ -machine [14, 15]. There, a process is viewed as a channel that communicates information from the past, $\overleftarrow{X} = \dots X_{-3}X_{-2}X_{-1}$, to the future, $\overrightarrow{X} = X_0X_1X_2\dots$ (X_t takes values in a finite measurement alphabet \mathcal{A} .) The excess entropy is the shared (or mutual) information between the past and the future: $\mathbf{E} = I[\overleftarrow{X}; \overrightarrow{X}]$. The amount of historical information that a process stores in the present is different. It is given by the Shannon information $C_\mu = H[S]$ of the distribution over the ϵ -machine’s *causal states* \mathcal{S} . C_μ is called the *statistical complexity* and the causal states are sets of pasts \overleftarrow{x} that are equivalent for prediction [14]:

$$\epsilon(\overleftarrow{x}) = \{\overleftarrow{x'} : \Pr(\overrightarrow{X}|\overleftarrow{x}) = \Pr(\overrightarrow{X}|\overleftarrow{x'})\}. \quad (1)$$

Causal states have a Markovian property that they render the past and future statistically independent; they *shield* the future from the past [15]:

$$\Pr(\overleftarrow{X}, \overrightarrow{X}|\mathcal{S}) = \Pr(\overleftarrow{X}|\mathcal{S})\Pr(\overrightarrow{X}|\mathcal{S}). \quad (2)$$

ϵ -machines are also *unifilar* [14, 16]: from the start state, each observed sequence $\dots x_{-3}x_{-2}x_{-1}\dots$ corresponds to one and only one sequence of causal states. The signature of unifilarity is that on knowing the current state and measurement, the uncertainty in the next state vanishes: $H[\mathcal{S}_{t+1}|\mathcal{S}_t, X_t] = 0$.

Although they are not the same, the basic relationship between these quantities is clear: \mathbf{E} is the process’s channel utilization and C_μ is the sophistication of that channel. Their difference, one of our main concerns in the following, indicates how a process stores, manipulates and hides internal state information.

Until recently, \mathbf{E} could not be as directly calculated from the ϵ -machine as the process’s entropy rate h_μ and its statistical complexity. [1, 2] solved this problem, giving a closed-form expression for the excess entropy:

$$\mathbf{E} = I[\mathcal{S}^+; \mathcal{S}^-], \quad (3)$$

where \mathcal{S}^+ are the causal states of the process scanned in the ‘forward’ direction and \mathcal{S}^- are the causal states of the process scanned in the ‘reverse’ time direction.

This result comes in a historical context. Some time ago, an explicit expression for the excess entropy had been developed from the Hamiltonian for one-dimensional spin chains with range- R interactions [8]:

$$\mathbf{E} = C_\mu - R h_\mu. \quad (4)$$

One-dimensional spin chains are special cases of order- R Markov processes. For this more general class of processes, a similar, but slightly less compact form is known:

$$\mathbf{E} = H[X_0^R] - R h_\mu, \quad (5)$$

where $X_0^R = X_0, \dots, X_{R-1}$. It has also been known for some time that the statistical complexity is an upper bound on the excess entropy [16]:

$$\mathbf{E} \leq C_\mu,$$

which follows from the equality derived there:

$$\mathbf{E} = C_\mu - H[S^+|\vec{X}].$$

Using forward and reverse ϵ -machines, [1] extended this, deriving the closed-form expression for \mathbf{E} in equation (3) and two new bounds on \mathbf{E} : $\mathbf{E} \leq C_\mu^-$ and $\mathbf{E} \leq C_\mu^+$. It also showed that

$$H[S^+|\vec{X}] = H[S^+|S^-] \quad (6)$$

and identified this quantity as controlling how a process hides its internal state information. For this reason, it is called the process's *crypticity*:

$$\chi^+ = H[S^+|\vec{X}]. \quad (7)$$

In the context of forward and reverse ϵ -machines, one must distinguish two crypticities; depending on the scan direction one has

$$\chi^+ = H[S^+|S^-] \quad \text{or} \quad \chi^- = H[S^-|S^+].$$

In the following we will not concern ourselves with reverse representations and so can simplify the notation, using C_μ for C_μ^+ and χ for χ^+ .

Here we show that, for a restricted class of processes, the crypticity in equation (6) can be systematically expanded to give an alternative closed form to the excess entropy in equation (3). One ancillary benefit is a new and, we argue, natural hierarchy of processes in terms of information accessibility.

2. k -crypticity

The process classifications based on spin-block length and order- R Markov are useful. They give some insight into the nature of the kinds of process we can encounter and, concretely, they allow for closed-form expressions for the excess entropy (and other system properties). In a similar vein, we wish to carve the space of processes with a new blade. We define the class of k -cryptic processes and develop their properties and closed-form expressions for their excess entropies.

For convenience, we need to introduce several shorthands. First, to denote a symbol sequence that begins at time t and is L symbols long, we write X_t^L . Note that X_t^L includes X_{t+L-1} , but not X_{t+L} . Second, to denote a symbol sequence that begins at time t and continues on to infinity, we write \vec{X}_t . Analogously, the causal state at time t is denoted S_t , and a sequence of states beginning at time t that is L states long is denoted S_t^L .

Definition. *The k -crypticity criterion is satisfied when*

$$H[S_k|\vec{X}_0] = 0. \quad (8)$$

Definition. *A k -cryptic process is one for which the process's ϵ -machine satisfies the k -crypticity criterion.*

Definition. An ∞ -cryptic process is one for which the process's ϵ -machine does not satisfy the k -crypticity criterion for any finite k .

Lemma 1. $H[S_k|\vec{X}_0]$ is a nonincreasing function of k .

Proof. This follows directly from stationarity and the fact that conditioning on more random variables cannot increase entropy:

$$H[S_{k+1}|\vec{X}_0] = H[S_k|\vec{X}_{-1}] \leq H[S_k|\vec{X}_0]. \quad \square$$

Lemma 2. If \mathcal{P} is k -cryptic, then \mathcal{P} is also j -cryptic for all $j > k$.

Proof. Being k -cryptic implies $H[S_k|\vec{X}_0] = 0$. Applying lemma 1, $H[S_j|\vec{X}_0] \leq H[S_k|\vec{X}_0] = 0$. By positivity of entropy, we conclude that \mathcal{P} is also j -cryptic. \square

This provides us with a new way of partitioning the space of processes. We create a parametrized class of sets $\{\chi_k : k = 0, 1, 2, \dots\}$, where $\chi_k = \{\mathcal{P} : \mathcal{P} \text{ is } k\text{-cryptic and not } (k-1)\text{-cryptic}\}$.

The following result provides a connection to a very familiar class of processes.

Proposition 1. If a process \mathcal{P} is order- k Markov, then it is k -cryptic.

Proof. If \mathcal{P} is order- k Markov, then $H[S_k|X_0^k] = 0$. Conditioning on more variables does not increase uncertainty, so

$$H[S_k|X_0^k, \vec{X}_k] = 0.$$

But the left-hand side is $H[S_k|\vec{X}_0]$. Therefore, \mathcal{P} is k -cryptic. \square

Note that the converse of proposition 1 is not true. For example, the Even Process (EP), the Random Noisy Copy Process (RnC) and the Random Insertion Process (RIP) (see [1, 2]), are all 1-cryptic, but are not order- R Markov for any finite R .

Note also that proposition 1 does not preclude an order- k Markov process from being j -cryptic, where $j < k$. Later we will show an example demonstrating this.

Given a process, in general one will not know its cryptic order. One way to investigate this is to study the sequence of estimates of χ at different orders. To this end, we define the k -cryptic approximation.

Definition. The k -cryptic approximation is defined as

$$\chi(k) = H[S_0|X_0^k, S_k].$$

2.1. The k -cryptic expansion

We will now develop a systematic expansion of χ to order k in which $\chi(k)$ appears directly and the k -crypticity criterion plays the role of an error term.

Theorem 1. The process crypticity is given by

$$\chi = \chi(k) + H[S_k|\vec{X}_0]. \quad (9)$$

Proof. We calculate directly, starting from the definition, adding and subtracting the k -crypticity criterion term from χ 's definition, equation (7):

$$\chi = H[S_0|\vec{X}_0] - H[S_k|\vec{X}_0] + H[S_k|\vec{X}_0].$$

We claim that the first two terms are $\chi(k)$. Expanding the conditionals in the purported $\chi(k)$ terms and then canceling, we get joint distributions:

$$H[\mathcal{S}_0|\vec{X}_0] - H[\mathcal{S}_k|\vec{X}_0] = H[\mathcal{S}_0, \vec{X}_0] - H[\mathcal{S}_k, \vec{X}_0].$$

Now, splitting the future into two pieces and using this to write conditionals, the right-hand side becomes

$$H[\vec{X}_k|\mathcal{S}_0, X_0^k] + H[\mathcal{S}_0, X_0^k] - H[\vec{X}_k|\mathcal{S}_k, X_0^k] - H[\mathcal{S}_k, X_0^k].$$

Appealing to the ϵ -machine's unifilarity, we then have

$$H[\vec{X}_k|\mathcal{S}_k] + H[\mathcal{S}_0, X_0^k] - H[\vec{X}_k|\mathcal{S}_k, X_0^k] - H[\mathcal{S}_k, X_0^k].$$

Now, applying causal shielding gives

$$H[\vec{X}_k|\mathcal{S}_k] + H[\mathcal{S}_0, X_0^k] - H[\vec{X}_k|\mathcal{S}_k] - H[\mathcal{S}_k, X_0^k].$$

Canceling terms, this simplifies to

$$H[\mathcal{S}_0, X_0^k] - H[\mathcal{S}_k, X_0^k].$$

We now re-expand, using unifilarity to give

$$H[\mathcal{S}_0, X_0^k, \mathcal{S}_k] - H[\mathcal{S}_k, X_0^k].$$

Finally, we combine these, using the definition of conditional entropy, to simplify again:

$$H[\mathcal{S}_0|\mathcal{S}_k, X_0^k].$$

Note that this is our definition of $\chi(k)$.

This establishes our original claim:

$$\chi = \chi(k) + H[\mathcal{S}_k|\vec{X}_0],$$

with the k -crypticity criterion playing the role of an approximation error. □

Corollary 1. *A process \mathcal{P} is k -cryptic if and only if*

$$\chi = \chi(k).$$

Proof. Given the order- k expansion of χ just developed, we now assume that the k -crypticity criterion is satisfied; namely, $H[\mathcal{S}_k|\vec{X}_0] = 0$. Thus, we have from equation (9):

$$\chi = \chi(k).$$

Likewise, assuming $\chi = \chi(k)$ requires, by equation (9) that $H[\mathcal{S}_k|\vec{X}_0] = 0$ and thus the process is k -cryptic. □

Corollary 2. *For any process, $\chi(0) = 0$.*

Proof.

$$\begin{aligned} \chi(0) &= H[\mathcal{S}_0|X_0^0, \mathcal{S}_0] \\ &= H[\mathcal{S}_0|\mathcal{S}_0] \\ &= 0. \end{aligned}$$

□

2.2. Convergence

Proposition 2. *The approximation $\chi(k)$ is a nondecreasing function of k .*

Proof. Lemma 1 showed that $H[\mathcal{S}_k|\vec{X}_0]$ is a nonincreasing function of k . By theorem 1, $\chi(k)$ must be a nondecreasing function of k . \square

Corollary 3. *Once $\chi(k)$ reaches the value χ , $\chi(j) = \chi$ for all $j > k$.*

Proof. If there exists such a k , then by theorem 1 the process is k -cryptic. By lemma 2, the process is j -cryptic for all $j > k$. Again, by theorem 1, $\chi(j) = \chi$. \square

Corollary 4. *If there is a $k \geq 1$ for which $\chi(k) = 0$, then $\chi(1) = 0$.*

Proof. By positivity of the conditional entropy $H[\mathcal{S}_0|X_0, \mathcal{S}_1]$, $\chi(1) \geq 0$. By the nondecreasing property of $\chi(k)$ from proposition 2, $\chi(1) \leq \chi(k) = 0$. Therefore, $\chi(1) = 0$. \square

Corollary 5. *If $\chi(1) = 0$, then $\chi(k) = 0$ for all k .*

Proof. Applying stationarity, $\chi(1) = H[\mathcal{S}_0|X_0, \mathcal{S}_1] = H[\mathcal{S}_k|X_k, \mathcal{S}_{k+1}]$. We are given $\chi(1) = 0$ and so $H[\mathcal{S}_k|X_k, \mathcal{S}_{k+1}] = 0$. We use this below. Expanding $\chi(k+1)$,

$$\begin{aligned} \chi(k+1) &= H[\mathcal{S}_0|X_0^{k+1}, \mathcal{S}_{k+1}] \\ &= H[\mathcal{S}_0|X_0^k, X_k, \mathcal{S}_{k+1}] \\ &= H[\mathcal{S}_0|X_0^k, \mathcal{S}_k, X_k, \mathcal{S}_{k+1}] \\ &\leq H[\mathcal{S}_0|X_0^k, \mathcal{S}_k] \\ &= \chi(k). \end{aligned}$$

The third line follows from $\chi(1) = 0$. By proposition 2, $\chi(k+1) \geq \chi(k)$. Therefore, $\chi(k+1) = \chi(k)$. Finally, using $\chi(1) = 0$, we have by induction that $\chi(k) = 0$ for all k . \square

Corollary 6. *If there is a $k \geq 1$ for which $\chi(k) = 0$, then $\chi(j) = 0$ for all $j \geq 1$.*

Proof. This follows by composing corollary 4 with corollary 5. \square

Together, the proposition and its corollaries show that $\chi(k)$ is a nondecreasing function of k which, if it reaches χ at a finite k , remains at that value for all larger k .

Proposition 3. *The cryptic approximation $\chi(k)$ converges to χ as $k \rightarrow \infty$.*

Proof. Note that $\chi = \lim_{k \rightarrow \infty} H[\mathcal{S}_0|X_0^k]$ and recall that $\chi(k) = H[\mathcal{S}_0|X_0^k, \mathcal{S}_k]$. We show that the difference approaches zero:

$$\begin{aligned} H[\mathcal{S}_0|X_0^k] - H[\mathcal{S}_0|X_0^k, \mathcal{S}_k] &= H[\mathcal{S}_0, X_0^k] - H[X_0^k] - H[\mathcal{S}_0, X_0^k, \mathcal{S}_k] + H[X_0^k, \mathcal{S}_k] \\ &= H[\mathcal{S}_0, X_0^k] - H[X_0^k] - H[\mathcal{S}_0, X_0^k] + H[X_0^k, \mathcal{S}_k] \\ &= H[X_0^k, \mathcal{S}_k] - H[X_0^k] \\ &= H[\mathcal{S}_k|X_0^k]. \end{aligned}$$

Moreover, $\lim_{k \rightarrow \infty} H[\mathcal{S}_k|X_0^k] = 0$ by the ϵ map from pasts to causal states of equation (1). Therefore, as $k \rightarrow \infty$, $\chi(k) \rightarrow \chi$. \square

2.3. Excess entropy for k -cryptic processes

Given a k -cryptic process, we can calculate its excess entropy in a form that involves a sum of $\propto |\mathcal{A}^k|$ terms, where each term involves products of k matrices. Specifically, we have the following.

Corollary 7. *A process \mathcal{P} is k -cryptic if and only if $\mathbf{E} = C_\mu - \chi(k)$.*

Proof. From [1], we have $\mathbf{E} = C_\mu - \chi$, and by corollary 1, $\chi = \chi(k)$. Together, these complete the proof. \square

The following proposition is a simple and useful consequence of the class of k -cryptic processes.

Corollary 8. *A process \mathcal{P} is 0-cryptic if and only if $\mathbf{E} = C_\mu$.*

Proof. If \mathcal{P} is 0-cryptic, then $\mathbf{E} = C_\mu - \chi(0)$ and corollary 2 says that $\chi(0) = 0$. To establish the opposite direction, note that $\mathbf{E} = C_\mu$ implies $\chi = 0$. Applying corollary 2 shows $\chi = \chi(0)$, and so the process is 0-cryptic by corollary 1. \square

2.4. Crypticity of spin chains

Now, we provide results on the crypticity of one-dimensional spin chains to complement prior results on Markovity and excess entropy. First recall equation (5), which gives the excess entropy for order- R Markov processes:

$$\mathbf{E} = H[X_0^R] - Rh_\mu.$$

By proposition 1, such processes are also R -cryptic and so

$$\mathbf{E} = C_\mu - \chi(R).$$

One-dimensional spin chains are precisely those order- R Markov processes for which the statistical complexity, $C_\mu \equiv H[S_R]$, equals the entropy over R -blocks, $H[X_0^R]$. Reference [8] stated a condition under which equality held in terms of transfer matrices. Here, we state a simpler condition by equating two chain-rule expansions of $H[X_0^R, S_R]$:

$$H[X_0^R | S_R] + H[S_R] = H[S_R | X_0^R] + H[X_0^R].$$

Since the process is Markov, $H[S_R | X_0^R] = 0$ and thus

$$H[X_0^R] = H[S_R] \iff H[X_0^R | S_R] = 0.$$

In words, spin chains are processes for which there exists a one-to-one correspondence between the R -blocks and the causal states, confirming the interpretation specified in [8].

The above equations also show that spin chains have $\chi(R) = Rh_\mu$. Here we provide another proof:

Proposition 4.

$$H[X_0^R | S_R] = 0 \iff \chi(R) = Rh_\mu, \tag{10}$$

where h_μ is the process's entropy rate.

Proof. The proof is a direct calculation:

$$\begin{aligned}
 \chi(R) &= H[\mathcal{S}_0 | X_0^R, \mathcal{S}_R] \\
 &= H[\mathcal{S}_0, X_0^R] - H[X_0^R, \mathcal{S}_R] \\
 &= H[\mathcal{S}_0, X_0^R] - H[X_0^R | \mathcal{S}_R] - H[\mathcal{S}_R] \\
 &= H[\mathcal{S}_0, X_0^R] - H[X_0^R | \mathcal{S}_R] - H[\mathcal{S}_0] \\
 &= H[X_0^R | \mathcal{S}_0] - H[X_0^R | \mathcal{S}_R] \\
 &= Rh_\mu - H[X_0^R | \mathcal{S}_R]. \quad \square
 \end{aligned}$$

Proposition 5. *Periodic processes are 0-cryptic.*

Proof. Periodic processes are order- R Markov spin chains, so $\mathbf{E} = C_\mu - Rh_\mu$. Since $h_\mu = 0$, $\mathbf{E} = C_\mu$. By corollary 8 the process is 0-cryptic. \square

Proposition 6. *An order- R spin chain with positive entropy rate is not $(R - 1)$ -cryptic.*

Proof. Assume that the order- R Markov spin chain is $(R - 1)$ -cryptic.

For $R \geq 1$, if the process is $(R - 1)$ -cryptic, then by corollary 1 $\chi(R - 1) = \chi$. Combining this with the above proposition 4, we have $\chi(R - 1) = (R - 1)h_\mu - H[X_0^{R-1} | \mathcal{S}_{R-1}]$. If it is an order- R Markov spin chain, then we also have from equation (4) that $\chi = Rh_\mu$. Combining this with the previous equation, we find that $H[X_0^{R-1} | \mathcal{S}_{R-1}] = -h_\mu$. By positivity of conditional entropies, we have reached a contradiction. Therefore an order- R Markov spin chain must not be $(R - 1)$ -cryptic.

For $R = 0$, the proof also holds since negative cryptic orders are not defined. \square

Proposition 7. *An order- R spin chain with positive entropy rate is not k -cryptic for any $0 \leq k < R$.*

Proof. By lemma 2, if the process were k -cryptic for some $0 \leq k < R$, then it would also be $(R - 1)$ -cryptic. By proposition 6, this is not true. Therefore, the primitive orders of Markovity and crypticity are the same. \square

3. Examples

It is helpful to see crypticity in action. We now turn to a number of examples to illustrate how various orders of crypticity manifest themselves in ϵ -machine structure and what kinds of processes are cryptic and so hide internal state information from an observer. For details (transition matrices, notation, and the like) not included in the following and for complementary discussions and analyses of them, see [1, 2, 17].

We start at the bottom of the crypticity hierarchy with a 0-cryptic process and then show examples of 1-cryptic and 2-cryptic processes. Continuing up the hierarchy, we generalize and give a parametrized family of processes that are k -cryptic. Finally, we demonstrate an example that is ∞ -cryptic.

It should be pointed out, though, that these examples were hand chosen to illustrate some of the range of possible processes in terms of cryptic and Markov orders. If one were to encounter a process in the wild, its cryptic order would not be known and the calculation of crypticity would require that one determines the cryptic order. One can estimate the cryptic order by calculating the cryptic approximation until it appears to have converged or computational power has run out. Alternatively, one might deduce the order exactly via some




Figure 1. A 0-cryptic process: Even Process. The transitions denote the probability p of generating symbol x as $p|x$.




Figure 2. A 1-cryptic process: Golden Mean Process.

other technique, as we do in the upcoming examples. Of course, we wish to note that [1] demonstrates how to calculate χ without any knowledge of the cryptic order.

3.1. Even Process: 0-cryptic

Figure 1 gives the ϵ -machine for the Even Process. The Even Process produces binary sequences in which all blocks of uninterrupted 1s are even in length, bounded by 0s. Further, after each even length is reached, there is a probability p of breaking the block of 1s by inserting one or more 0s.

Reference [2] showed that the Even Process is 0-cryptic with a statistical complexity of $C_\mu = H(1/(2-p))$, an entropy rate of $h_\mu = H(p)/(2-p)$ and a crypticity of $\chi = 0$. Note that $H(p)$ is the binary entropy function. If $p = \frac{1}{2}$, then $\mathbf{E} = C_\mu = \log_2(3) - \frac{2}{3}$ bits. (As [2] notes, these closed-form expressions for C_μ and \mathbf{E} have been known for some time.)

To see why the Even Process is 0-cryptic, first note that the semi-infinite string $\vec{X}_0 = 1, 1, 1 \dots$ occurs with probability 0. So with probability 1, a given future will have only a finite number of 1s before a 0 is seen. Once the 0 is seen, it is straightforward to count the number of 1s preceding it. If the number of 1s is even, then \mathcal{S}_0 , the causal state that preceded this future, is A. Otherwise, it is B. In either case, we know the causal state with certainty, and so, $H[\mathcal{S}_0|\vec{X}_0] = 0$.

It is important to note that this process is *not* order- R Markov for any finite R [17]. Nonetheless, our new expression for \mathbf{E} is valid. This shows the broadening of our ability to calculate \mathbf{E} even for low complexity processes that are, in effect, infinite-order Markov.

3.2. Golden Mean Process: 1-cryptic

Figure 2 shows the ϵ -machine for the Golden Mean Process [17]. The Golden Mean Process is one in which no two 0s occur consecutively. After each 1, there is a probability p of generating a 0. As the sequence length grows, the ratio of the number of allowed words of length L to the number of allowed words at length $L - 1$ approaches the golden ratio; hence, its name. The Golden Mean Process ϵ -machine looks remarkably similar to that for the Even Process. The informational analysis, however, shows that they have markedly different properties.

Reference [2] showed that the Golden Mean Process has the same statistical complexity and entropy rate as the Even Process: $C_\mu = H(1/(2-p))$ and $h_\mu = H(p)/(2-p)$.




Figure 3. A 2-cryptic process: Butterfly Process over a 6-symbol alphabet.

However, the crypticity is not zero (for $0 < p < 1$). From corollary 1 we calculate

$$\begin{aligned} \chi &= \chi(1) \\ &= H[\mathcal{S}_0 | X_0^1, \mathcal{S}_1] \\ &= H[\mathcal{S}_0 | X_0^1] \\ &= \Pr(0)H[\mathcal{S}_0 | X_0 = 0] + \Pr(1)H[\mathcal{S}_0 | X_0 = 1] \\ &= H(p)/(2 - p). \end{aligned}$$

If $p = \frac{1}{2}$, $C_\mu = \log_2(3) - \frac{2}{3}$ bits, excess entropy $\mathbf{E} = \log_2(3) - \frac{4}{3}$ bits, and crypticity $\chi = \frac{2}{3}$ bits. Thus, the excess entropy differs from that of the Even Process. (As with the Even Process, these closed-form expressions for C_μ and \mathbf{E} have been known for some time.)

The Golden Mean Process is 1-cryptic. To see why, it is enough to note that it is order-1 Markov. By proposition 1, it is 1-cryptic. We know it is not 0-cryptic since any future beginning with 1 could have originated in either state A or B. In addition, the spin-block expression for excess entropy of [17], equation (4) here, applies for an $R = 1$ Markov chain.

3.3. Butterfly Process: 2-cryptic

The next example, the Butterfly Process of figure 3, illustrates, in a more explicit way than possible with the previous processes, the role that crypticity plays and how it can be understood in terms of an ϵ -machine’s structure. Most of the explanation does not require calculating much, if anything.

It is first instructive to see why the Butterfly Process is *not* 1-cryptic.

If we can find a family $\{\vec{x}_0\}$ such that $H[\mathcal{S}_1 | \vec{X}_0 = \vec{x}_0] \neq 0$, then the total conditional entropy will be positive and, thus, the machine will not be 1-cryptic. To show that this can happen, consider the future $\vec{x}_0 = (0, 1, 2, 4, 4, 4, \dots)$. It is clear that the state following 1 must be A. Thus, in order to generate 0 or 1 before arriving at A, the state pair $(\mathcal{S}_0, \mathcal{S}_1)$ can be either (B, C) or (D, E) . This uncertainty in \mathcal{S}_1 is enough to break the criterion, and this occurs for the family of futures beginning with 01.

To see that the process is 2-cryptic, note that the two paths (B, C) and (D, E) converge on A. Therefore, there is no uncertainty in \mathcal{S}_2 given this future. It is reasonably straightforward to see that indeed *any* two-symbol word (X_0, X_1) will lead to a unique causal state. This

is because the Butterfly Process is a very limited version of an 8-symbol, order-2 Markov process.

Note that the transition matrix is doubly stochastic and so the stationary distribution is uniform. The statistical complexity is rather direct in this case: $C_\mu = \log_2 5$. We now can calculate χ using corollary 1:

$$\begin{aligned} \chi &= \chi(2) \\ &= H[\mathcal{S}_0 | X_0^2, \mathcal{S}_2] \\ &= H[\mathcal{S}_0 | X_0^2] \\ &= \Pr(01) \cdot H[\mathcal{S}_0 | X_0^2 = 01] \\ &\quad + \Pr(12) \cdot H[\mathcal{S}_0 | X_0^2 = 12] \\ &\quad + \Pr(13) \cdot H[\mathcal{S}_0 | X_0^2 = 13] \\ &= \frac{1}{10} \cdot 1 + \frac{1}{10} \cdot 1 + \frac{1}{10} \cdot 1 \\ &= \frac{3}{10} \text{ bits.} \end{aligned}$$

From corollary 7, we get an excess entropy of

$$\begin{aligned} \mathbf{E} &= C_\mu - \chi(2) \\ &= \log_2 5 - \frac{3}{10} \\ &\approx 2.0219 \text{ bits.} \end{aligned}$$

For comparison, if we had assumed the Butterfly Process was 1-cryptic, then we would have

$$\begin{aligned} \mathbf{E} &= C_\mu - \chi(1) \\ &= C_\mu - (H[\mathcal{S}_0, X_0] - H[\mathcal{S}_1, X_0]) \\ &\approx \log 2(5) - (3.3219 - 2.5062) \\ &= \log 2(5) - 0.8156 \approx 1.5063 \text{ bits.} \end{aligned}$$

We can see that this is substantially below the true value: a 25% error.

3.4. Restricted Golden Mean: k -cryptic

Now, we turn to illustrate a crypticity-parametrized family of processes, giving examples of k -cryptic processes for any k . We call this family the Restricted Golden Mean as its support is a restriction of the Golden Mean support. (See figure 4 for its ϵ -machines.) The $k = 1$ member of the family is exactly the Golden Mean.

It is straightforward to see that this process is order- k Markov since each word of length k induces just one causal state. Proposition 1 then implies it is (at most) k -cryptic. In order to show that it is not $(k - 1)$ -cryptic, consider the case $\vec{x}_0 = 1^k 0^\infty$. The first $(k - 1)$ 1s will induce a mixture over states k and 0. The following future $\vec{x}_k = 10^\infty$ is consistent with both states k and 0. Therefore, the $(k - 1)$ -crypticity criterion is not satisfied. Therefore, it is k -cryptic.

For arbitrary k , there are $k + 1$ causal states and the stationary distribution is

$$\pi = \left(\frac{2}{k+2}, \frac{1}{k+2}, \frac{1}{k+2}, \dots, \frac{1}{k+2} \right).$$

The statistical complexity is

$$C_\mu = \log_2(k+2) - \frac{2}{k+2}.$$




Figure 4. k -cryptic processes: Restricted Golden Mean Family.




Figure 5. k -cryptic processes: Stretched Golden Mean Family.

For the k th member of the family, we have for the crypticity:

$$\chi = \chi(k) = \frac{2k}{k+2}.$$

And the excess entropy follows directly from corollary 7:

$$\mathbf{E} = C_\mu - \chi = \log_2(k+2) - \frac{2(k+1)}{k+2},$$

which diverges with k . (Calculational details are found in [18].)

3.5. Stretched Golden Mean

The Stretched Golden Mean is a family of processes that does not occupy the same support as the Golden Mean. Instead of requiring that blocks of 0s should be of length 1, we require that they should be of length k . The ϵ -machine for this process is shown in figure 5.

Again, it is straightforward to see that this process is order- k Markov. To see that it is not 0-cryptic, note that

$$\begin{aligned} H[\mathcal{S}_0|\vec{X}_0] &= H[\mathcal{S}_0|X_0 = 0, \vec{X}_1] + H[\mathcal{S}_0|X_0 = 1, \vec{X}_1] \\ &\geq H[\mathcal{S}_0|X_0 = 1, \vec{X}_1] \\ &= \frac{2}{k+2} \sum_{\vec{x}_1} H[\mathcal{S}_0|X_0 = 1, \vec{X}_1 = \vec{x}_1] \\ &\geq \frac{2}{k+2} H[\mathcal{S}_0|\vec{X}_1 = 1^\infty] \\ &= \frac{2}{k+2} \\ &> 0. \end{aligned}$$

To see that this family is 1-cryptic, first note that if $X_0 = 1$, then $\mathcal{S}_1 = 0$. Next, consider the case when $X_0 = 0$. If the future $\vec{x}_1 = 1^\infty$, then $\mathcal{S}_1 = k$. Similarly, if the future $\vec{x}_1 = 0^n 1^\infty$, then $\mathcal{S}_1 = k - n$.

This family provides an example for which the cryptic order is strictly less than the Markov order. In this case, the cryptic order is fixed at 1 for all k , while the Markov order is k . Note that the separation between the Markov and cryptic order can grow arbitrarily large and, thus, the two properties are clearly not redundant.

The stationary distribution is the same as for the Restricted Golden Mean and so, then, is the statistical complexity. In addition, we have

$$\begin{aligned} \chi &= \chi(1) \\ &= H[\mathcal{S}_0|X_0, \mathcal{S}_1] \\ &= h_\mu. \end{aligned}$$

Consequently,

$$\mathbf{E} = C_\mu - \chi = C_\mu - h_\mu.$$

3.6. Nemo Process: ∞ -cryptic

We close our cryptic process bestiary with a (very) finite-state process that has infinite cryptic order: the three-state Nemo Process. Over no finite-length sequence will all of the internal state information be present in the observations. The Nemo Process ϵ -machine is shown in figure 6.

Its stationary state distribution is

$$\Pr(\mathcal{S}) \equiv \pi = \frac{1}{3-2p} \begin{pmatrix} A & B & C \\ 1 & 1-p & 1-p \end{pmatrix},$$

from which one calculates the statistical complexity:

$$C_\mu = \log_2(3-2p) - \frac{2(1-p)}{3-2p} \log_2(1-p).$$

The Nemo Process is not a finite-cryptic process. That is, there exists no finite k for which $H[\mathcal{S}_k|\vec{X}_0] = 0$. To show this, we must demonstrate that there exists a family of futures such that for each future $H[\mathcal{S}_k|\vec{X}_0 = \vec{x}] > 0$. The family of futures we use begins with all 0s and then has a 1. Intuitively, the 1 is chosen because it is a synchronizing word for the process—after observing a 1, the ϵ -machine is always in state A. Then, causal shielding will




Figure 6. The ∞ -cryptic Nemo Process.

decouple the infinite future from the first few symbols, thereby allowing us to compute the conditional entropies for the entire family of futures.

First, recall the shorthand:

$$\Pr(\mathcal{S}_k | \vec{X}_0) = \lim_{L \rightarrow \infty} \Pr(\mathcal{S}_k | X_0^L).$$

Without loss of generality, assume $k < L$. Then,

$$\begin{aligned} \Pr(\mathcal{S}_k | X_0^L) &= \frac{\Pr(X_0^k, \mathcal{S}_k, X_k^L)}{\Pr(X_0^L)} \\ &= \frac{\Pr(X_k^L | X_0^k, \mathcal{S}_k) \Pr(X_0^k, \mathcal{S}_k)}{\Pr(X_0^L)} \\ &= \frac{\Pr(X_k^L | \mathcal{S}_k) \Pr(X_0^k, \mathcal{S}_k)}{\Pr(X_0^L)}, \end{aligned}$$

where the last step is possible since the causal states are Markovian [15], shielding the past from the future. Each of these quantities is given by

$$\begin{aligned} \Pr(X_k^L = w | \mathcal{S}_k = \sigma) &= [T^{(w)} \mathbf{1}]_\sigma \\ \Pr(X_0^k = w, \mathcal{S}_k = \sigma) &= [\pi T^{(w)}]_\sigma \\ \Pr(X_0^L = w) &= \pi T^{(w)} \mathbf{1}, \end{aligned}$$

where $T^{(w)} \equiv T^{(x_0)} T^{(x_1)} \dots T^{(x_{L-1})}$, $\mathbf{1}$ is a column vector of 1s, and $T_{\sigma\sigma'}^{(x)} = \Pr(\mathcal{S}' = \sigma', X = x | \mathcal{S} = \sigma)$. To establish $H[\mathcal{S}_k | \vec{X}_0] > 0$ for any k , we rely on using values of k that are multiples of 3. So, we concentrate on the following for $n = 0, 1, 2, \dots$:

$$H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1}] > 0.$$

Since 1 is a synchronizing word, we can greatly simplify the conditional probability distribution. First, we freely include the synchronized causal state A and rewrite the conditional distribution as a fraction:

$$\begin{aligned} \Pr(\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1}) &= \Pr(\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A, \vec{X}_{3n+1}) \\ &= \frac{\Pr(\mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A, \vec{X}_{3n+1})}{\Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A, \vec{X}_{3n+1})}. \end{aligned}$$

Then, we factor everything except \vec{X}_{3n+1} out of the numerator and make use of causal shielding to simplify the conditional. For example, the numerator becomes

$$\begin{aligned} & \Pr(\mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A, \vec{X}_{3n+1}) \\ &= \Pr(\vec{X}_{3n+1} | \mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A) \\ & \quad \times \Pr(\mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A) \\ &= \Pr(\vec{X}_{3n+1} | \mathcal{S}_{3n+1} = A) \\ & \quad \times \Pr(\mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A) \\ &= \Pr(\vec{X}_{3n+1} | \mathcal{S}_{3n+1} = A) \Pr(\mathcal{S}_{3n}, X_0^{3n+1} = 0^{3n} \mathbf{1}). \end{aligned}$$

Similarly, the denominator becomes

$$\begin{aligned} & \Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}, \mathcal{S}_{3n+1} = A, \vec{X}_{3n+1}) \\ &= \Pr(\vec{X}_{3n+1} | \mathcal{S}_{3n+1} = A) \Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}). \end{aligned}$$

Combining these results, we obtain a finite form for the entropy of \mathcal{S}_{3n} conditioned on a family of infinite futures, first noting:

$$\Pr(\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1}) = \Pr(\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}).$$

Thus, for all \vec{x}_{3n+1} , we have

$$\begin{aligned} & H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1} = \vec{x}_{3n+1}] \\ &= H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}]. \end{aligned}$$

Now, we are ready to compute the conditional entropy for the entire family. First, note that $T^{(0)}$ raised to the third power is a diagonal matrix with each element equal to $(1 - p)(1 - q)$. Thus, for $j = 1, 2, 3 \dots$,

$$[T^{(0)}]_{\sigma\sigma}^{3j} = (1 - p)^j (1 - q)^j.$$

Using all of the above relations, we can easily calculate

$$\Pr(\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}) = \frac{1}{3 - 2p} \begin{pmatrix} A & B & C \\ p & 0 & q(1 - p) \end{pmatrix}.$$

Thus, for $p, q \in (0, 1)$, we have

$$\begin{aligned} & H[\mathcal{S}_{3n} | \vec{X}_0] \geq H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1}] \\ &= \sum_{\vec{x}_{3n+1}} \Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1} = \vec{x}_{3n+1}) \\ & \quad \times H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1} = \vec{x}_{3n+1}] \\ &= H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}] \\ & \quad \times \sum_{\vec{x}_{3n+1}} \Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}, \vec{X}_{3n+1} = \vec{x}_{3n+1}) \\ &= H[\mathcal{S}_{3n} | X_0^{3n+1} = 0^{3n} \mathbf{1}] \Pr(X_0^{3n+1} = 0^{3n} \mathbf{1}) \\ &= \left(\frac{p}{3 - 2p} \log_2 \frac{3 - 2p}{p} + \frac{q(1 - p)}{3 - 2p} \log_2 \frac{3 - 2p}{q(1 - p)} \right) \\ & \quad \times [(1 - p)(1 - q)]^{3n} \\ &> 0. \end{aligned}$$

So, any time k is a multiple of 3, $H[S_k|\vec{X}_0] > 0$. Finally, suppose $(k \bmod 3) = i$, where $i \neq 0$. That is, suppose k is not a multiple of 3. By lemma 1, $H[S_k|\vec{X}_0] \geq H[S_{k+i}|\vec{X}_0]$ and, since we just showed that the latter quantity is always strictly greater than zero, we conclude that $H[S_k|\vec{X}_0] > 0$ for every value of k .

The above establishes that the Nemo Process does not satisfy the k -crypticity criterion for any finite k . Thus, the Nemo Process is ∞ -cryptic. This means that we cannot make use of the k -cryptic approximation to calculate χ or \mathbf{E} .

Fortunately, the techniques introduced in [1, 2] do not rely on an approximation method. To avoid ambiguity, denote the statistical complexity we just computed as C_μ^+ . When those techniques are applied to the Nemo Process, we find that the process is causally reversible ($C_\mu^+ = C_\mu^-$) and has the following forward-reverse causal-state conditional distribution:

$$\Pr(S^+|S^-) = \frac{1}{p+q-pq} \begin{matrix} & \begin{matrix} A & B & C \end{matrix} \\ \begin{matrix} D \\ E \\ F \end{matrix} & \begin{pmatrix} p & 0 & q(1-p) \\ 0 & q & p(1-q) \\ q & p(1-q) & 0 \end{pmatrix} \end{matrix}.$$

With this, one can calculate \mathbf{E} , in a closed form, via

$$\mathbf{E} = C_\mu^+ - H[S^+|S^-].$$

(Again, calculational details are provided in [18].)

4. Conclusion

Calculating the excess entropy $I[\vec{X}; \vec{X}]$ is, at first blush, a daunting task. We ask for a mutual information between two infinite sets of random variables. Appealing to $\mathbf{E} = I[S; \vec{X}]$, we use the compact representation of the ϵ -machine to reduce one infinite set (the past) to a (usually) finite set. A process's k -crypticity captures something similar about the infinite set of future variables and allows us to further compact our form for excess entropy, reducing an infinite variable set to a finite one. The resulting stratification of process space is a novel way of thinking about its *structure* and, as long as we know in which stratum we lie, we can rapidly calculate many quantities of interest.

Unfortunately, in the general case, one will not know *a priori* a process's crypticity order. Worse, as far as we are aware, there is no known finite method for calculating the cryptic order. This strikes us as an interesting open problem and challenge.

If, by construction or by some other means, one does know it, then, as we showed, crypticity and \mathbf{E} can be calculated using the crypticity expansion. Failing this, though, one might consider using the expansion to search for the order. There is no known stopping criterion, so this search may not find k in finite time. Moreover, the expansion is a calculation that grows exponentially in computational complexity with cryptic order, as we noted. Devising a stopping criterion would be very useful to such a search.

Even without knowing the k -crypticity, the expansion is often still useful. For use in estimating \mathbf{E} , it provides us with a bound from above. This is complementary to the lower bound one finds using the typical expansion $\mathbf{E}(L) = H[X_0^L] - h_\mu L$ [17]. Using these upper and lower bounds, one may determine that for a given purpose, the estimate of χ or \mathbf{E} is within an acceptable tolerance.

The crypticity hierarchy is a revealing way to carve the space of processes in that it concerns how they hide internal state information from an observer. The examples were chosen to illustrate several features of this new view. The Even Process, a canonical example of order- ∞ Markov, resides instead at the very bottom of this ladder. The two example families

show us how k -cryptic is neither a parallel nor an independent concept to order- R Markov. Finally, we see in the last example an apparently simple process with ∞ -crypticity.

The general lesson is that internal state information need not be immediately available in measurement values, but instead may be spread over long measurement sequences. If a process is k -cryptic and k is finite, then internal state information is accessible over sequences of length k . The existence, as we demonstrated, of processes that are ∞ -cryptic is rather sobering. Interpreted as a statement of the impossibility of extracting state information, it reminds us of earlier work on hidden spatial dynamical systems that exhibit a similar encrypting of internal structure in observed spacetime patterns [19].

Due to the exponentially growing computational effort to search for the cryptic order and, concretely, the existence of ∞ -cryptic processes, the general theory introduced in [1, 2] is seen to be necessary. It allows one to directly calculate \mathbf{E} and crypticity and to do so efficiently.

Acknowledgments

Chris Ellison was partially supported on a GAANN fellowship. The Network Dynamics Program funded by Intel Corporation also partially supported this work.

References

- [1] Crutchfield J P, Ellison C J and Mahoney J R 2009 Time's barbed arrow: irreversibility, crypticity, and stored information *Phys. Rev. Lett.* in press
- [2] Ellison C J, Mahoney J R and Crutchfield J P 2009 Prediction, retrodiction, and the amount of information stored in the present *J. Stat. Phys.* in press
- [3] Fraser A 1991 Chaotic data and model building *Information Dynamics (NATO ASI Series Series B: Physics vol 256)* ed H Atmanspacher and H Scheingraber (New York: Plenum) p 125
- [4] Casdagli M and Eubank S (ed) 1992 Nonlinear modeling *SFI Studies in the Sciences of Complexity* (Reading, MA: Addison-Wesley)
- [5] Spratt J C 2003 *Chaos and Time-Series Analysis* 2nd edn (Oxford: Oxford University Press)
- [6] Kantz H and Schreiber T 2006 *Nonlinear Time Series Analysis* 2nd edn (Cambridge: Cambridge University Press)
- [7] Arnold D 1996 Information-theoretic analysis of phase transitions *Complex Syst.* **10** 143–55
- [8] Crutchfield J P and Feldman D P 1997 Statistical complexity of simple one-dimensional spin systems *Phys. Rev. E* **55** 1239R–43R
- [9] Feldman D P and Crutchfield J P 1998 Discovering non-critical organization: Statistical mechanical, information theoretic, and computational views of patterns in simple one-dimensional spin systems *Santa Fe Institute Working Paper* 98-04-026
- [10] Tononi G, Sporns O and Edelman G M 1994 A measure for brain complexity: relating functional segregation and integration in the nervous system *Proc. Natl. Acad. Sci. USA* **91** 5033–7
- [11] Bialek W, Nemenman I and Tishby N 2001 Predictability, complexity, and learning *Neural. Comput.* **13** 2409–63
- [12] Ebeling W and Poschel T 1994 Entropy and long-range correlations in literary English *Europhys. Lett.* **26** 241–6
- [13] Debowski L 2008 On the vocabulary of grammar-based codes and the logical consistency of texts *IEEE Trans. Info. Th.* (arXiv:0810.3125 [cs.IT])
- [14] Crutchfield J P and Young K 1989 Inferring statistical complexity *Phys. Rev. Lett.* **63** 105–8
- [15] Crutchfield J P and Shalizi C R 1999 Thermodynamic depth of causal states: objective complexity via minimal representations *Phys. Rev. E* **59** 275–283
- [16] Shalizi C R and Crutchfield J P 2001 Computational mechanics: pattern and prediction, structure and simplicity *J. Stat. Phys.* **104** 817–79
- [17] Crutchfield J P and Feldman D P 2003 Regularities unseen, randomness observed: levels of entropy convergence *Chaos* **13** 25–54
- [18] Mahoney J R, Ellison C J and Crutchfield J P 2009 Information accessibility and cryptic processes: linear combinations of causal states arXiv:0906.5099 [cond-mat]
- [19] Crutchfield J P 1992 Unreconstructible at any radius *Phys. Lett. A* **171** 52–60