

# Computational Mechanics of Input-Output Processes: Structured transformations and the $\epsilon$ -transducer

Nix Barnett<sup>1,2,\*</sup> and James P. Crutchfield<sup>1,2,3,†</sup>

<sup>1</sup>*Complexity Sciences Center*  
<sup>2</sup>*Mathematics Department*  
<sup>3</sup>*Physics Department*  
*University of California at Davis*  
*One Shields Avenue, Davis, CA 95616*  
(Dated: February 6, 2016)

Computational mechanics quantifies structure in a stochastic process via its causal states, leading to the process's minimal, optimal predictor—the  $\epsilon$ -machine. We extend computational mechanics to communication channels coupling two processes, obtaining an analogous optimal model—the  $\epsilon$ -transducer—of the stochastic mapping between them. Here, we lay the foundation of a structural analysis of communication channels, treating joint processes and processes with input. The result is a principled structural analysis of mechanisms that support information flow between processes. It is the first in a series on the structural information theory of memoryful channels, channel composition, and allied conditional information measures.

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

Keywords: sequential machine, communication channel, finite-state transducer, statistical complexity, causal state, minimality, optimal prediction, subshift endomorphism

## CONTENTS

I. Introduction	1	VIII. Reproducing a Channel	13
II. Processes and their Presentations	2	IX. Optimality	14
A. Stationary, Ergodic Processes	2	X. Global $\epsilon$ -Machine versus $\epsilon$ -Transducer	18
B. Examples	3	XI. History $\epsilon$ -Transducer versus Generator $\epsilon$ -Transducer	18
1. Biased Coin Process	3	XII. Examples Revisited	18
2. Period-2 Process	3	XIII. Infinite-state $\epsilon$ -Transducers: The Simple Nonunifilar Channel	22
3. Golden Mean Process	3	XIV. Discussion	23
4. Even Process	3	A. Related Work: Modeling	24
C. Optimal Presentations	3	B. Related Work: Classification	24
D. Example Process $\epsilon$ -Machines	5	C. Applications	25
III. Input-Output Processes and Channels	5	1. Inference versus Experimentation	25
IV. Example Channels and their Classification	7	2. Information Flow within and between Systems	25
A. Memorylessness	8	3. Process Decomposition	25
B. Finite Feedforward	8	4. Perception-Action Cycles	26
C. Infinite Feedforward	9	XV. Conclusion	26
D. Finite Feedback	9	Acknowledgments	26
E. Infinite Feedback	10	A. Equivalence of Two $\epsilon$ -Transducer Definitions	27
F. Infinite Feedforward-Feedback	10	References	28
G. Irreducible Feedforward-Feedback	10		
H. Causal Channel Markov Order Hierarchy	11		
I. Causal-State Channels	11		
V. Global $\epsilon$ -Machine	11		
VI. $\epsilon$ -Transducer	12		
VII. Structural Complexity	13		

## I. INTRODUCTION

Arguably, the distinctive character of natural and engineered systems lies in their organization. This is in contrast to differences, say, in how random they are or in their temperature. Computational mechanics provides

\* nix@math.ucdavis.edu; <http://math.ucdavis.edu/~nix/>

† chaos@ucdavis.edu; <http://csc.ucdavis.edu/~chaos/>

an analytical and constructive way to determine a system's organization [1], supplementing the well developed statistical physics view of systems in terms of disorder—via thermodynamic entropy and free energy. The contrast begs a question, though, how does organization arise?

As a step to an answer, we extend computational mechanics from describing individual systems to analyze organization in transformations between systems. We build on its well established methods to describe a calculus for detecting the emergence of structure. Indeed, natural systems are nothing, if not the result of transformations of energy, information, or both. There is no lack of examples: Filtering measurement time series to discover the temporal behavior of hidden states [2]; Maxwell's Demon translating measurement information to control actions that rectify thermal fluctuations into work, locally defeating Thermodynamics' Second Law [3–5]; sensory transduction in the retina that converts light intensity and spectra into neural spike trains [6]; perception-action cycles in which an agent decides its future behavior based on its interpretation of its environment's state [7, 8]; and, finally, firms that transform raw materials into finished goods [9].

We refer to our objects of study as *structured transformations* to emphasize the particular focus on basic questions of organization. How complex is a transformation? What and how many resources are required to implement it? What does it add to an input in producing an output? Randomness, structure, or both? What is thrown away? How is its own organization reflected in that of its output?

The framework addresses these questions, both quantitatively and from first principles. It is the first in a series. Foremost, it's burden is to lay the groundwork necessary to answer these questions. Sequels introduce information measures to classify the kinds of information processing in joint input-output processes and in structured transformations. To delineate the underlying mechanisms that produce organization, they analyze a range of examples and describe a number of interesting, even counterintuitive, properties of structured transformations.

The questions posed are basic, so there is a wide range of applications and of historical precedents. Due to this diversity and to avoid distracting from the development, we defer reviewing them and related work until later, once the context has been set and the focus, benefits, and limitations of our approach are clear.

The following analyzes communication channels and channel composition in terms of intrinsic computation [1, 10]. As such, it and the entire series, for that matter, assume familiarity with stochastic processes at the level of Ref. [11], information theory at the level of Refs. [12, 13], and computational mechanics at the level of Refs. [14, 15]. These serve as the default sources for statements in our review.

We first present a brief overview and relevant notation for how computational mechanics describes pro-

cesses. We extend this to describe controlled processes—processes with input. Several examples that illustrate the basic kinds of input-output process are then given, by way of outlining an organizational classification scheme. At that point, we describe the global  $\epsilon$ -machine for joint input-output processes. Using this we introduce the  $\epsilon$ -transducer, defining its structural complexity and establishing its optimalities. We close with a thorough analysis of the example input-output processes and compare and contrast  $\epsilon$ -transducers with prior work.

## II. PROCESSES AND THEIR PRESENTATIONS

### A. Stationary, Ergodic Processes

The temporal *stochastic process* we consider is a one-dimensional chain  $\dots Y_{-2}Y_{-1}Y_0Y_1Y_2\dots$  of discrete random variables  $\{Y_t\}_{t \in \mathbb{Z}}$  that takes values  $\{y_t\}_{t \in \mathbb{Z}}$  over a finite or countable alphabet  $\mathcal{Y}$ . A finite block  $Y_iY_{i+1}\dots Y_{j-1}$  of variables with  $t \in [i, j)$  is denoted  $Y_{i:j}$ . The left index is always inclusive, the right always exclusive. Let  $\overleftarrow{Y}$  denote the bi-infinite chain  $Y_{-\infty:\infty}$ , and let  $\overleftarrow{\mathcal{Y}}$  denote the set of all bi-infinite sequences  $\overleftarrow{y}$  with alphabet  $\mathcal{Y}$ .  $\overleftarrow{Y}_t = \dots Y_{t-2}Y_{t-1}$  is the *past* leading up to time  $t$ , not including  $t$ , and  $\overrightarrow{Y}_t = Y_tY_{t+1}\dots$  is the *future* leading from it, including  $t$ .

We can also use the time index notation above to specify the time origin of a process or bi-infinite sequence. This is useful when we are comparing two such processes or sequences. For example, if we wish to say that  $\overleftarrow{Y}$  is  $\overleftarrow{X}$  delayed by three time steps, this can be written as  $\overleftarrow{Y}_3 = \overleftarrow{X}_0$ . To indicate the time origin in *specific* realized symbol values in a chain, we place a period before the symbol occurring at  $t = 0$ : e.g.,  $\overleftarrow{y} = \dots acbb.caba\dots$ , where  $y_0 = c$ . Finally, if the word  $abcd$  follows a sequence of random variables  $Y_{i:j}$  we denote this by  $Y_{i:j+4} = Y_{i:j}abcd$ , for example. A word occurring before a sequence of random variables is denoted by a similar concatenation rule; e.g.,  $Y_{i-4:j} = abcdY_{i:j}$ .

A stochastic process is defined by its *word distributions*:

$$\mathbb{P}(Y_{t:t+L}) \equiv \{\mathbb{P}(Y_{t:t+L} = y_{t:t+L})\}_{y_{t:t+L} \in \mathcal{Y}^L}, \quad (1)$$

for all  $L \in \mathbb{Z}^+$  and  $t \in \mathbb{Z}$ .

We will often use an equivalent definition of a stochastic process as a random variable defined over the set of bi-infinite sequences  $\overleftarrow{\mathcal{Y}}$ . In this case, a stochastic process is defined by an indexed set of probabilities of the form:

$$\mathbb{P}(\overleftarrow{Y}) \equiv \{\mathbb{P}(\overleftarrow{Y} \in \sigma)\}_{\sigma \subseteq \overleftarrow{\mathcal{Y}}}, \quad (2)$$

where  $\sigma$  is a measurable set of bi-infinite sequences. We can obtain a process' word probabilities by selecting appropriate measurable subsets—*cylinder sets*—that correspond to holding the values of a contiguous subset of

random variables fixed.

In the following, we consider only *stationary processes*—those invariant under time translation:

$$\begin{aligned}\mathbb{P}(Y_{t:t+L}) &= \mathbb{P}(Y_{0:L}) \text{ and} \\ \mathbb{P}(\overrightarrow{Y}_t) &= \mathbb{P}(\overrightarrow{Y}_0),\end{aligned}$$

for all  $t \in \mathbb{Z}$  and  $L \in \mathbb{Z}^+$ . This property ensures that a process's behavior has no *explicit* dependence on time origin.

We will also primarily limit the discussion to *ergodic* stationary processes—processes where any realization  $\overleftarrow{y}$  gives good empirical estimates  $\widehat{\mathbb{P}}(Y_{0:L})$  of the process's true word probabilities  $\mathbb{P}(Y_{0:L})$  [16]. That is, for any finite realization  $y_{0:M}$ , the empirical estimate  $\widehat{\mathbb{P}}(w)$  of a word  $w = w_0 w_1 \dots w_{L-1}$  of length  $L$ , converges almost surely to the true process probability  $\mathbb{P}(Y_{0:L} = w)$  as  $M \rightarrow \infty$ , where:

$$\widehat{\mathbb{P}}(w) \equiv \sum_{t=0}^{M-L} \frac{\mathbf{I}_w(y_{t:t+L})}{M-L+1},$$

and the indicator function  $\mathbf{I}_w(y_{t:t+L})$  equals 1 when  $y_{t:t+L} = w$ , and 0 otherwise. This property ensures, among other things, that any particular realization of the process reflects the behavior of the process in general—a useful property when we do not have freedom to re-initialize the system at will.

## B. Examples

To illustrate key ideas in the development, we use several example processes, all with the binary alphabet  $\mathcal{Y} = \{0, 1\}$ . They are used repeatedly in the following and in the sequels.

### 1. Biased Coin Process

The *Biased Coin Process* is an independent, identically distributed (IID) process, where word probabilities factor into a product of single-variable probabilities:

$$\mathbb{P}(Y_{0:L}) = \mathbb{P}(Y_0)\mathbb{P}(Y_1)\cdots\mathbb{P}(Y_{L-1}),$$

where  $\mathbb{P}(Y_t) = \mathbb{P}(Y_0)$  for all  $t$ . If  $n$  is the number of 0s in  $y_{0:L}$ , then  $\mathbb{P}(y_{0:L}) = p^n(1-p)^{L-n}$ , where  $p = \mathbb{P}(Y_0 = 0)$ .

### 2. Period-2 Process

The *Period-2 Process* endlessly repeats the word 01, starting with either a 0 or a 1 with equal probability. It

is specified by the word distributions:

$$\begin{aligned}\mathbb{P}(Y_{i:j} = (\dots 1010.1010 \dots)_{i:j}) &= \frac{1}{2} \text{ and} \\ \mathbb{P}(Y_{i:j} = (\dots 0101.0101 \dots)_{i:j}) &= \frac{1}{2},\end{aligned}$$

where  $i, j \in \mathbb{Z}, i < j$ .

### 3. Golden Mean Process

The *Golden Mean Process* generates all binary sequences, except those with consecutive 0s. After a 1 is generated, the next 0 or 1 appears with equal likelihood. Its word distributions are determined by a Markov chain with states  $Y_0 = 0$  and  $Y_0 = 1$  having probabilities  $\mathbb{P}(Y_0 = 0) = 1/3$  and  $\mathbb{P}(Y_0 = 1) = 2/3$ , respectively, and transition probabilities:

$$\begin{aligned}\mathbb{P}(Y_1 = 0|Y_0 = 0) &= 0, \\ \mathbb{P}(Y_1 = 1|Y_0 = 0) &= 1, \\ \mathbb{P}(Y_1 = 0|Y_0 = 1) &= \frac{1}{2}, \text{ and} \\ \mathbb{P}(Y_1 = 1|Y_0 = 1) &= \frac{1}{2}.\end{aligned}\tag{3}$$

### 4. Even Process

The *Even Process* generates all binary sequences, except that a 1 always appears in an even-length block of 1s bordered by 0s. After an even number of 1s are generated, the next 0 or 1 appears with equal likelihood. Notably, the Even Process cannot be represented by a Markov chain of any finite order or, equivalently, by any finite list of conditional probabilities over words. As we will see, though, it can be represented concisely by a finite-state hidden Markov model. To see this, we must first introduce alternative models—here called *presentations*—for a given process.

## C. Optimal Presentations

We can completely, and rather prosaically, specify a stationary process by listing its set of word probabilities, as in Eq. (1), or conditional probabilities, as in Eq. (3). As with the Even Process, however, generally these sets are infinite and so one prefers to use finite or, at least, more compact descriptions. Fortunately, there is a canonical *presentation* for any stationary process—the  $\epsilon$ -machine of computational mechanics [10, 14]—that is its unique, optimal, unifilar generator of minimal size; which we now define.

Given a process's word distribution, its  $\epsilon$ -machine is constructed by regarding any two infinite pasts  $\overleftarrow{y}$  and  $\overleftarrow{y}'$  as equivalent when they lead to the same distribution over infinite futures—the same *future morphs*, of the form  $\mathbb{P}(\overrightarrow{Y}|\overleftarrow{y})$ . This grouping is given by the *causal equivalence*

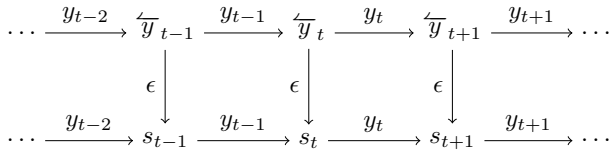


FIG. 1. Process lattice: The dynamic inherited by the causal states  $s_t = \sigma_i \in \mathcal{S}$  from a process's pasts via the  $\epsilon$ -map:  $s_t = \epsilon(\dots y_{t-2}y_{t-1}) \xrightarrow{y_t} s_{t+1} = \epsilon(\dots y_{t-2}y_{t-1}y_t)$ .

relation  $\sim_\epsilon$ :

$$\overleftarrow{y} \sim_\epsilon \overleftarrow{y}' \iff \mathbb{P}(\overrightarrow{Y} | \overleftarrow{Y} = \overleftarrow{y}) = \mathbb{P}(\overrightarrow{Y} | \overleftarrow{Y} = \overleftarrow{y}') . \quad (4)$$

The equivalence classes of  $\sim_\epsilon$  partition the set  $\overleftarrow{\mathcal{Y}}$  of all allowed pasts. The classes are the process's *causal states*. The indexed set of causal states is denoted by  $\mathcal{S}$  and has elements  $\sigma_i$ . Note that  $i$  is not a time index, but an element of an index set with the same cardinality as  $\mathcal{S}$ . The associated random variable over alphabet  $\mathcal{S}$  is denoted by  $S$ . The  $\epsilon$ -map is a function  $\epsilon : \overleftarrow{\mathcal{Y}} \rightarrow \mathcal{S}$  that takes a given past to its corresponding causal state or, equivalently, to the set of pasts to which it is causally equivalent:

$$\begin{aligned} \epsilon(\overleftarrow{y}) &= \sigma_i \\ &= \{ \overleftarrow{y}' : \overleftarrow{y} \sim_\epsilon \overleftarrow{y}' \} . \end{aligned}$$

The  $\epsilon$ -map induces a dynamic over pasts: Appending a new symbol  $y_t$  to past  $\overleftarrow{y}_t$  produces a new past  $\overleftarrow{y}_{t+1} = \overleftarrow{y}_t y_t$ . This, in turn, defines a stochastic process—the *causal-state process*—with random variable chain  $\overleftarrow{S} = S_{-\infty:\infty} = \dots S_{-1}S_0S_1\dots$ , where at time  $t$  each  $S_t$  takes on some value  $s_t = \sigma_i \in \mathcal{S}$ . The relationship between a process's pasts and its causal states is summarized in Fig. 1. The map from the observed chain  $\overrightarrow{Y}$  to the internal state chain  $\overleftarrow{S}$  is the *causal state filter*. (We return to this mapping later on.) The dynamic over causal states is specified by an indexed set  $\mathcal{T}$  of *symbol-labeled transition matrices*:

$$\mathcal{T} \equiv \{ T^{(y)} \}_{y \in \mathcal{Y}} ,$$

where  $T^{(y)}$  has elements:

$$T_{ij}^{(y)} = \mathbb{P}(S_1 = \sigma_j, Y_0 = y | S_0 = \sigma_i) .$$

The  $\epsilon$ -map also induces a distribution  $\pi$  over causal states, with elements:

$$\begin{aligned} \pi(i) &= \mathbb{P}(S_0 = \sigma_i) \\ &= \mathbb{P}(\epsilon(\overleftarrow{Y}) = \sigma_i) . \end{aligned}$$

Since the process is stationary and the  $\epsilon$ -map is time independent,  $\pi$  is also stationary. We therefore refer to  $\pi$  as the process's *stationary distribution*. Note that for ergodic processes, the stationary distribution can be calculated directly from the transition matrices alone [17].

The tuple  $(\mathcal{Y}, \mathcal{S}, \mathcal{T})$  consisting of the process' alphabet, causal state set, and transition matrix set is the process'  $\epsilon$ -machine.

The  $\epsilon$ -machine is a process's unique, maximally predictive, minimal-size unifilar presentation [10, 18, 19]. In other words, the causal states are as predictive as any rival partition  $\mathcal{R}$  of the pasts  $\overleftarrow{\mathcal{Y}}$ . In particular, any given causal state  $\sigma_i$  is as predictive as any of its pasts  $\overleftarrow{y} \in \epsilon^{-1}(\sigma_i)$ . Measuring predictive ability via the mutual information between states and future observations, this translates to the statement that:

$$\begin{aligned} \mathbb{I}[S_0; \overrightarrow{Y}_0] &= \mathbb{I}[\overleftarrow{Y}_0; \overrightarrow{Y}_0] \\ &\geq \mathbb{I}[R_0; \overrightarrow{Y}_0] , \end{aligned}$$

where  $R_0 \in \mathcal{R}$ . Moreover, among all equally predictive (*prescient rival*) partitions  $\widehat{\mathcal{R}}$  of the past, the causal states minimize the state Shannon entropy:  $\mathbb{H}[S] = \mathbb{H}[\pi] \leq \mathbb{H}[\widehat{\mathcal{R}}]$ . Due to the  $\epsilon$ -machine's minimality, the *statistical complexity*  $C_\mu = \mathbb{H}[S] = \mathbb{H}[\pi]$  measures the amount of historical information a process stores.

A process's  $\epsilon$ -machine presentation has several additional properties that prove useful. First, the causal states form a Markov chain. This means that the  $\epsilon$ -machine is a type of *hidden Markov model*. Second, the causal states are *unifilar*:

$$\mathbb{H}[S_{t+1} | Y_t, S_t] = 0 .$$

That is, the current state and symbol uniquely determine the next state. This is necessary for an observer to maintain its knowledge of a process's current causal state while scanning a sequence of symbols. Third, unlike general (that is, nonunifilar) hidden Markov models, one can calculate a process's key informational properties directly from its  $\epsilon$ -machine. For example, a process's *entropy rate*  $h_\mu$  can be written in terms of the causal states:

$$h_\mu = \mathbb{H}[Y_0 | S_0] .$$

And, using the methods of Refs. [14, 20], a process's past-future mutual information—the *excess entropy*  $\mathbf{E}$ —is given by its forward  $\mathcal{S}^+$  and reverse  $\mathcal{S}^-$  causal states:

$$\begin{aligned} \mathbf{E} &\equiv \mathbb{I}[\overleftarrow{Y}; \overrightarrow{Y}] \\ &= \mathbb{I}[\mathcal{S}^-; \mathcal{S}^+] . \end{aligned} \quad (5)$$

Generally, the excess entropy is only a lower bound on the information that must be stored in order to predict a process:  $\mathbf{E} \leq C_\mu$ . This difference is captured by the process's *crypticity*  $\chi = C_\mu - \mathbf{E}$ .

Since they are conditioned on semi-infinite pasts, the causal states defined above correspond to *recurrent* states in a process's complete  $\epsilon$ -machine presentation. They capture a process's time-asymptotic behavior. An  $\epsilon$ -machine also has *transient* causal states that arise when conditioning on finite-length pasts, as well as a unique

start state, which can be either transient or recurrent. When the underlying process is ergodic, they can be derived from the recurrent causal states using the mixed-state method of Ref. [14]. In general, they can be obtained from a modified causal equivalence relation, extended to include finite pasts. We omit them, unless otherwise noted, from our development.

The preceding introduced what is known as the *history* specification of an  $\epsilon$ -machine: From a stationary, ergodic process, one derives its  $\epsilon$ -machine by applying equivalence relation Eq. (4). There is also the complementary *generator* specification: As a generator, rather than using equivalence classes over a process's histories, the  $\epsilon$ -machine is defined as a strongly connected hidden Markov model whose transitions are unifilar and whose states are probabilistically distinct. Such an  $\epsilon$ -machine generates a unique, ergodic, stationary process, and is the same  $\epsilon$ -machine that we would obtain by applying the causal equivalence relation to said generated process [21]. The following uses both history and generator  $\epsilon$ -machines; which will be clear in context.

#### D. Example Process $\epsilon$ -Machines

The cardinality of a process's causal state set  $\mathcal{S}$  need not be finite or even countable; see, e.g., Figs. 7, 8, 10, and 17 in Ref. [22]. For simplicity in the following, though, we restrict our study to processes with a finite or countably infinite number of causal states. This allows us to depict a process graphically, showing its  $\epsilon$ -machine as an *edge-labeled directed graph*. Nodes in the graph are causal states and edges, transitions between them. A transition from state  $\sigma_i$  to state  $\sigma_j$  while emitting symbol  $y$  is represented as an edge connecting the corresponding nodes that is labeled  $\mathbf{y} : p$ . (Anticipating our needs later on, this differs slightly from prior notation.) Here,  $p = T_{ij}^{(y)}$  is the state transition probability and  $y$  is the symbol emitted on the transition. Figure 2 displays  $\epsilon$ -machine state-transition diagrams for the example processes.

Since the Biased Coin Process's current output is independent of the past, all pasts are causally equivalent. This leads to a single causal state  $A$ , occupied with probability 1. (See Fig. 2(a).) Therefore, the Biased Coin Process has a statistical complexity of  $C_\mu = 0$  bits. Its excess entropy  $\mathbf{E}$  also vanishes. This example illustrates the general property that IID processes lack causal structure. They can be quite random; the example here has an entropy rate of  $h_\mu = H(2/3) \approx 0.918$  bits per step, where  $H(p)$  is the binary entropy function.

In contrast, the Period-2 Process has two causal states, call them  $A$  and  $B$ , that correspond to pasts that end in either a 1 or a 0, respectively. (See Fig. 2(b).) Since the states are occupied with equal probability, the Period-2 Process has  $C_\mu = 1$  bit of stored information. In this case,  $\mathbf{E} = C_\mu$ . It is perfectly predictable, with  $h_\mu = 0$  bits per step.

The two causal states of the Golden Mean Process also

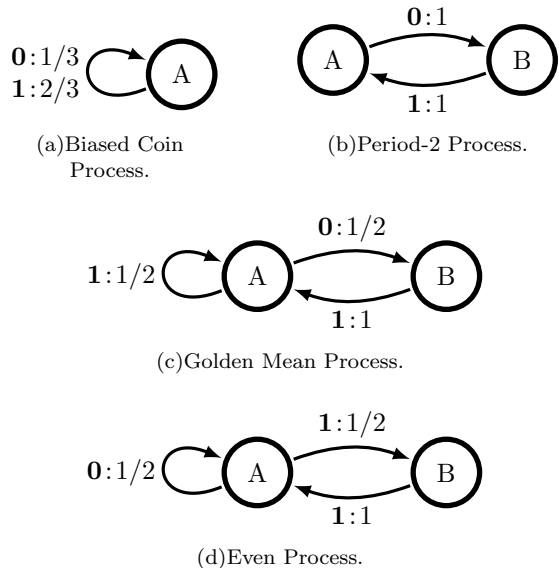


FIG. 2.  $\epsilon$ -Machines for the example processes. Transitions labeled  $\mathbf{y} : p$ , where  $p = T_{ij}^{(y)}$  is the state transition probability and  $y$  is the symbol emitted on the transition.

correspond to pasts ending with either a 0 or 1, but for it the state transition structure ensures that no consecutive 0s are generated. (See Fig. 2(c).) Causal state  $A$  is occupied with  $\mathbb{P}(A) = 2/3$  and state  $B$  with  $\mathbb{P}(B) = 1/3$ , giving  $C_\mu = H(2/3) \approx 0.918$  bits. Note that the excess entropy is substantially less— $\mathbf{E} \approx 0.2516$  bits—indicating that the process is *cryptic*. We therefore must store additional state information above and beyond  $\mathbf{E}$  in order to predict the process [20]. Its entropy rate is  $h_\mu = H(2/3) \approx 0.918$  bits per step.

As mentioned already, the Even Process cannot be represented by a finite Markov chain. It can be represented by a finite hidden Markov model, however. In particular, its  $\epsilon$ -machine provides the most compact presentation—a two-state hidden Markov model. Causal state  $A$  corresponds to pasts that end with an even block of 1s, and state  $B$  corresponds to pasts that end with an odd block of 1s. (See Fig. 2(d).) Since the probability distribution over states is the same as that of the Golden Mean Process, the Even process has  $C_\mu \approx 0.918$  bits of statistical complexity and  $h_\mu \approx 0.918$  bits per step. In contrast, the excess entropy and statistical complexity are equal:  $\mathbf{E} = C_\mu$ . The Even Process is not cryptic.

### III. INPUT-OUTPUT PROCESSES AND CHANNELS

Up to this point, we focused on a stochastic process and a particular canonical presentation of a mechanism—the  $\epsilon$ -machine—that can generate it. We now turn to our main topic, generalizing the  $\epsilon$ -machine presentation of a given process to a presentation of a controlled process—

that is, to input-output (I/O) processes. I/O processes are related to probabilistic extensions of Moore's sequential machines [23, 24], probabilistic extensions of codes from symbolic dynamics [25], and, perhaps more naturally, to Shannon's communication channels [26]. And so, we refer to them simply as *channels*. There are important differences from how channels are developed in standard treatments [12], though. The principal difference is that we consider channels with memory, while the latter in its elementary treatment, at least, typically considers memoryless channels or channels with restricted forms of memory [27]. In addition, with an eye to applications, the framework here is adapted to *reconstruct* channels from observations of an input-output process. (A topic to which we return at the end.) Finally, the development places a unique emphasis on detecting and analyzing structure inherent in I/O processes. Our development parallels that for  $\epsilon$ -machines; see Ref. [1, and citations therein].

To begin, we define I/O processes. To make these concrete, we provide several example channels, leveraging the example processes and their  $\epsilon$ -machines already described. The examples naturally lead to the desired extension—the  $\epsilon$ -transducer. The development then turns to the main properties of  $\epsilon$ -transducers.

Loosely speaking, a channel defines a coupling between stochastic processes. It can be memoryful, probabilistic, and even anticipate the future. In other words, the channel's current output symbol may depend probabilistically upon its past, current, and future input and output symbols. As such, we will be led to generalize the memoryless and anticipation-free communication channels primarily studied in elementary information theory [12].

**Definition 1.** A channel  $\overleftarrow{Y}|\overleftarrow{X}$  with input alphabet  $\mathcal{X}$  and output alphabet  $\mathcal{Y}$  is a collection of stochastic processes over alphabet  $\mathcal{Y}$ , where each such process  $\overleftarrow{Y}|\overleftarrow{x}$  corresponds to a bi-infinite input sequence in  $\overleftarrow{\mathcal{X}}$ :

$$\overleftarrow{Y}|\overleftarrow{X} \equiv \{\overleftarrow{Y}|\overleftarrow{x}\}_{\overleftarrow{x} \in \overleftarrow{\mathcal{X}}} \quad (6)$$

That is, each fixed realization  $\overleftarrow{x} = \dots x_{-1}x_0x_1\dots$  over input alphabet  $\mathcal{X}$  is mapped to a stochastic process  $\overleftarrow{Y}|\overleftarrow{x} = \overleftarrow{Y}|\dots x_{-1}x_0x_1\dots$  over alphabet  $\mathcal{Y}$ .

If we are given a process  $\overleftarrow{X}$ —an *input process*—then a channel maps this distribution over sequences to a joint process  $(\overleftarrow{X}, \overleftarrow{Y})$ , which can be marginalized to obtain a process  $\overleftarrow{Y}$ —the *output process*. We can characterize a channel by an indexed set  $\mathcal{P}$  of *conditional word probabilities*:

$$\begin{aligned} \mathbb{P}(Y_{t:t+L}|\overleftarrow{x}) \\ \equiv \{\mathbb{P}(Y_{t:t+L} = y_{t:t+L} | \overleftarrow{X} = \overleftarrow{x})\}_{y_{t:t+L} \in \mathcal{Y}^L, \overleftarrow{x} \in \overleftarrow{\mathcal{X}}} \end{aligned}$$

Equivalently, we can represent a channel as a distribution over bi-infinite sequences (output) conditioned on a particular bi-infinite sequence (input). In other words, a

channel can be characterized by an indexed set:

$$\mathbb{P}(\overleftarrow{Y}|\overleftarrow{x}) \equiv \{\mathbb{P}(\overleftarrow{Y} \in \sigma | \overleftarrow{X} = \overleftarrow{x})\}_{\sigma \subseteq \overleftarrow{\mathcal{Y}}, \overleftarrow{x} \in \overleftarrow{\mathcal{X}}}.$$

Given an input process' distribution  $\mathbb{P}(\overleftarrow{X})$  and a channel's distribution  $\mathbb{P}(\overleftarrow{Y}|\overleftarrow{x})$  for all inputs  $\overleftarrow{x}$ , we obtain the output process' distribution as follows:

$$\begin{aligned} \mathbb{P}(\overleftarrow{Y}) &= \int \mathbb{P}(\overleftarrow{Y}, \overleftarrow{x}) d\overleftarrow{x} \\ &= \int \mathbb{P}(\overleftarrow{Y}|\overleftarrow{x})\mathbb{P}(\overleftarrow{x}) d\overleftarrow{x}, \end{aligned}$$

where the first integrand shows the appearance of the intermediate joint process  $(\overleftarrow{X}, \overleftarrow{Y})$ .

Let's say a few words about definitional choices made up to this point. As defined, the channels we consider are *total* (defined for every possible input sequence). One could extend the definition to allow for *partial* channels (defined only for a subset of possible sequences), but we do not consider such channels in what follows. This is the primary reason for defining a channel's domain in terms of bi-infinite sequences rather than say, collections of finite words. Such channels would not necessarily be total. Also, requiring that channels be defined for *every finite* input word is restrictive, as even the simplest channels may not have well defined behavior for say, a single symbol input word. We could instead define channels over a *subset* of all finite input words and explicitly add in the restriction that the channel be total, but this is still more restrictive than the  $\overleftarrow{\mathcal{X}}$  definition above. Consider, for example, the channel that outputs all 1s if its input sequence contains at least one 1 and outputs all 0s otherwise. Such a channel is undefined for any finite input word that consists of all 0s, but is well defined for any bi-infinite binary sequence. It is also true that any total channel defined over finite input words can be trivially defined over bi-infinite sequences by appending an arbitrary infinite past and future to each finite word (without changing the channel's output behavior).

Stationarity is as useful a property for channels as it is for processes.

**Definition 2.** A channel is stationary if and only if its probability distributions are invariant under time translation:

$$\begin{aligned} \mathbb{P}(Y_{t:t+L}|\overleftarrow{x}_t) &= \mathbb{P}(Y_{0:L}|\overleftarrow{x}_0) \text{ and} \\ \mathbb{P}(\overleftarrow{Y}_t|\overleftarrow{x}_t) &= \mathbb{P}(\overleftarrow{Y}_0|\overleftarrow{x}_0), \end{aligned}$$

for all  $t \in \mathbb{Z}$ ,  $L \in \mathbb{Z}^+$ , and every input sequence  $\overleftarrow{x}$ .

An immediate consequence is that stationary channels map stationary input processes to stationary output processes:

**Proposition 1** (Stationarity Preservation). *When a stationary channel is applied to a stationary input process, the resulting output process is also stationary.*

*Proof.* One calculates directly:

$$\begin{aligned}\mathbb{P}(\overrightarrow{Y}_t) &= \int \mathbb{P}(\overrightarrow{Y}_t | \overrightarrow{x}_t) \mathbb{P}(\overrightarrow{x}_t) d\overrightarrow{x}_t \\ &= \int \mathbb{P}(\overrightarrow{Y}_0 | \overrightarrow{x}_0) \mathbb{P}(\overrightarrow{x}_0) d\overrightarrow{x}_0 \\ &= \mathbb{P}(\overrightarrow{Y}_0),\end{aligned}$$

where the second equality follows from stationarity of both input process and channel, once we note that  $\dots dx_{t-1} dx_t dx_{t+1} \dots = \dots dx_{-1} dx_0 dx_1 \dots$ , since under shifted indexing the relationships between  $x_i$  and infinitesimals  $dx_i$  are preserved and the latter themselves do not change.  $\square$

We also primarily restrict our discussion to ergodic channels [28].

**Definition 3.** An ergodic channel is a channel that maps any ergodic input process  $\overrightarrow{X}$  to an ergodic joint process  $(\overrightarrow{X}, \overrightarrow{Y})$ .

Another important channel property is that of causality.

**Definition 4.** A causal channel is anticipation-free:

$$\mathbb{P}(Y_{t:t+L} | \overrightarrow{X}) = \mathbb{P}(Y_{t:t+L} | \overrightarrow{X}_{t+L}).$$

That is, the channel has well defined behavior on semi-infinite input pasts and is completely characterized by that behavior.

Channel causality is a reasonable assumption when a system has no access to future inputs. However, as a note of caution in applying the following to analyze, say, large-scale systems with many components, the choice of observables may lead to input-output processes that violate causality. For example, treating spatial configurations of one-dimensional spin systems or cellular automata as if they were time series—a somewhat common strategy—violates causality. In the following, though, we assume channel stationarity and causality, unless stated otherwise.

It is worth noting that causality is often not a severe restriction. Specifically, the following results extend to channels with finite anticipation—channels whose current output depends upon  $N$  future inputs. When both the input process and channel are stationary, one delays the appearance of the channel output by  $N$  time indices. This does not change the output process, but converts *finite anticipation* to finite channel memory and delayed output. In this way, it is possible to apply the analysis to follow to channels with anticipation directly, but the optimality theorems established must be modified slightly.

#### IV. EXAMPLE CHANNELS AND THEIR CLASSIFICATION

To motivate our main results, consider several example stationary causal channels with binary input and output. Figure 3 illustrates the mapping from input words  $\{x_{0:L} = x_0 x_1 \dots x_{L-1}\}$  to output words  $\{y_{0:L} = y_0 y_1 \dots y_{L-1}\}$  for the example channels. The *wordmaps* there treat each input-output pair  $(x_{0:L}, y_{0:L})$  as a point  $p = (p_x, p_y)$  in the unit square, where the input (output) word forms the binary expansion of  $p_x = 0.x_0 x_1 \dots x_{L-1}$  ( $p_y = 0.y_0 y_1 \dots y_{L-1}$ ). Since the channels are defined for all inputs, the plots show every input word.

We organize the examples around a classification scheme paralleling that used in signal processing to highlight memory and the nature of feedback [29].

We describe channel behavior using recurrence relations of the form:

$$Y_t \sim r(\overrightarrow{X}_t, \overrightarrow{Y}_t, X_t),$$

where  $r(\cdot)$  is either a logical function, in simple cases, or a distribution. Note that the next output  $Y_t$  depends only on the past—the channels are causal.

In general, such a recurrence relation only captures the channel's *present* behavior. Fortunately, a causal channel's future behavior is summarized by its present behavior:

$$\begin{aligned}\mathbb{P}(Y_{0:L} = ab\dots yz | \overrightarrow{(X, Y)}_0, X_{0:L} = \alpha\beta\dots\psi\omega) \\ &= \mathbb{P}(Y_0 = a | \overrightarrow{(X, Y)}_0, X_{0:L} = \alpha\dots\omega) \\ &\quad \cdot \mathbb{P}(Y_1 = b | \overrightarrow{(X, Y)}_0, X_{0:L} = \alpha\dots\omega, Y_0 = a) \\ &\quad \cdots \mathbb{P}(Y_{L-1} = z | \overrightarrow{(X, Y)}_0, X_{0:L} = \alpha\dots\omega, Y_{0:L-1} = a\dots y) \\ &\stackrel{(a)}{=} \mathbb{P}(Y_0 = a | \overrightarrow{(X, Y)}_0, X_{0:L} = \alpha\dots\omega) \\ &\quad \cdot \mathbb{P}(Y_1 = b | \overrightarrow{(X, Y)}_1, X_{1:L} = \beta\dots\omega) \\ &\quad \cdots \mathbb{P}(Y_{L-1} = z | \overrightarrow{(X, Y)}_{L-1}, X_{L-1} = \omega) \\ &\stackrel{(b)}{=} \mathbb{P}(Y_0 = a | \overrightarrow{(X, Y)}_0, X_0 = \alpha) \\ &\quad \cdot \mathbb{P}(Y_1 = b | \overrightarrow{(X, Y)}_1, X_1 = \beta) \\ &\quad \cdots \mathbb{P}(Y_{L-1} = z | \overrightarrow{(X, Y)}_{L-1}, X_{L-1} = \omega) \\ &\stackrel{(c)}{=} \mathbb{P}(Y_0 = a | \overrightarrow{(X, Y)}_0, X_0 = \alpha) \\ &\quad \cdot \mathbb{P}(Y_0 = b | \overrightarrow{(X, Y)}_0, X_0 = \beta) \\ &\quad \cdots \mathbb{P}(Y_0 = z | \overrightarrow{(X, Y)}_0, X_0 = \omega),\end{aligned}$$

where in (a) we merge individual variables into pasts to obtain new pasts, in (b) we remove input variables that have no effect due to causality, and (c) follows from stationarity.

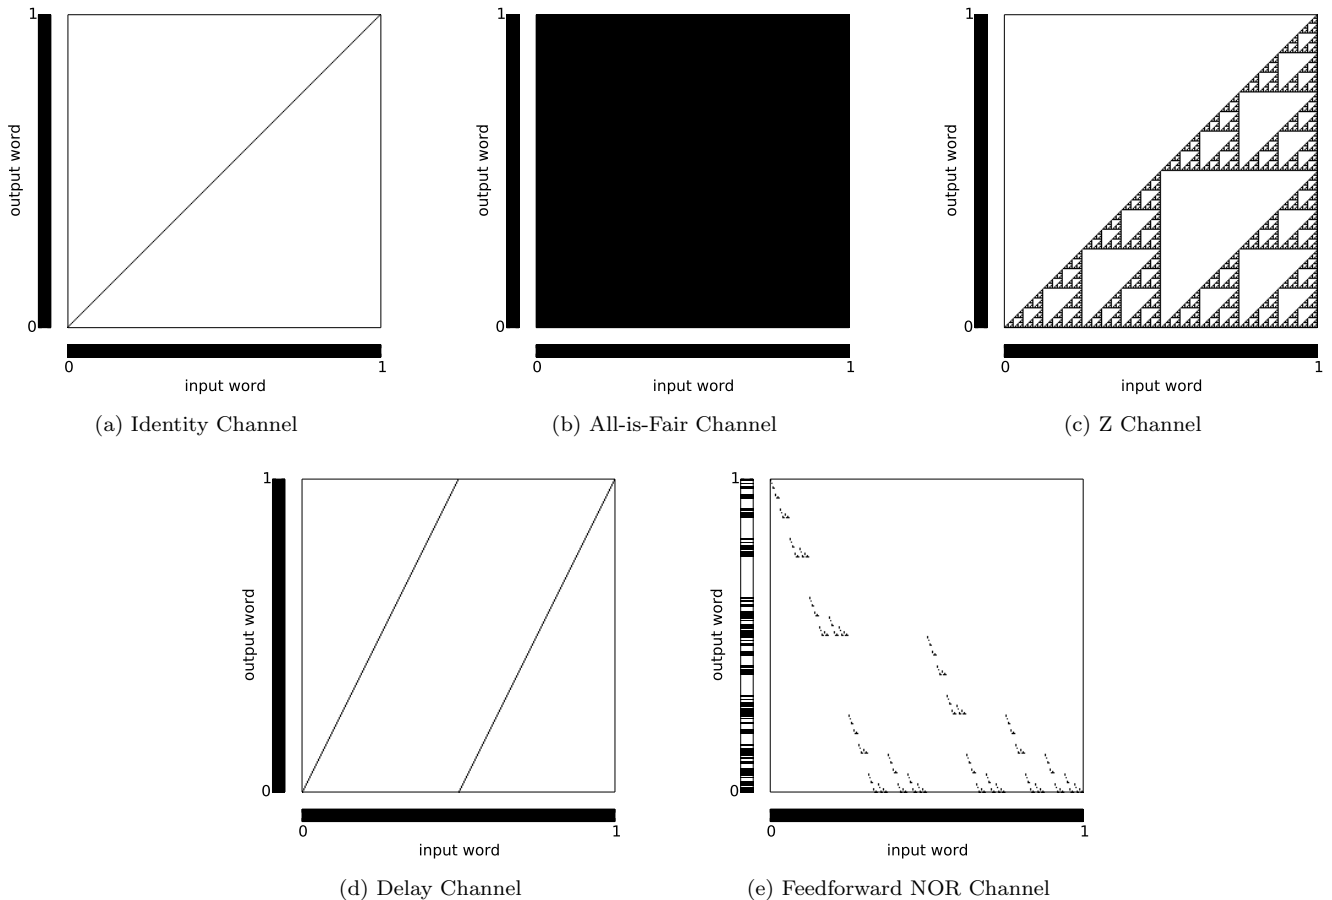


FIG. 3. Wordmaps for simple example channels display the map from finite input words to finite output words: Allowed input-output pairs  $(x_{0:s}, y_{0:s})$  correspond to a dot placed in the unit square at  $(0.x_0x_1 \cdots x_7, 0.y_0y_1 \cdots y_7)$ . The possible observed input (output) words are displayed below (left) of the horizontal (vertical) axis.

### A. Memorylessness

The *Memoryless Binary Channel's* (MBC's) current output depends only on its current input; the analog of an IID process in that its behavior at time  $t$  is independent of that at other times. The MBC includes as special cases the *Binary Symmetric Channel* (BSC) and the *Z Channel* [12]. We can summarize the MBC's present behavior with a simplified recurrence relation of the form  $Y_t \sim r(X_t)$  and its conditional probabilities factor as follows:

$$\begin{aligned} \mathbb{P}(Y_{t:t+L} | \overleftarrow{X}) \\ = \mathbb{P}(Y_t | X_t) \mathbb{P}(Y_{t+1} | X_{t+1}) \cdots \mathbb{P}(Y_{t+L-1} | X_{t+L-1}) . \end{aligned}$$

The first three wordmaps of Fig. 3 illustrate the behavior of memoryless channels. We see that the *Identity Channel* (Fig. 3(a)) always maps a word to itself. Whereas, the *All-is-Fair Channel* (Fig. 3(b)) maps each input word uniformly to every output word. We immediately see that deterministic channels have wordmaps with a single filled pixel per plot column.

The *Z Channel* wordmap is shown in Fig. 3(c). It

transmits all 0s with no noise, but adds noise to all 1s transmitted. The wordmap shows the maximal noise case, where all 1s are replaced with the output of a fair coin. We see that even memoryless channels have non-trivial word mappings. In this case, the latter forms a self-similar Sierpinski right triangle [30] in the unit square.

### B. Finite Feedforward

Now, consider channels whose behavior depends only on a finite input history. Their behavior on input is analogous to those of order- $R$  Markov chains. They can also be thought of as stochastic, anticipation-free *sliding block codes* [25] with finite memory or as generalized *finite impulse response filters* [29]. These channels' present behavior can be summarized with a recurrence relation of the form  $Y_t \sim r(X_{t-M:t}, X_t)$ , where  $M$  is a finite input history length such that  $M > 0$ .

The *Delay Channel* simply stores its input at time  $t-1$  and outputs it at time  $t$ . Its wordmap is shown in



Fig. 3(d). Those familiar with one-dimensional iterated maps of the interval will recognize the word mapping as the *shift map*, capturing the fact that delayed output corresponds to a binary shift applied to the input word. Note that when viewed as a function on the space of processes, the Delay Channel acts as the identity.

The *Feedforward NOR Channel*'s output at time  $t$  results from performing a logical NOR ( $\downarrow$ ) on its inputs at times  $t$  and  $t - 1$ :

$$y_t = x_{t-1} \downarrow x_t .$$

The wordmap for the Feedforward NOR Channel is shown in Fig. 3(e). There, we see that although the channel is deterministic, the wordmap's self-similarity makes it difficult to see that the mapping is a function—that there is, in fact, a single output word for each input.

The additional complexity of the remaining examples does not lead to new types of apparent graphical structure beyond that seen in the existing wordmaps. So, wordmaps will be omitted for now. With additional theory developed, we return to illustrate these channels, but using a more advanced form of wordmap.

### C. Infinite Feedforward

Generally, channels depend on infinitely long input histories. This behavior is analogous to the long-range dependence seen in typical hidden Markov models [31, 32] or in *strictly sofic* subshifts [25]. Channels with dependence upon infinitely long input histories *alone* can also be interpreted as generalized *infinite impulse response filters* [29]. The present behavior of such channels can be summarized by a recurrence relation of the form  $Y_t \sim r(\bar{X}_t, X_t)$ .

The *Odd NOT Channel* stores the parity (even or odd) of the number of ones observed in its input since the last zero observed; much like the Even Process. If the parity is currently even, it behaves as the Identity Channel. If the parity is odd, it outputs the bitwise NOT (bit flip) of its input. Since the channel's behavior depends on the parity of its input, it cannot be characterized by finite input histories alone.

A channel can depend, however, on past *outputs* as well as inputs. Such *feedback* can allow one to replace the infinite-history recurrence relation with one that includes only a finite history of inputs and outputs. Consider again the Odd NOT Channel described above. Note that its behavior is determined entirely by the current input value, as well as the parity of the number of ones observed on input. In fact, the parity at time  $t$  can be summarized by the input and output at time  $t - 1$ . If  $x_{t-1} = 0$ , the parity will always be even. If  $x_{t-1} = 1$ , and  $y_{t-1} = 0$  we know that the parity *was* odd, since the bit was flipped, but since a 1 was just observed on input, the parity is now even. Finally, if  $x_{t-1} = 1$  and  $y_{t-1} = 1$ , we know that the parity was previously even, and the newly observed

1 makes the current parity *odd*. Summarizing, we have that:

$$\begin{aligned} (x, y)_{t-1} = (0, 0) &\Leftrightarrow \text{Even input history parity,} \\ (x, y)_{t-1} = (0, 1) &\Leftrightarrow \text{Even input history parity,} \\ (x, y)_{t-1} = (1, 0) &\Leftrightarrow \text{Even input history parity, and} \\ (x, y)_{t-1} = (1, 1) &\Leftrightarrow \text{Odd input history parity .} \end{aligned} \quad (7)$$

By allowing feedback, we can therefore summarize the behavior of the Odd NOT Channel with a recurrence relation that depends only upon finite history (of length  $M = 1$ ):  $Y_t \sim r(X_{t-1}, Y_{t-1}, X_t)$ .

An interesting observation is that the Odd NOT channel maps the Even Process to a bit-flipped Golden Mean Process. However, a single process-to-process mapping does *not* uniquely define a channel. There are an infinite number of channels, in fact, that map the Even Process to the bit-flipped Golden Mean Process.

### D. Finite Feedback

As seen in the previous section, allowing for even a finite amount of output feedback can lead to substantial simplifications in the description of a channel. Let's consider channels that depend on a finite output history on their own terms.

The most trivial case would be channels that depend solely on output histories, with no dependence on inputs. Since there is effectively no input, these channels reduce to the output-only stochastic processes (generators) discussed earlier. Consider, for example, the *All is Golden Channel* that outputs the Golden Mean Process, regardless of what input it receives.

A less trivial example is the *Feedback NOR Channel*, similar to the Feedforward NOR Channel, except the output at time  $t$  is the logical NOR of its current input and *previous output*:

$$y_t = x_t \downarrow y_{t-1} .$$

This channel's behavior is clear in this feedback form. It might be desirable, however, to find a purely feedforward presentation for the channel. The recurrence relation for the Feedback NOR Channel can be solved recursively to give a feedforward presentation that is defined for almost every input history. We recurse in the following way:

$$\begin{aligned} y_t &= x_t \downarrow y_{t-1} \\ &= x_t \downarrow (x_{t-1} \downarrow y_{t-2}) \\ &= x_t \downarrow (x_{t-1} \downarrow (x_{t-2} \downarrow y_{t-3})) \\ &= \dots , \end{aligned}$$

and so on, until reaching sufficiently far into the input past that a 1 is observed. When this happens, the recursion terminates as the output of the NOR function is always 0 when either argument is 1. We can therefore con-

struct a purely feedforward recurrence relation, but the input histories can be arbitrarily long—corresponding to arbitrarily long input histories consisting entirely of 0s. The resulting feedforward recurrence relation is defined for all histories, except for the infinite history of all 0s.

Note that such ill-defined behavior for certain infinite histories is typical in systems that have infinite memory lengths and is not a problem specific to channels. One can be careful to explicitly *define* behavior for such cases, but this is beyond the scope of our current work, and these pathological histories typically occur with zero probability.

In contrast, consider replacing the logical NOR in the Feedback NOR Channel with an exclusive OR (XOR or  $\oplus$ ), thus giving the *Feedback XOR Channel*:

$$y_t = x_t \oplus y_{t-1}. \quad (8)$$

In this case, solving for a pure-feedforward relation fails since the output of a logical XOR is never determined by a single argument. This channel illustrates the fact that a presentation which includes feedback cannot always be reduced to a pure-feedforward presentation.

### E. Infinite Feedback

Just as we can define channels whose behavior depends upon infinite input histories, we can define channels whose behavior depends upon infinite *output* histories. In fact, we have already studied a channel that can be represented this way. Consider a channel that stores the parity (even or odd) of the number of ones observed in its *output* since the last zero observed. If this parity is currently even, it behaves as the Identity Channel. If the parity is odd, it outputs the bitwise NOT (bit flip) of its input. This appears to be very similar to the Odd NOT Channel defined above, but with a dependence on infinite output histories and the present input, rather than infinite input histories. In fact, this is simply a different *presentation* of the Odd NOT Channel.

It suffices to show that the feedback presentation of the Odd NOT Channel can be reduced to the same finite history presentation as with its feedforward presentation. Proceeding as before, we observe that if  $y_{t-1} = 0$ , the output parity is always even. If  $x_{t-1} = 0$  and  $y_{t-1} = 1$ , the previous parity was odd (the bit was flipped), but the 1 observed on output makes the output parity even. Finally, if  $x_{t-1} = 1$  and  $y_{t-1} = 1$ , the output parity was even, and the 1 observed makes the output parity odd. Summarizing, we obtained the same presentation as the feedforward presentation specified by Eq. (7).

### F. Infinite Feedforward-Feedback

We just examined an example channel whose presentation depends upon infinite histories when only feed-

forward or feedback is allowed, but only a finite history when both feedforward and feedback are allowed. The following example shows that finding a finite history presentation is not always possible. Channels of this form are perhaps the most natural channel generalization of infinite Markov order (strictly sofic) processes.

The *Odd Random Channel* stores the parity of its input history just as the Odd NOT Channel does, and it again behaves as the identity when the parity is currently even. When the parity is odd, the channel outputs a 0 or 1 with equal probability. Like the Odd NOT Channel, this has an infinite feedback presentation that stores the channel’s output history parity. The channel does *not* have any finite history presentation, however. If one attempts to construct a finite presentation via the recursion unrolling procedure used for the Odd NOT Channel, it is simple to obtain the following relationships:

$$\begin{aligned} (x, y)_{t-1} = (0, 0) &\Leftrightarrow \text{Even input history parity,} \\ (x, y)_{t-1} = (0, 1) &\Leftrightarrow \text{Even input history parity, and} \\ (x, y)_{t-1} = (1, 0) &\Leftrightarrow \text{Even input history parity.} \end{aligned}$$

The problem arises from the fact that when  $x_{t-1} = 1$  and  $y_{t-1} = 1$ , the input history parity is uncertain. The channel could have been operating as the identity (even parity) or giving random output (odd parity). Looking at progressively longer histories can resolve this uncertainty, but only once a 0 has been observed (on either input, output, or both). This ambiguity requires that we specify arbitrarily long joint histories to determine the behavior of the channel in general. It is therefore not possible to construct any finite-history presentation of the Odd Random Channel.

### G. Irreducible Feedforward-Feedback

As a final example, consider a channel that has a finite presentation when both feedforward and feedback are allowed, but has *no* pure-feedforward or pure-feedback presentation. The *Period-2 Identity NOT Channel* alternates between the identity and bit flipped identity at each time step. This channel’s present behavior is completely determined by whether the previous input and output bits,  $x_{t-1}$  and  $y_{t-1}$ , match. When the bits match, the channel was in its “identity” state and is therefore now in its “NOT” state. The opposite is clearly true when the bits do not match. The channel therefore has a recurrence relation of the form  $Y_t \sim r(X_{t-1}, Y_{t-1}, X_t)$ .

Since the behavior does not depend on the particular values of the input or output, but whether or not they *match*, there is no way to construct a pure-feedforward or pure-feedback presentation of the channel. This channel therefore illustrates the notion of *irreducible* output (or input) memory—dependence upon past output (input) that cannot be eliminated even by including dependence upon infinite past outputs (inputs).

## H. Causal Channel Markov Order Hierarchy

It turns out that the set of examples above outlines a classification scheme for causal channels in terms of their Markov orders [33] that we now make explicit. Just as Markov order plays a key role in understanding the organization of processes, it is similarly helpful for channels. Channel Markov orders are the history lengths required to *completely* specify a causal channel's behavior, given certain constraints on knowledge of other histories.

### Definition 5.

1. The pure feedforward Markov order  $R_{\text{pff}}$  is the smallest  $M$  such that  $\forall L$   
 $\mathbb{P}(Y_{0:L} | (\overline{X}, \overline{Y})_0, X_{0:L}) = \mathbb{P}(Y_{0:L} | X_{-M:0}, X_{0:L})$ .
2. The pure feedback Markov order  $R_{\text{pfb}}$  is the smallest  $M$  such that  $\forall L$   
 $\mathbb{P}(Y_{0:L} | (\overline{X}, \overline{Y})_0, X_{0:L}) = \mathbb{P}(Y_{0:L} | Y_{-M:0}, X_{0:L})$ .
3. The channel Markov order  $R$  is the smallest  $M$  such that  $\forall L$   
 $\mathbb{P}(Y_{0:L} | (\overline{X}, \overline{Y})_0, X_{0:L}) = \mathbb{P}(Y_{0:L} | (X, Y)_{-M:0}, X_{0:L})$ .
4. The irreducible feedforward Markov order  $R_{\text{iff}}$  is the smallest  $M$  such that  $\forall L$   
 $\mathbb{P}(Y_{0:L} | (\overline{X}, \overline{Y})_0, X_{0:L}) = \mathbb{P}(Y_{0:L} | X_{-M:0}, \overleftarrow{Y}_0, X_{0:L})$ .
5. The irreducible feedback Markov order  $R_{\text{ifb}}$  is the smallest  $M$  such that  $\forall L$   
 $\mathbb{P}(Y_{0:L} | (\overline{X}, \overline{Y})_0, X_{0:L}) = \mathbb{P}(Y_{0:L} | \overleftarrow{X}_0, Y_{-M:0}, X_{0:L})$ .

For example, we showed that the Odd NOT Channel's presentation requires an infinite history when only feedforward or feedback is allowed. And so, it has  $R_{\text{pff}} = R_{\text{pfb}} = \infty$ . However, it only requires finite history when both are allowed:  $R = 1$ . If we have full knowledge of the output past, we still need one symbol of input history in order to characterize the channel, so  $R_{\text{iff}} = 1$ . Similarly, we have  $R_{\text{ifb}} = 1$ .

Note that the irreducible feedforward Markov order  $R_{\text{iff}}$  will only be nonzero if the pure feedback order  $R_{\text{pfb}}$  is undefined. Similarly, the irreducible feedback Markov order  $R_{\text{ifb}}$  will only be nonzero when the pure feedforward order  $R_{\text{pff}}$  is undefined. In words, if a channel has irreducible feedback (feedforward), the channel has no pure feedforward (feedback) presentation. Moreover, the channel Markov order  $R$  bounds the pure Markov orders from below and the smallest of the two irreducible Markov orders bounds the channel Markov order from below:

$$\min(R_{\text{iff}}, R_{\text{ifb}}) \leq R \leq \min(R_{\text{pff}}, R_{\text{pfb}}) .$$

In a sequel, we address input and output memory using information-theoretic quantities. There, we characterize different *amounts* of input and output memory and how they relate, whereas here in discussing Markov orders we considered *lengths* of input and output sequences. We also focused on causal channels, which allowed us to

restrict our discussion to present behavior and memory lengths. In the case of anticipatory channels, we must explicitly consider *future* behavior, as well as *anticipation* lengths. However, this is best left to another venue, so that we do not deviate too far from the path to our goal.

## I. Causal-State Channels

There is a natural and quite useful channel embedded in any  $\epsilon$ -machine presentation of a stationary process—the *causal-state channel*—that identifies structure embedded in a process via the  $\epsilon$ -machine's causal states.

Consider a process and its  $\epsilon$ -machine  $M$ . Previously, we described  $M$  as a generator of the process. An  $\epsilon$ -machine, however, is also a *recognizer* of its process's sequences. Briefly,  $M$  reads a sequence and follows the series of transitions determined by the symbols it encounters. The output of the causal-state channel is then the sequence of causal states. The operation of this channel is what we call *causal-state filtering*. Notably, the induced mapping is a function due to the  $\epsilon$ -machine's unifilarity.

In this way, the channel filters observed sequences, returning step-by-step associated causal states. Given an  $\epsilon$ -machine, the causal-state filter has the same topology as the  $\epsilon$ -machine, but input symbols match the  $\epsilon$ -machine transition symbols and output symbols are the state to which the transition goes.

The recurrence relation for causal-state filtering is:

$$S_t \sim r(X_{t-1}, S_{t-1}) .$$

As just noted,  $r(\cdot)$  is a (nonprobabilistic) function, determined by the  $\epsilon$ -map:

$$\begin{aligned} S_t &= \epsilon(\overleftarrow{X}_t) \\ &= \epsilon(X_{t-1}, \overleftarrow{X}_{t-1}) \\ &= \epsilon(X_{t-1}, S_{t-1}) . \end{aligned}$$

Thus, the pure feedback order is  $R_{\text{pfb}} = 1$ , as is the channel Markov order  $R = 1$ . The pure feedforward order  $R_{\text{pff}}$ , however, is the original process's Markov order. For methods to determine the latter see Ref. [33].

In this way, an  $\epsilon$ -machine can be used to detect the hidden structures captured by the causal states. For example, causal-state filtering has been used to great effect in detecting emergent domains, particles, and particle interactions in spatially extended dynamical systems [34], single-molecule conformational states [35], and polycrystalline and fault structures in complex materials [36].

## V. GLOBAL $\epsilon$ -MACHINE

Before introducing channel presentations, we first frame the question in the global setting of the  $\epsilon$ -machine of joint (input-output) processes. This, then, grounds

the  $\epsilon$ -transducer in terms of an  $\epsilon$ -machine.

Given a stationary process  $\overleftarrow{X}$  and a channel  $\overleftarrow{Y}|\overleftarrow{X}$  with output process  $\overleftarrow{Y}$ , form a joint random variable  $Z_t = (X, Y)_t$  over input-output symbol pairs with alphabet  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ . For example, if  $\mathcal{X} = \{a, b\}$  and  $\mathcal{Y} = \{c, d\}$ , then  $\mathcal{X} \times \mathcal{Y} = \{ac, ad, bc, bd\}$ .

**Definition 6.** The process  $\overleftarrow{Z}$  over  $\mathcal{Z}$  defines the channel's I/O process.

**Definition 7.** A stationary channel's global  $\epsilon$ -machine is the  $\epsilon$ -machine of its I/O process.

In this setting, the next section asks for a particular decomposition of the global  $\epsilon$ -machine, when one has selected a portion of  $Z_t$  as “input” and another as “output”. The input process  $\overleftarrow{X}$  is then described by the marginal distribution of the joint process that projects onto “input” sequences. The same also holds for the output process  $\overleftarrow{Y}$ . In this way, the following results not only provide an analysis for specified input and output processes, but also an analysis of possible input-to-output mappings embedded in any process or its  $\epsilon$ -machine. Leveraging this observation and anticipating the sequels, we also note here that the global  $\epsilon$ -machine also provides the proper setting for posing questions about information storage and flow *within* any given process.

## VI. $\epsilon$ -TRANSDUCER

Computational mechanics' fundamental assumption—only prediction matters—applies as well to channels as to processes. In the case of channels, though, we wish to predict the channel's future *output* given the channel's past inputs *and* outputs *and* the channel's future *input*. This leads to a new causal equivalence relation  $\sim_\epsilon$  over joint pasts  $\overleftarrow{z} = \overleftarrow{(x, y)}$ :

$$\begin{aligned} \overleftarrow{(x, y)} \sim_\epsilon \overleftarrow{(x, y)'} &\iff \\ \mathbb{P}(\overleftarrow{Y}|\overleftarrow{X}, \overleftarrow{(X, Y)} = \overleftarrow{(x, y)}) &= \mathbb{P}(\overleftarrow{Y}|\overleftarrow{X}, \overleftarrow{(X, Y)} = \overleftarrow{(x, y)'}) . \end{aligned} \quad (9)$$

Compare Eq. (4) applied to the I/O process. The equivalence classes of  $\sim_\epsilon$  partition the set  $\overleftarrow{\mathcal{Z}} = \overleftarrow{(\mathcal{X}, \mathcal{Y})}$  of all input-output pasts. These classes are the channel's *causal states*, denoted  $\mathcal{S}$ . The  $\epsilon$ -map is a function  $\epsilon : \overleftarrow{(\mathcal{X}, \mathcal{Y})} \rightarrow \mathcal{S}$  that maps each joint past to its corresponding channel causal state or, equivalently, to the set of joint pasts to which it is causally equivalent:

$$\epsilon(\overleftarrow{(x, y)}) = \sigma_i = \{ \overleftarrow{(x, y)'} : \overleftarrow{(x, y)} \sim_\epsilon \overleftarrow{(x, y)'} \} .$$

The dynamic over causal states is again inherited from the implicit dynamic over joint pasts via the  $\epsilon$ -map, resulting from appending the joint symbol  $z_t = (x, y)_t$ , as shown in Fig. 4. Since state transitions now depend upon the current input symbol, we specify the dynamic by an

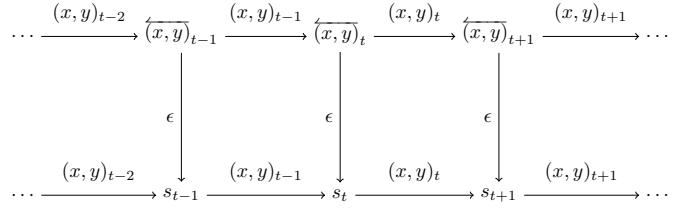


FIG. 4.  $\epsilon$ -Transducer dynamic induced by the causal states  $s_t = \sigma_i \in \mathcal{S}$  from a channel's joint pasts via the  $\epsilon$ -map:  $s_t = \epsilon(\dots(x, y)_{t-2}(x, y)_{t-1}) \xrightarrow{(x, y)_t} s_{t+1} = \epsilon(\dots(x, y)_{t-2}(x, y)_{t-1}(x, y)_t)$ .

indexed set of *conditional-symbol transition matrices*:

$$\mathcal{T} \equiv \left\{ T^{(y|x)} \right\}_{x \in \mathcal{X}, y \in \mathcal{Y}} ,$$

where  $T^{(y|x)}$  has elements:

$$T_{ij}^{(y|x)} = \mathbb{P}(S_1 = \sigma_j, Y_0 = y | S_0 = \sigma_i, X_0 = x) .$$

While the causal states for a stationary process have a *unique* stationary distribution, each stationary input to a channel can drive its causal states into a *different* stationary state distribution. The  $\epsilon$ -map from joint histories to the channel's causal states can also be seen as a function that maps a *distribution* over joint histories to a distribution over the channel's causal states. Since each input history specifies a particular distribution over *output* histories (via the channel's conditional word probabilities), it follows that a *distribution* over input histories also specifies a distribution over output histories. When this input history distribution is specified via a particular input process, we obtain a unique distribution over causal states via its  $\epsilon(\cdot)$  function. We write this input-dependent state distribution  $\pi_X$  as:

$$\begin{aligned} \pi_X(i) &= \mathbb{P}_X(S_0 = \sigma_i) \\ &= \mathbb{P}_X(\epsilon(\overleftarrow{(X, Y)}) = \sigma_i) , \end{aligned}$$

where the subscript  $X$  indicates that the input process has a specific, known distribution. The distribution over joint histories is stationary by assumption here and, since the  $\epsilon$ -map is time independent,  $\pi_X$  is stationary. We therefore refer to  $\pi_X$  as the (input-dependent) *stationary distribution*.

When both the input process and channel are stationary and ergodic, we can calculate this stationary distribution from the input and channel's causal-state transition matrices using a generalization of the algorithm found in Ref. [17]. We save an in-depth discussion of this algorithm for a sequel.

**Definition 8.** The tuple  $(\mathcal{X}, \mathcal{Y}, \mathcal{S}, \mathcal{T})$ —consisting of the channel's input and output alphabets, causal states, and conditional-symbol transition probabilities, respectively—is the channel's  $\epsilon$ -transducer.

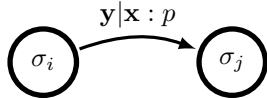


FIG. 5. Generic transition from  $\epsilon$ -transducer causal state  $\sigma_i$  to state  $\sigma_j$  while accepting input symbol  $x$  and emitting symbol  $y$  with probability  $p = T_{ij}^{(y|x)}$ .

Note that in the causal equivalence relation for channels, we condition on the input future  $\overline{X}$ , as a channel is defined by its output behavior *given* input. Requiring causal equivalence for the output future alone (or the joint future for that matter) requires knowledge of a particular *input process* as well. In particular, if we *do* have knowledge of the input process we can extend the standard causal equivalence relation to an equivalence relation involving joint pasts and joint future morphs, giving us the *global* (joint)  $\epsilon$ -machine of the preceding section.

As with  $\epsilon$ -machines, it is useful to consider channels whose  $\epsilon$ -transducers have a finite (or countable) number of causal states. This restriction again allows us to represent an  $\epsilon$ -transducer as a labeled-directed graph. Since transitions between causal states now depend on inputs as well as outputs, we represent a transition from state  $\sigma_i$  to state  $\sigma_j$  while accepting input symbol  $x$  and emitting symbol  $y$  as a directed edge from node  $\sigma_i$  to node  $\sigma_j$  with edge label  $\mathbf{y|x} : p$ , where  $p = T_{ij}^{(y|x)}$  is the edge transition probability. This is illustrated in Fig. 5.

## VII. STRUCTURAL COMPLEXITY

To monitor the degree of structuredness in an  $\epsilon$ -transducer, we use the Shannon information captured by its causal states, paralleling the definition of  $\epsilon$ -machine statistical complexity. The stationary distribution over  $\epsilon$ -transducer states allows one to define an  $\epsilon$ -transducer's *input-dependent* statistical complexity:

$$C_X = H[\pi_X] .$$

(Note that  $X$  replaces the previously subscripted measure  $\mu$  in  $\epsilon$ -machine statistical complexity to specify the now-relevant measure.) While quantifying structural complexity,  $C_X$ 's dependence on input requires a new interpretation. Some processes drive a transducer into simple, compressible behavior, while others will lead to complex behavior. Figure 6 illustrates this.

Input dependence can be removed by following the standard definition of channel capacity [12], giving a single number characterizing an  $\epsilon$ -transducer. We take the supremum of the statistical complexity over input processes. This gives an upper bound on  $\epsilon$ -transducer complexity—the *channel complexity*:

$$\overline{C}_\mu = \sup_X C_X ,$$

where the maximizing input measure  $\mu$  is implicitly defined. Note that not all transducers can be driven to a uniform distribution over states. Thus, recalling that uniform distributions maximize Shannon entropy, in general  $\overline{C}_\mu \leq C_0 \equiv \log_2 |\mathcal{S}|$ —the *topological state complexity*.

## VIII. REPRODUCING A CHANNEL

To establish that the  $\epsilon$ -transducer is an exact presentation of the causal channel it models, we must show that it reproduces the channel's conditional word probabilities. We first establish some needed notation.

Recall that the  $\epsilon$ -map takes a distribution over joint histories to a distribution over the  $\epsilon$ -transducer's causal states and that a particular input history defines a distribution over a channel's output history. It follows that a particular input history defines a distribution over joint histories and, therefore, also defines a distribution over an  $\epsilon$ -transducer's causal states via the  $\epsilon$ -map. We call this distribution  $\tau$ :

$$\begin{aligned} \tau(i) &= \mathbb{P}_X(S_0 = \sigma_i | \overline{X}_0 = \overline{x}) \\ &= \mathbb{P}_X(\epsilon(\overline{(X, Y)}_0) = \sigma_i | \overline{X}_0 = \overline{x}) . \end{aligned}$$

Note that while the  $\epsilon$ -map takes a particular *joint* history to a unique state, the distribution over states induced by a particular *input* history need not be concentrated on a single state.

When the input process is known and stationary, the  $\epsilon$ -transducer's stationary distribution  $\pi_X$  can be determined. This provides a starting distribution for the  $\epsilon$ -transducer, which is updated by the  $\epsilon$ -transducer's symbol transition matrices as each input symbol is observed and each output symbol is generated. We can therefore calculate the final state distribution  $\mathbb{P}(S_L = \sigma_m | x_{0:L}, S_0 \sim \pi_X)$  that results from starting states in distribution  $\pi_X$  and observing finite input word  $x_{0:L}$ :

$$\begin{aligned} \mathbb{P}(S_L = \sigma_m | x_{0:L}, S_0 \sim \pi_X) \\ = \sum_{y_{0:L}} \sum_{i, j, k, \dots, l} \pi_X(i) T_{ij}^{(y_0|x_0)} T_{jk}^{(y_1|x_1)} \dots T_{lm}^{(y_{L-1}|x_{L-1})} . \end{aligned}$$

We then obtain  $\tau$  from the  $\epsilon$ -transducer by shifting this distribution by  $L$  and taking the limit as  $L \rightarrow \infty$ :

$$\begin{aligned} \tau(i) &= \mathbb{P}_X(S_0 = \sigma_i | \overline{X}_0 = \overline{x}) \\ &= \lim_{L \rightarrow \infty} \mathbb{P}(S_0 = \sigma_i | x_{-L:0}, S_{-L} \sim \pi_X) . \end{aligned}$$

We can now establish that the  $\epsilon$ -transducer is an exact presentation of the causal channel that it models.

**Proposition 2** (Presentation). *A causal channel's  $\epsilon$ -transducer exactly (and only) reproduces the channel's conditional word probabilities  $\mathbb{P}(Y_{0:L} | \overline{X}_L)$ .*

*Proof.* Recall that a *causal* channel's output words do

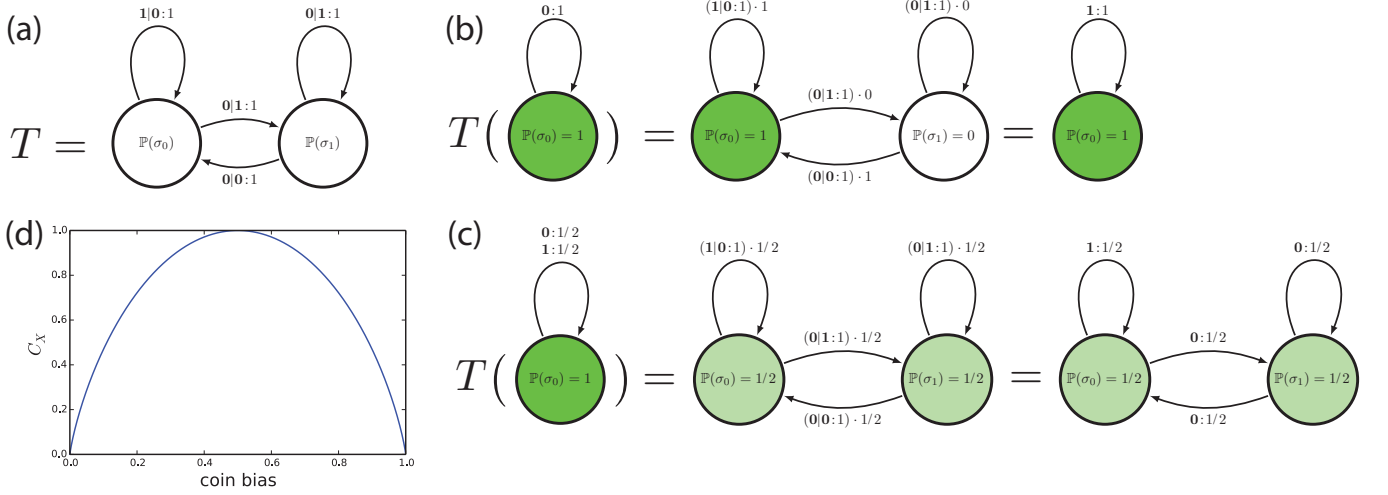


FIG. 6. A channel’s statistical complexity  $C_X$  depends on the input process  $X$ : (a) An example transducer  $T$ —the Feedforward NOR Channel described later—driven by two example inputs, showing the intermediate stage calculation  $T(X)$  and the  $\epsilon$ -machine of the output process  $Y = T(X)$ . (b) When driven by an input process consisting of all 0s, the statistical complexity vanishes:  $C_X(T) = \log_2(\mathbb{P}(\sigma_0) = 1) = 0$  bits. (c) When driven by the Fair Coin Process, the distribution over causal states is uniform and there is a positive statistical complexity:  $C_X(T) = H(1/2) = 1$  bit. (d) Statistical complexity  $C_X$  as a function of the bias of an input Biased Coin Process.

not depend on any inputs occurring *after* the output word. Therefore, the  $\epsilon$ -transducer must reproduce all of a channel’s conditional word probabilities of the form  $\mathbb{P}(y_{0:L}|\overleftarrow{x}_L)$ . As discussed above, an input history induces a distribution  $\tau$  over the  $\epsilon$ -transducer’s causal states. So, we calculate the word probabilities directly via repeated application of the  $\epsilon$ -transducer’s symbol transition matrices:

$$\begin{aligned} \mathbb{P}(y_{0:L}|\overleftarrow{x}_L) &= \sum_{i,j,k,\dots,l,m} \tau(i) T_{ij}^{(y_0|x_0)} T_{jk}^{(y_1|x_1)} \dots T_{lm}^{(y_{L-1}|x_{L-1})}. \end{aligned} \quad (10)$$

□

In fact, we can transduce a *finite* input word even when the channel’s behavior depends upon arbitrarily long input histories. By simply starting the  $\epsilon$ -transducer in its stationary distribution  $\pi_X$ , we have:

$$\begin{aligned} \mathbb{P}(y_{0:L}|x_{0:L}) &= \sum_{i,j,k,\dots,l,m} \pi_X(i) T_{ij}^{(y_0|x_0)} T_{jk}^{(y_1|x_1)} \dots T_{lm}^{(y_{L-1}|x_{L-1})}. \end{aligned}$$

In either case, we can multiply these conditional word probabilities by the input’s word probabilities to obtain joint word probabilities:

$$\mathbb{P}((x, y)_{0:L}) = \mathbb{P}(y_{0:L}|x_{0:L})\mathbb{P}(x_{0:L}).$$

Summing over input words then gives output word prob-

abilities:

$$\mathbb{P}(y_{0:L}) = \sum_{x_{0:L} \in \mathcal{X}^L} \mathbb{P}((x, y)_{0:L}).$$

We can also start the  $\epsilon$ -transducer with other, arbitrary state distributions when certain initial behavior is desired or if the current internal configuration of the  $\epsilon$ -transducer is known. This can be very useful in practice, but the resulting generated behavior is no longer guaranteed to be stationary or to match the original channel’s behavior. One use of arbitrary state distributions is *real-time* transduction of symbols, where a state distribution  $\nu$  is repeatedly updated each time-step after a single input symbol  $x_t$  is transduced to an output symbol  $y_t$ :

$$\mathbb{P}(S_{t+1} = \sigma_i | (x, y)_t, S_t \sim \nu) = \sum_{i,j} \nu(j) T_{ij}^{(y_t|x_t)}.$$

## IX. OPTIMALITY

We now establish that the  $\epsilon$ -transducer is a channel’s unique, maximally predictive, minimal statistical complexity unifilar presentation. Other properties, analogous to those of the  $\epsilon$ -machine, are also developed. Several proofs parallel those in Refs. [18, 19], but are extended from  $\epsilon$ -machines to the  $\epsilon$ -transducer. For this initial development, we also adopt a caveat from there concerning the use of infinite pasts and futures. For example, the semi-infinite pasts’ entropy  $H[\overleftarrow{Y}]$  is typically infinite. And so, to properly use such quantities, one first intro-

duces finite-length chains (e.g.,  $H[Y_{0:L}]$ ) and at the end of an argument one takes infinite-length limits, as appropriate. Here, as previously, we do not include these extra steps, unless there is subtlety that requires attention using finite-length chains.

**Proposition 3** (Causal States Proxy the Past). *When conditioned on causal states, the future output given input is independent of past input and past output:*

$$\mathbb{P}(\vec{Y}_0 | \vec{X}_0, \overleftarrow{(X, Y)}_0, S_0) = \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) .$$

*Proof.* By construction, the causal states have the same future morphs as their corresponding pasts:

$$\mathbb{P}(\vec{Y}_0 | \vec{X}_0, \overleftarrow{(X, Y)}_0) = \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) .$$

Since the causal states are a function of the past— $S_0 = \epsilon(\overleftarrow{(X, Y)}_0)$ —we also have that:

$$\mathbb{P}(\vec{Y}_0 | \vec{X}_0, \overleftarrow{(X, Y)}_0, S_0) = \mathbb{P}(\vec{Y}_0 | \vec{X}_0, \overleftarrow{(X, Y)}_0) .$$

Combining these two equalities gives the result.  $\square$

In other words, when predicting a channel’s future behavior from its past behavior, it suffices to use the causal states instead.

**Proposition 4** (Causal Shielding). *Past output  $\vec{Y}_0$  and future output  $\vec{Y}_1$  given future input  $\vec{X}_0$  are independent given the current causal state  $S_0$ :*

$$\mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) = \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) \mathbb{P}(\vec{Y}_1 | \vec{X}_0, S_0) .$$

*Proof.* We directly calculate:

$$\begin{aligned} \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) &= \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) \mathbb{P}(\vec{Y}_1 | \vec{X}_0, \vec{Y}, S_0) \\ &= \mathbb{P}(\vec{Y}_0 | \vec{X}_0, S_0) \mathbb{P}(\vec{Y}_1 | \vec{X}_0, S_0) . \end{aligned}$$

Where the second equality follows from applying Prop. 3 to the second factor.  $\square$

In the following, depending on use, we refer to either of the previous propositions as *causal shielding*.

**Proposition 5** (Joint Unifilarity). *The current causal state  $S_0$  and current input-output symbol pair  $(X, Y)_0$  uniquely determine the next causal state. In this case:*

$$H[S_1 | (X, Y)_0, S_0] = 0 .$$

*Proof.* If two pasts are causally equivalent, then either (i) appending a new symbol pair  $(x, y)$  to both pasts results in two *new* pasts that are also causally equivalent or (ii) such a symbol pair is never observed when in  $S_0$ . We must show that the two new pasts have the same future morph:

$$\begin{aligned} \overleftarrow{(x, y)} \sim_\epsilon \overleftarrow{(x, y)}' \\ \implies \mathbb{P}(\vec{Y}_1 | \vec{X}_1, \overleftarrow{(x, y)}(a, b)) = \mathbb{P}(\vec{Y}_1 | \vec{X}_1, \overleftarrow{(x, y)}'(a, b)) , \end{aligned}$$

where we have  $a \in \mathcal{X}$  and  $b \in \mathcal{Y}$  and  $\overleftarrow{(x, y)}(a, b) = (\overleftarrow{x}a, \overleftarrow{y}b)$ , and the futures  $\vec{Y}_1$  and  $\vec{X}_1$  denote those immediately following the associated conditioning pasts  $\overleftarrow{(x, y)}(a, b)$  and  $\overleftarrow{(x, y)}'(a, b)$ , respectively. Or, we must show that the input-output pair  $(x, y)$  is forbidden.

First, let  $\overleftarrow{(x, y)} \sim_\epsilon \overleftarrow{(x, y)}'$ . Since causal equivalence applies for *any* joint future, it applies to the particular future beginning with symbol pair  $(a, b)$ :

$$\mathbb{P}(b\vec{Y}_1 | a\vec{X}_1, \overleftarrow{(x, y)}) = \mathbb{P}(b\vec{Y}_1 | a\vec{X}_1, \overleftarrow{(x, y)}') .$$

Factoring:

$$\mathbb{P}(b\vec{Y}_1 | \cdot) = \mathbb{P}(\vec{Y}_1 | Y_0 = b, \cdot) \mathbb{P}(Y_0 = b | \cdot)$$

gives:

$$\begin{aligned} \mathbb{P}(\vec{Y}_1 | Y_0 = b, a\vec{X}_1, \overleftarrow{(x, y)}) \mathbb{P}(Y_0 = b | a\vec{X}_1, \overleftarrow{(x, y)}) \\ = \mathbb{P}(\vec{Y}_1 | Y_0 = b, a\vec{X}_1, \overleftarrow{(x, y)}') \mathbb{P}(Y_0 = b | a\vec{X}_1, \overleftarrow{(x, y)}') . \end{aligned}$$

The second factors on both sides are equal by causal equivalence. So, there are two cases: These factors either vanish or they do not. If they are positive, then we have:

$$\mathbb{P}(\vec{Y}_1 | Y_0 = b, a\vec{X}_1, \overleftarrow{(x, y)}) = \mathbb{P}(\vec{Y}_1 | Y_0 = b, a\vec{X}_1, \overleftarrow{(x, y)}') .$$

Rewriting the conditional variables with the symbol pair  $(a, b)$  attached to the joint past then gives the first part of the result:

$$\mathbb{P}(\vec{Y}_1 | \vec{X}_1, \overleftarrow{(x, y)}(a, b)) = \mathbb{P}(\vec{Y}_1 | \vec{X}_1, \overleftarrow{(x, y)}'(a, b)) .$$

In the other case, when the factors vanish, we have:

$$\mathbb{P}(Y_0 = b | a\vec{X}_1, \overleftarrow{(x, y)}) = \mathbb{P}(Y_0 = b | a\vec{X}_1, \overleftarrow{(x, y)}') = 0 .$$

This implies that:

$$\begin{aligned} \mathbb{P}(Y_0 = b | X_0 = a, \overleftarrow{(x, y)}) &= \mathbb{P}(Y_0 = b | X_0 = a, \overleftarrow{(x, y)}') \\ &= 0 . \end{aligned}$$

In other words,  $Y_0 = b$  is never observed following either past, given  $X_0 = a$ . That is,  $(a, b)$  is forbidden.

It then follows that:

$$H[S_1 | S_0, (X, Y)_0] = 0 ,$$

which is equivalent to joint unifilarity when there is a finite number of causal states.  $\square$

Unifilarity guarantees that once we know the process is in a particular causal state—we are “synchronized” [33]—we do not lose synchronization over time. This is an important property when using causal states to simulate or predict a system’s behavior. Using presentations that are nonunifilar, typically it is necessary to keep track of a *distribution* over states.

Unifilarity is also a useful property to have when attempting to infer an  $\epsilon$ -transducer from data. Inference

of nonunifilar transducers can be challenging, partly due to the existence of multiple possible state paths given a particular start state. Unifilar transducers avoid this problem, effectively reducing the difficulty to that of inferring a Markov chain from data [37]. Finally, unifilarity plays a key role, as a sequel shows, in calculating channel information quantities.

The next theorem shows that  $\epsilon$ -transducers are input-dependent hidden Markov models.

**Proposition 6** (Markovity). *A channel's causal states satisfy the conditional Markov property:*

$$\mathbb{P}(S_t | X_{t-1}, \overleftarrow{S}_t) = \mathbb{P}(S_t | X_{t-1}, S_{t-1}) .$$

*Proof.* Since the causal-state transitions are unifilar, there is a well defined set of output symbols  $\mathcal{Z} \subseteq \mathcal{Y}$  that causes a transition from state  $\sigma_j$  to state  $\sigma_k$ . We therefore have:

$$\begin{aligned} \mathbb{P}(S_t = \sigma_k | X_{t-1}, S_{t-1} = \sigma_j, \overleftarrow{S}_{t-1}) \\ = \mathbb{P}(Y_{t-1} \in \mathcal{Z} | X_{t-1}, S_{t-1} = \sigma_j, \overleftarrow{S}_{t-1}) . \end{aligned}$$

Causal shielding applies to finite futures as well as infinite. This, combined with the observation that  $\overleftarrow{S}_{t-1}$  is purely a function of the past, allows us to use  $S_{t-1}$  to causally shield  $Y_{t-1}$  from  $\overleftarrow{S}_{t-1}$ , giving:

$$\begin{aligned} \mathbb{P}(Y_{t-1} \in \mathcal{Z} | X_{t-1}, S_{t-1} = \sigma_j, \overleftarrow{S}_{t-1}) \\ = \mathbb{P}(Y_{t-1} \in \mathcal{Z} | X_{t-1}, S_{t-1} = \sigma_j) \\ = \mathbb{P}(S_t = \sigma_k | X_{t-1}, S_{t-1} = \sigma_j) . \end{aligned}$$

The final equality is again possible due to unifilarity.  $\square$

The following theorem shows that the causal states store as much information as possible (from the past) about a channel's future behavior—a desirable property for any predictive model.

**Definition 9.** *The prescience of a set  $\mathcal{R}$  of rival states—equivalence classes of an alternative partition of pasts—reflects how well the rival states predict a channel's future behavior. Quantitatively, this is monitored by the amount of information they share with future output, given future input:*

$$\mathbb{I}[R ; \overrightarrow{Y} | \overrightarrow{X}] ,$$

where  $R$  is the associated rival-state random variable.

Note that it is sometimes simpler to prove statements about conditional entropy than it is for mutual information. Due to this, we will transform statements about prescience into statements about uncertainty in prediction in several proofs that follow. Specifically, we will make use of the identity:

$$\begin{aligned} \mathbb{I}[R_0 ; \overrightarrow{Y}_0 | \overrightarrow{X}_0] &= \lim_{L \rightarrow \infty} \mathbb{I}[R_0 ; Y_{0:L} | \overrightarrow{X}_0] \\ &= \lim_{L \rightarrow \infty} \left( \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0] - \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, R_0] \right) , \end{aligned}$$

where  $\mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, R_0]$  is the finite-future *prediction uncertainty*. Note that the infinite-future prediction uncertainty  $\mathbb{H}[\overrightarrow{Y}_0 | \overrightarrow{X}_0, R_0]$  typically will be infinite, but rewriting the prescience in terms of the limit of finite-future prediction uncertainty allows us to continue to work with finite quantities.

**Theorem 1** (Maximal Prescience). *Among all rival partitions  $\mathcal{R}$  of joint pasts, the causal states have maximal prescience and they are as prescient as pasts:*

$$\begin{aligned} \mathbb{I}[S ; \overrightarrow{Y} | \overrightarrow{X}] &= \mathbb{I}[\overleftarrow{(X, Y)} ; \overrightarrow{Y} | \overrightarrow{X}] \\ &\geq \mathbb{I}[R ; \overrightarrow{Y} | \overrightarrow{X}] . \end{aligned}$$

*Proof.* We will prove the equivalent statement that the causal states *minimize* finite-future prediction uncertainty for futures of any length  $L$  and have the same finite-future prediction uncertainty as pasts; i.e., that for all  $L$ :

$$\begin{aligned} \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, S_0] &= \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, \overleftarrow{(X, Y)}_0] \\ &\leq \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, R_0] , \end{aligned}$$

Like the causal states, rival states—equivalence classes of an alternative partition  $\mathcal{R}$  of pasts—are a function of the past:

$$R = \eta(\overleftarrow{(X, Y)}) .$$

By the Data Processing Inequality [12], we have:

$$\mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, \overleftarrow{(X, Y)}_0] \leq \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, R_0] .$$

The causal states share future morphs with their corresponding pasts. By simple marginalization, the same is true for finite-future morphs:

$$\begin{aligned} \mathbb{P}(Y_{0:L} | \overrightarrow{X}_0, S_0) &= \mathbb{P}(Y_{0:L} | \overrightarrow{X}_0, \overleftarrow{(X, Y)}_0) \\ \implies \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, S_0] &= \mathbb{H}[Y_{0:L} | \overrightarrow{X}_0, \overleftarrow{(X, Y)}_0] . \end{aligned}$$

$\square$

Note that this proof also shows that the causal states are maximally prescient for all length- $L$  futures:  $\mathbb{I}[S ; Y_{0:L} | X_{0:L}] = \mathbb{I}[\overleftarrow{(X, Y)}_0 ; Y_{0:L} | X_{0:L}]$ .

The causal equivalence relation can also be applied directly to a channel that anticipates its future inputs, but the resulting  $\epsilon$ -transducer has outputs that depend not only upon the current state and input symbol, but some set of *future* input symbols. If the set is finite, then one uses the previous construction that transforms finite anticipation into additional transducer memory (statistical complexity). In this case, the causal states still capture all of the information from the past needed for prediction. That is, they have maximal prescience. Any additional future input dependence, though, must be encoded in the machine's transitions.

**Definition 10.** *A prescient rival is an indexed set  $\widehat{\mathcal{R}}$  of*



states (with elements  $\hat{\rho}_i$  and random variable  $\hat{R}$ ) that is as predictive as any past:

$$I[\hat{R}; \vec{Y} | \vec{X}] = I[(\overline{X}, Y); \vec{Y} | \vec{X}].$$

**Lemma 1** (Refinement). *The partition of a prescient rival  $\hat{\mathcal{R}}$  is a refinement (almost everywhere) of the causal-state partition of the joint input-output pasts.*

*Proof.* Since the causal states' future morphs include every possible future morph of a channel, we can always express a prescient rival's future morph as a (convex) combination of the causal states' future morphs. This allows us to rewrite the entropy over a prescient rival (finite-length) future morph as:

$$\begin{aligned} H[Y_{0:L} | \vec{x}, \hat{\rho}_k] &= H[\mathbb{P}(Y_{0:L} | \vec{x}, \hat{\rho}_k)] \\ &= H\left[\sum_j \mathbb{P}(Y_{0:L} | \vec{x}, \sigma_j) \mathbb{P}(\sigma_j | \hat{\rho}_k)\right]. \end{aligned} \quad (11)$$

Since entropy is convex, we also have:

$$\begin{aligned} H\left[\sum_j \mathbb{P}(Y_{0:L} | \vec{x}, \sigma_j) \mathbb{P}(\sigma_j | \hat{\rho}_k)\right] \\ \geq \sum_j \mathbb{P}(\sigma_j | \hat{\rho}_k) H[Y_{0:L} | \vec{x}, \sigma_j]. \end{aligned} \quad (12)$$

Therefore:

$$\begin{aligned} H[Y_{0:L} | \vec{X}_0, \hat{R}_0] &= \sum_k \mathbb{P}(\hat{\rho}_k) H[Y_{0:L} | \vec{X}_0, \hat{\rho}_k] \\ &\geq \sum_k \mathbb{P}(\hat{\rho}_k) \sum_j \mathbb{P}(\sigma_j | \hat{\rho}_k) H[Y_{0:L} | \vec{X}_0, \sigma_j] \\ &= \sum_{j,k} \mathbb{P}(\sigma_j, \hat{\rho}_k) H[Y_{0:L} | \vec{X}_0, \sigma_j] \\ &= \sum_j \mathbb{P}(\sigma_j) H[Y_{0:L} | \vec{X}_0, \sigma_j] \\ &= H[Y_{0:L} | \vec{X}_0, S_0], \end{aligned}$$

where the inequality follows from Eqs. (11) and (12). Since the rival states  $\hat{R}$  are prescient, we know that equality must be attained in this inequality for each  $L$ . Equality is only possible when  $\mathbb{P}(\sigma_j | \hat{\rho}_k) = 1$  for exactly one value of  $j$  and vanishes for every other  $j$ . That is, if a rival state is prescient, it is contained entirely within a single causal state, aside from a set of measure zero. Thus, the partition of the prescient rival states is a refinement of the causal-state partition almost everywhere.  $\square$

**Theorem 2** (Minimality). *For any given input process  $X$ , causal states have the minimal conditional statistical complexity among all prescient rival partitions  $\hat{R}$ :*

$$C_X(S) \leq C_X(\hat{R}).$$

*Proof.* Since a prescient rival partition is a refinement almost everywhere, there exists a function  $f$  defined almost everywhere that maps each rival state to the causal state that (almost everywhere) contains it:

$$f(\hat{\rho}_i) = \sigma_j.$$

Then, we have:

$$\begin{aligned} H_X[\hat{R}] &\geq H_X[f(\hat{R})] \\ &= H_X[S]. \end{aligned}$$

$\square$

**Corollary 1.** *Causal states minimize the channel complexity  $\overline{C}_\mu$ .*

*Proof.* Immediate from the preceding theorem.  $\square$

In words, we established the fact that the causal states store all of the information contained in the past that is necessary for predicting a channel's future behavior and as little of the remaining information "overhead" contained in the past as possible. Given an input process,  $\epsilon$ -transducer causal states maximize  $I[\vec{Y}; S | \vec{X}]$  while minimizing  $I[(\overline{X}, Y); S]$ .

The final optimality theorem shows that any states which have these properties are in fact the causal states.

**Theorem 3** (Uniqueness). *The  $\epsilon$ -transducer is the unique prescient, minimal partition of pasts. If  $C_X(\hat{R}) = C_X(S)$  for every input process  $\vec{X}$ , then the corresponding states  $\hat{R}$  and  $S$  are isomorphic to one another almost everywhere. And, their equivalence relations  $\sim_\eta$  and  $\sim_\epsilon$  are the same almost everywhere.*

*Proof.* Again, the Refinement Lemma (Lemma 1) says that  $S = f(\hat{R})$  almost everywhere. It therefore follows that  $H_X[S | \hat{R}] = 0$ . Moreover, by assumption  $H_X[S] = H_X[\hat{R}]$ . Combining these with the symmetry of mutual information gives:

$$\begin{aligned} I_X[S; \hat{R}] &= I_X[\hat{R}; S] \\ H_X[S] - H_X[S | \hat{R}] &= H_X[\hat{R}] - H_X[\hat{R} | S] \\ H_X[S] - 0 &= H_X[S] - H_X[\hat{R} | S] \\ H_X[\hat{R} | S] &= 0. \end{aligned}$$

The latter holds if and only if there is a function  $g$  such that  $\hat{R} = g(S)$  almost everywhere. By construction  $g$  is the inverse  $f^{-1}$  of  $f$  almost everywhere. We have that  $f \circ \eta = \epsilon$  and  $f^{-1} \circ \epsilon = \eta$ .

Finally, the two equivalence relations  $\sim_\epsilon$  and  $\sim_\eta$  are

the same almost everywhere:

$$\begin{aligned}
& \overline{(x, y)} \sim_\epsilon \overline{(x, y)'} \\
& \implies \epsilon(\overline{(x, y)}) = \epsilon(\overline{(x, y)'}) \\
& \implies f^{-1} \circ \epsilon(\overline{(x, y)}) = f^{-1} \circ \epsilon(\overline{(x, y)'}) \\
& \implies \eta(\overline{(x, y)}) = \eta(\overline{(x, y)'}) \\
& \implies \overline{(x, y)} \sim_\eta \overline{(x, y)'},
\end{aligned}$$

and

$$\begin{aligned}
& \overline{(x, y)} \sim_\eta \overline{(x, y)' } \\
& \implies \eta(\overline{(x, y)}) = \eta(\overline{(x, y)'}) \\
& \implies f \circ \eta(\overline{(x, y)}) = f \circ \eta(\overline{(x, y)'}) \\
& \implies \epsilon(\overline{(x, y)}) = \epsilon(\overline{(x, y)'}) \\
& \implies \overline{(x, y)} \sim_\epsilon \overline{(x, y)'}.
\end{aligned}$$

□

$\epsilon$ -Transducer uniqueness means that  $C_X$  is *the* conditional complexity of a channel and, therefore, justifies calling  $\overline{C}_\mu$  *the* channel complexity.

## X. GLOBAL $\epsilon$ -MACHINE VERSUS $\epsilon$ -TRANSDUCER

Given a particular joint process or its global  $\epsilon$ -machine, it is possible (provided that the input process satisfies certain requirements) to construct the  $\epsilon$ -transducer that maps input  $\overline{X}$  to output  $\overline{Y}$ , such that  $(\overline{X}, \overline{Y}) = (\overline{X}, f(\overline{X}))$ , where  $f$  is the transducer and  $f(\overline{X})$  designates the output of the transducer, given input process  $\overline{X}$ . Sequels address the relationship between a joint process' global  $\epsilon$ -machine and corresponding  $\epsilon$ -transducer at both the process (channel) level and at the automata ( $\epsilon$ -machine and  $\epsilon$ -transducer) level. There, we provide algorithms for “conditionalizing” a joint process or  $\epsilon$ -machine to obtain the corresponding channel or  $\epsilon$ -transducer, as well as algorithms for obtaining input or output marginals, applying an  $\epsilon$ -transducer to an input  $\epsilon$ -machine, composing multiple  $\epsilon$ -transducers, and inverting an invertible  $\epsilon$ -transducer.

Note that the ability to construct an  $\epsilon$ -transducer from a joint process can be useful when attempting to infer an  $\epsilon$ -transducer from data, as such data will typically come from a system driven by some *particular* (possibly controllable) input; i.e., the data is a sample of a *joint process*.

## XI. HISTORY $\epsilon$ -TRANSDUCER VERSUS GENERATOR $\epsilon$ -TRANSDUCER

The preceding focused on the *history* specification of an  $\epsilon$ -transducer, where a machine is obtained by partitioning a channel's histories (joint pasts). We can also consider the *generator* specification, where we instead *start* with a machine that produces a stationary, ergodic channel. Taking this perspective, an  $\epsilon$ -transducer is an input-dependent, strongly connected, aperiodic hidden Markov model with unifilar transitions and probabilistically distinct states. The history and generator specifications of an  $\epsilon$ -transducer are likely equivalent—as they are with  $\epsilon$ -machines [21]—but we leave such a proof to future work.

## XII. EXAMPLES REVISITED

With the  $\epsilon$ -transducer defined, we can revisit the example channels examined above. We display channel structure via its  $\epsilon$ -transducer's state transition diagram, as well as a wordmap that now colors each joint history based on its corresponding causal state. Input and output history projections are also shown. Histories that are not mapped to a single causal state are colored black. Recall that causal states partition *joint* histories, so input or output histories alone need not correspond to a unique causal state. We will discuss Markov orders for these channels, but we leave it to the reader to construct an exhaustive list of Markov orders for each channel.

Since the Identity (Fig. 7), All is Fair (Fig. 8), and Z (Fig. 9) Channels are memoryless, their behavior does not depend on the past. As a result, there is a single causal state containing every past and so their wordmaps are monochromatic. Since they have a single causal state,  $C_X = \overline{C}_\mu = 0$ .

In contrast, the Delay Channel (Fig. 10) has two causal states and so two colors corresponding to pasts in which the input ends on a 0 (left half of the wordmap) or a 1 (right half of the wordmap). This partitioning into halves is a characteristic of channels with a pure feedforward Markov order  $R_{\text{pff}} = 1$ . We also see that the output words are colored black, illustrating the fact that the output tells us nothing about the current causal state. The channel's pure feedforward Markov order of 1 can be seen in the channel's  $\epsilon$ -transducer state-transition diagram by observing that all transitions on input symbol 0 lead to state  $A$  and all transitions on input symbol 1 lead to state  $B$ . Since the Delay Channel is undefined for outputs alone ( $R_{\text{pfb}}$  is undefined), it is the first example channel with a nontrivial irreducible feedforward order:  $R_{\text{iff}} = 1$ . The causal states of the Delay Channel simply store a single bit of input, their entropy therefore matches the length-1 block entropy of the input process:  $C_X = H[X_0]$ . Maximizing the latter over input process gives us  $\overline{C}_\mu = 1$ , attained with Fair Coin Process input.

The Feedforward NOR Channel (Fig. 11) also stores

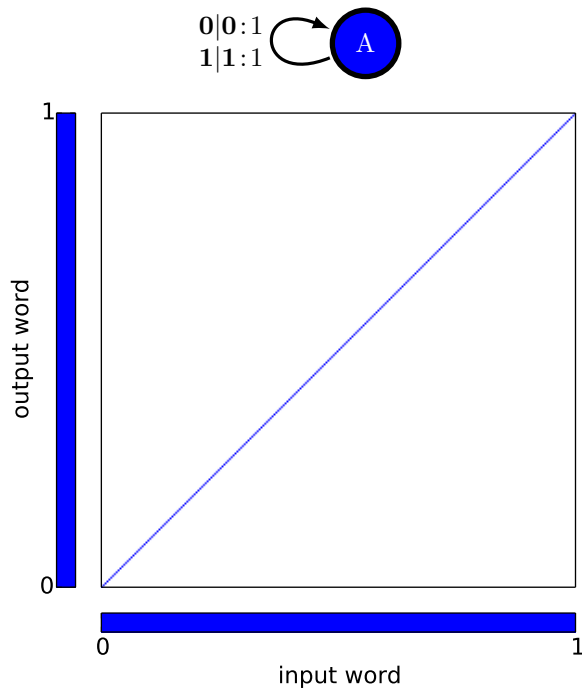


FIG. 7. Identity Channel:  $\epsilon$ -transducer and causal-state colored wordmap. See text for explanation.

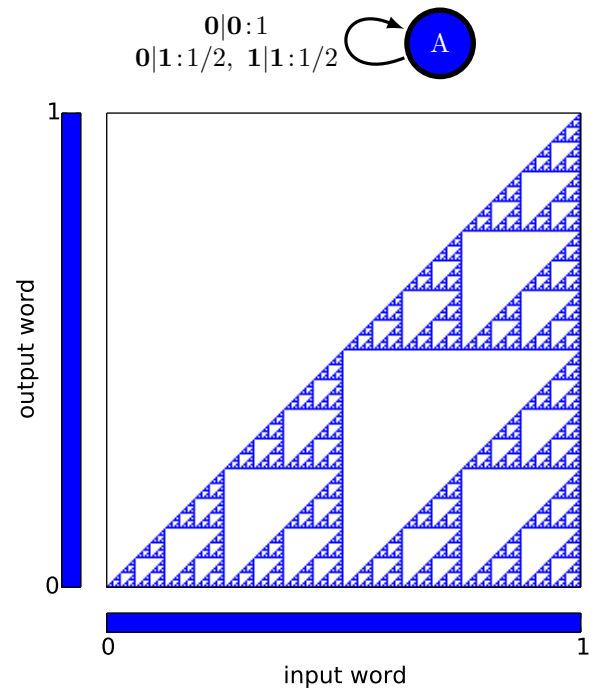


FIG. 9. Z Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

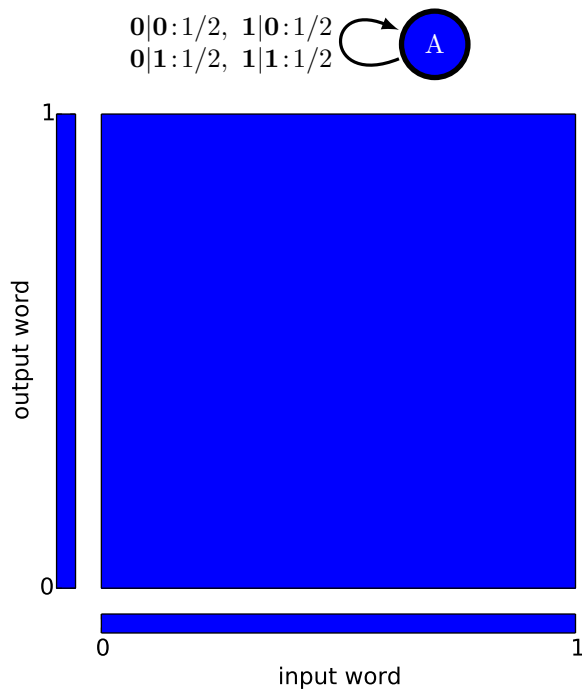


FIG. 8. All-Is-Fair Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

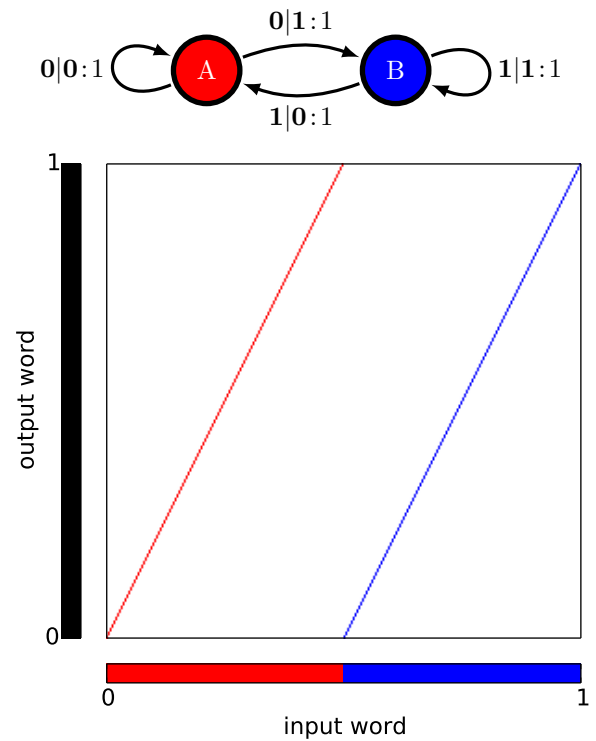


FIG. 10. Delay Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

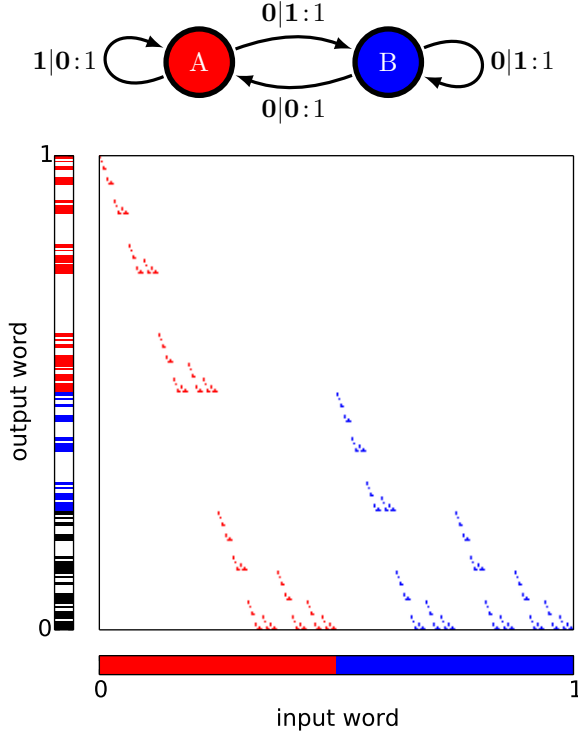


FIG. 11. Feedforward NOR Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

the previous input symbol and, therefore, partitions input histories the same way ( $R_{\text{pff}} = 1$ ). We see in the wordmap that there is again an ambiguity of causal state given output histories alone and there is, therefore, no pure feedback presentation for the channel. Specifically, we see that the ambiguity arises for histories where the output ends on two 0s (lower quarter of the wordmap). This can be verified in the channel's  $\epsilon$ -transducer by observing that a 1 on output always leads to state A, but a 0 on output only leads to a unique state if it is followed by a 1 on output. A single symbol of input is *always* needed to guarantee well defined behavior ( $R_{\text{iff}} = 1$ ). Since the Feedforward NOR Channel's causal states store the same information as the Delay Channel, we can again drive the channel with the Fair Coin Process to attain  $\overline{C}_\mu = 1$  bit.

The wordmap for the Odd NOT Channel (Fig. 12) has projected input (and output) partitions with structure at all scales. This is the signature of states that depend upon infinite histories—one must provide an arbitrarily long binary expansion to specify the location of the causal state boundaries and, therefore, the causal states themselves. If we observe both inputs and outputs, we only need to specify in which quadrant a joint history lies in order to determine its causal state. That is, the Odd NOT Channel is sofic on both input and output alone (infinite pure feedforward and pure feedback Markov orders,  $R_{\text{pff}}$  and  $R_{\text{pfb}}$ , respectively), but Markovian when both input and output are considered (finite Markov order  $R$ ). We also see that the causal

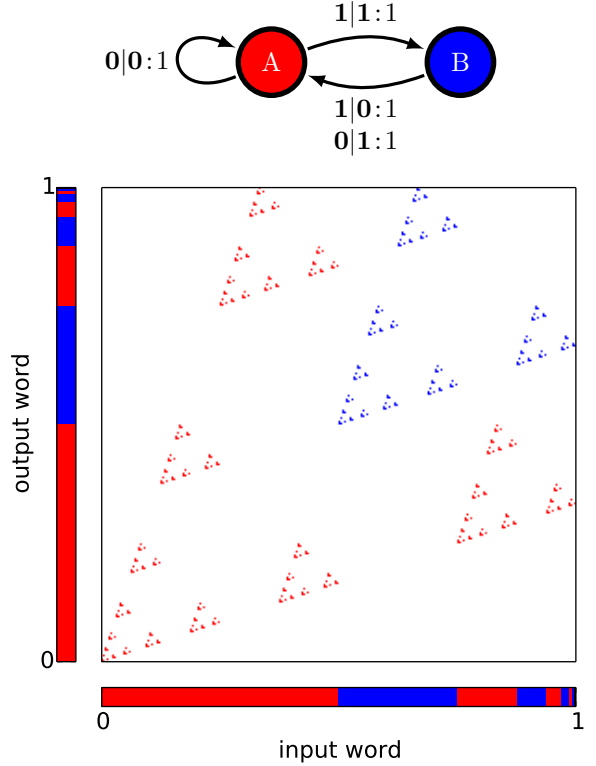


FIG. 12. Odd NOT Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

states store the *same information* (parity) about input histories as they do output histories, by observing the symmetry along the diagonal. Since the Period-2 Process generates sequences that always alternate between even and odd parity, we can drive the channel with this process to induce a uniform distribution over its causal states. Therefore, we have  $\overline{C}_\mu = 1$  bit, again.

We see that the All-Is-Golden Channel (Fig. 13) has a pure feedback Markov order of  $R_{\text{pfb}} = 1$ . Since it ignores its input, however, input histories tell us nothing about in which causal state the channel is. We also see that the wordmap is horizontally symmetric due to this lack of input dependence. Since the state transitions depend only on output, the state distribution and, therefore, the statistical complexity are independent of input. In particular, the channel's statistical complexity is that of the Golden Mean Process (GMP):  $C_X = \overline{C}_\mu = C_\mu(\text{GMP}) \approx 0.918$  bits.

The wordmap for the Feedback NOR Channel (Fig. 14) clearly shows that it has infinite pure feedforward Markov order  $R_{\text{pff}}$ , but finite Markov  $R$  and pure feedback Markov  $R_{\text{pfb}}$  orders. Contrast this with the wordmap for the Feedback XOR Channel (Fig. 15) clearly showing that the causal state, and so the channel's behavior, cannot be determined by input alone. Observe that the Feedback NOR Channel is in state A with probability 1 when a 1 is observed on input and oscillates between states A and B if the channel is driven with a

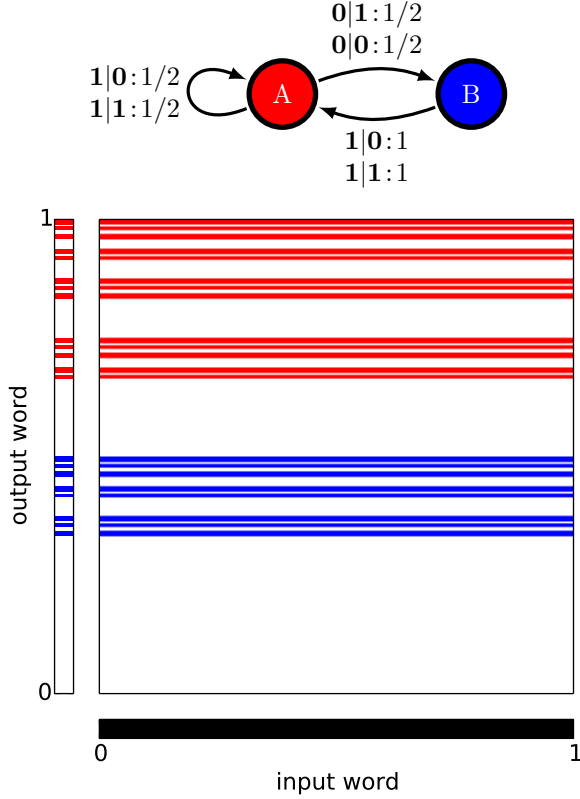


FIG. 13. All-Is-Golden Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

period-2 cycle of 0s and 1s from that point on. We can therefore induce a uniform distribution over causal states by driving the channel with the Period-2 Process. We can also induce a uniform distribution over the Feedback XOR Channel’s causal states by driving the channel with the Fair Coin Process, which causes all state transitions to occur with equal probability. In both cases,  $\overline{C}_\mu = 1$  bit.

The Odd Random Channel (Fig. 16) has infinite pure feedforward and pure feedback Markov orders ( $R_{\text{pff}} = R_{\text{pfb}} = \infty$ ), but unlike the Odd NOT channel, the Markov order  $R$  is infinite. In the Odd NOT channel, we saw structure at all scales in the input and output projections of the wordmap, but a partitioning into quadrants in the complete wordmap. Now, we see that there is no such simple partition in the complete wordmap, and there is structure at all scales in the coloring of histories. In other words, one must specify arbitrarily long *pairs* of binary expansions (input and output) in order to specify a causal state. Specifically, knowing that a past ended in  $(x, y)_t = (1, 1)$  does not uniquely determine a causal state. This can be seen as multiple colors appearing in the upper-right quadrant of the wordmap. If, however, we know that the *previous* pair was  $(0, 0)$ ,  $(0, 1)$ , or  $(1, 0)$ , we see that the history maps to causal state  $B$ ; corresponding to the upper-left, lower-left, and lower-right subquadrants of the upper-right quadrant, respectively.

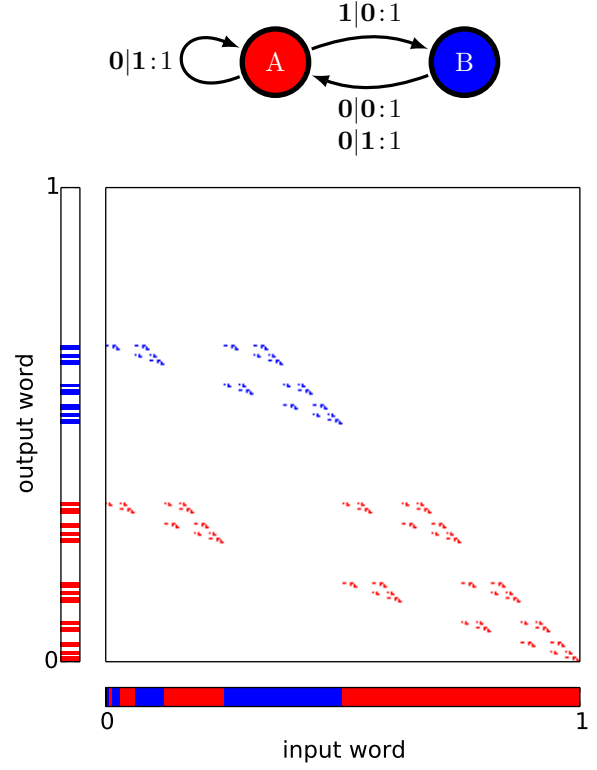


FIG. 14. Feedback NOR Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

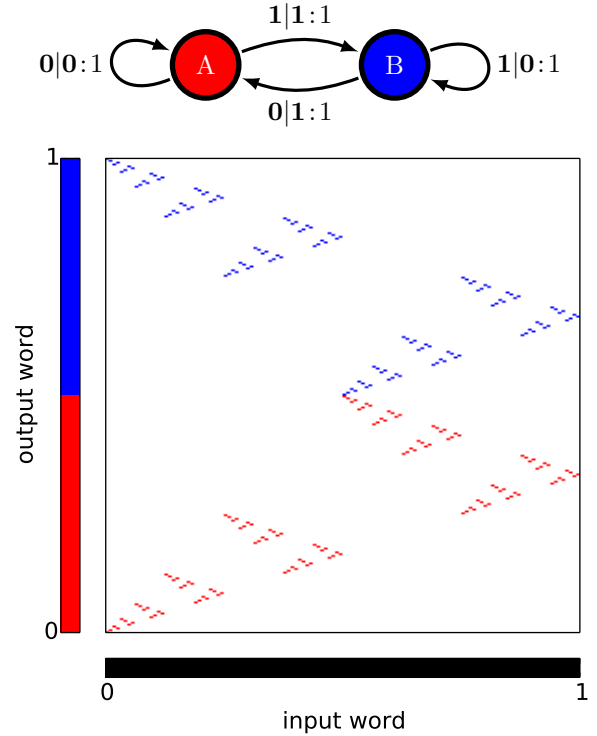


FIG. 15. Feedback XOR Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

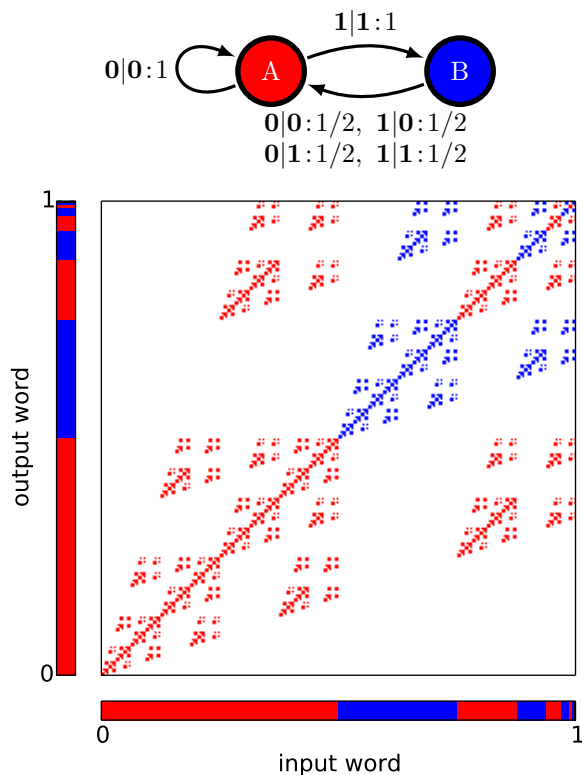


FIG. 16. Odd Random Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

Similarly, if the previous pair was  $(1, 1)$ , we are left with an ambiguity in state; corresponding to the upper-right subquadrant of the upper-right quadrant. Therefore, we require an arbitrarily long past to determine the channel's causal state in general. Since the causal states store the same parity as the Odd NOT channel, we can again drive the channel with the Period-2 Process to induce a uniform distribution, giving us  $\overline{C_\mu} = 1$  bit.

The Period-2 Identity NOT Channel (Fig. 17) has a Markov order of  $R = 1$ , but we clearly see that neither input nor output alone determines the channel's causal state ( $R_{\text{iff}} = R_{\text{ifb}} = 1$ ). Since the states have a uniform distribution regardless of input, we have  $C_X = \overline{C_\mu} = 1$  bit.

### XIII. INFINITE-STATE $\epsilon$ -TRANSDUCERS: THE SIMPLE NONUNIFILAR CHANNEL

The example channels were chosen to have a finite number of causal states (typically two), largely to keep the analysis of their structure accessible. We can see that even with a few states,  $\epsilon$ -transducers capture a great deal of behavioral richness. Nonetheless, many channels have an infinite number of causal states. Consider, for example, the *Simple Nonunifilar Channel*. This channel's behavior is captured simply by the finite-state presentation shown in Fig. 18. When in state  $A$ , the channel behaves

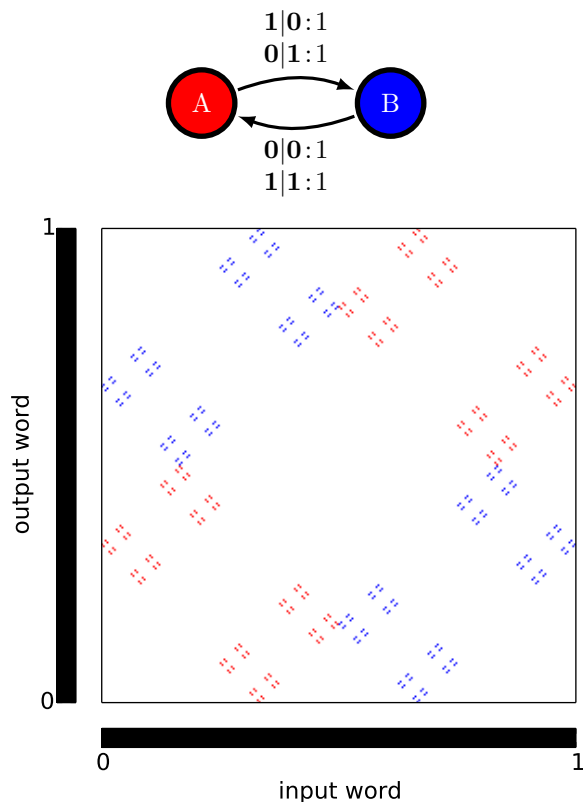


FIG. 17. Period-2 Identity NOT Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

as the identity and has an equal probability of staying in state  $A$  or transitioning to state  $B$ . When in state  $B$ , the channel will either behave as the identity and transition back to state  $B$  or behave as the bit flipped identity and transition to state  $A$ , each with equal probability.

Observe that the transducer shown is *nonunifilar* and is, therefore, *not* the  $\epsilon$ -transducer for the channel. For example, the joint symbol  $(0, 0)$  can cause state  $A$  to transition to either itself or state  $B$ . This nonunifilarity manifests in the wordmap as large blocks of black points. These indicate joint histories that lead to a *mixture* of transducer states. This illustrates the fact that an observer cannot typically retain synchronization to a *particular* state of a nonunifilar transducer—a problem not present when using unifilar transducers.

It is possible to construct the  $\epsilon$ -transducer for the Simple Nonunifilar Channel, but doing so results in a transducer with a countably infinite set of states. This minimal, unifilar  $\epsilon$ -transducer can be seen in Fig. 19. Since any channel with a finite Markov order (and finite alphabet) will have a finite number of causal states, a channel with an infinite number of causal states will have infinite Markov order. This is also evident in the causal-state wordmap for the Simple Nonunifilar Channel, as one needs infinite resolution (in general) to determine to which causal state a joint past leads. In fact, even if we know either the infinite input or output past, we still need

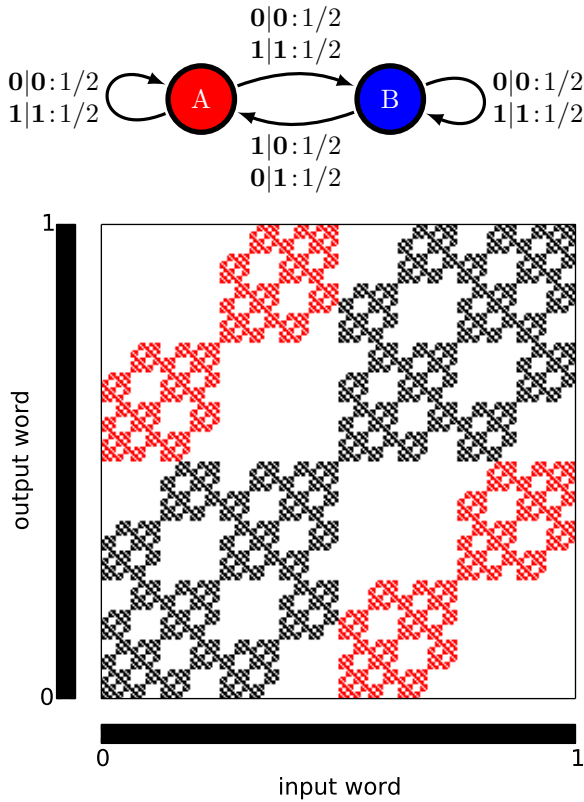


FIG. 18. Simple Nonunifilar Channel: *Nonunifilar* transducer presentation and state colored wordmap.

to know the full output or input past, respectively, in order to characterize the channel’s behavior. This is therefore the first example we have seen with  $R_{\text{iff}} = R_{\text{ifb}} = \infty$ .

Observe that while the Simple Nonunifilar Channel’s output clearly depends upon its input, its state-to-state transitions do *not*. Its statistical complexity is therefore independent of the input process chosen. In fact, the causal states and transitions between them are identical to the Simple Nonunifilar Source [22]. The statistical complexity is therefore equal to the statistical complexity of the Simple Nonunifilar source:  $C_X = \overline{C}_\mu \approx 2.71$  bits. Even though there are an infinite number of states, the  $B_i$  states are occupied with probability that decreases quickly with  $i$ , thus allowing for a finite Shannon state entropy. Note that if one were to use Fig. 18’s *nonunifilar* presentation for the Simple Nonunifilar Channel, the statistical complexity would be underestimated as  $C_X = \overline{C}_\mu = 1$  bit.

XIV. DISCUSSION

Previously, we described computational mechanics in the setting of either generating or controlling processes [15]. As noted there, generation and control are complementary. Here, we developed computational mechanics in a way that merges both control (the input process) and

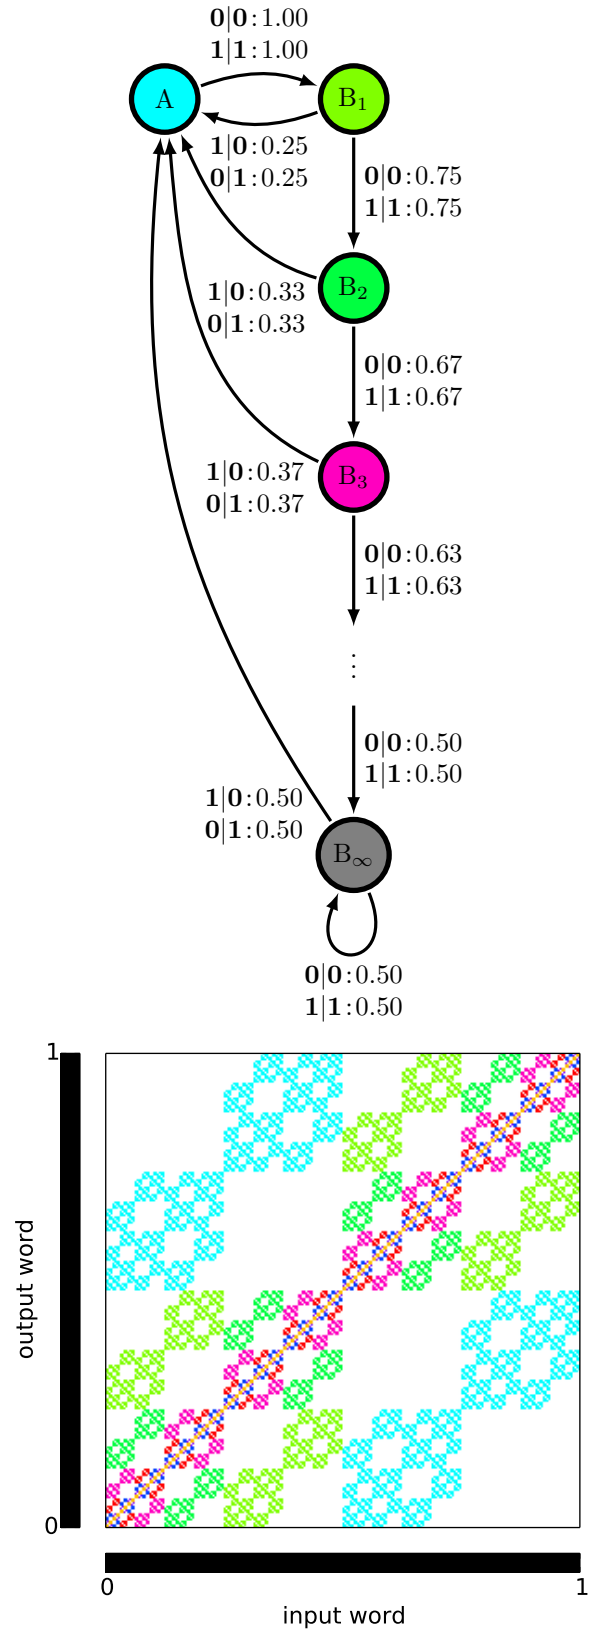


FIG. 19. Simple Nonunifilar Channel:  $\epsilon$ -transducer and causal-state colored wordmap.

generation (the output process), extending the  $\epsilon$ -machine to the  $\epsilon$ -transducer. With this laid out, we describe how the  $\epsilon$ -transducer overlaps and differs from alternatives to modeling input-output processes. We then turn to discuss applications, which incidentally elucidate our original motivations, and suggest future directions.

### A. Related Work: Modeling

Following the signposts of earlier approaches to modeling complex, nonlinear dynamical systems [2, 38], we are ultimately concerned with reconstructing a transducer when given a general channel or given a joint process, either analytically or via statistical inference. And so, when discussing related efforts, we distinguish between those whose goal is to extract a model, which we review now, and those that analyze types of transductions, which we review next. After this, we turn to applications.

For statistical estimation we note that the recently introduced Bayesian Structural Inference (BSI) [37] allows one to estimate the posterior probability that  $\epsilon$ -machines generate a given, even relatively short, data series. BSI’s generality allows it to be readily adapted to infer  $\epsilon$ -transducers from samples of an input-output process. This turns on either developing an enumeration of  $\epsilon$ -transducers which parallels that developed for  $\epsilon$ -machines in Ref. [39] or on developing a list of candidate  $\epsilon$ -transducers for a given circumstance. And, these are also readily accomplished. A sequel provides the implementations. Previously, inferring causal states, and so causal-state filters, had also been addressed; see, for example, Refs. [37, 40, 41].

Optimal transducers were originally introduced as structure-based filters to define hierarchical  $\epsilon$ -machine reconstruction [22, 42] in terms of causal-state filtering, to detect emergent spatiotemporal patterns [43–46], and to explore the origins of evolutionary selection pressure [47] and the evolution of language structure [48]. These causal-state transducers were first formalized in Ref. [49] and several of those results are reproduced in Ref. [50]. Appendix A shows that the definition there, which makes additional assumptions compared to that here, are equivalent. The more-general development is more elegant, in that it establishes unifilarity, for example, rather than assume such a powerful property. Likely, in addition, the generality will allow  $\epsilon$ -transducers to be more widely used.

Throwing the net wider—beyond these, most directly related, prior efforts—there have been many approaches to modeling input-output mappings. We will use the fact that most do not focus on quantitatively analyzing the mapping’s intrinsic structure to limit the scope of our comments. We mention a few and then only briefly. Hopefully the list, nonetheless, suggests directions for future work in these areas.

Today, many fall under the rubric of learning, though they are rather more accurately described as statistical

parameter estimation within a fixed model class. Probably, the most widely used and developed methods to model general input-output mappings are found in artificial neural networks [51, 52] and in the more modern approaches that employ kernel methods [53], statistical physics [54], and information theory [55, 56]. Often these methods require IID-drawn samples and so do not directly concern mappings from one temporal process to another. Unlike  $\epsilon$ -transducers, they are also typically limited to model classes—e.g., feedforward and directed acyclic graph structures—that do not allow internal feedback or dynamics.

That said, neural networks that are recurrent are universal approximators of dynamical systems and, per force, are channels with feedback and feedforward memory [57]. They are well known to be hard to train and, in any case, rarely quantitatively analyzed for the structures they capture when successfully trained. In the mathematical statistics of time series, for comparison, AutoRegressive-Moving-Average model with eXogenous inputs model (ARMAX models) are channels with feedback and feedforward memory, but they are *linear*—current output is a linear combination of past inputs and outputs. The nonlinear generalization is the Nonlinear AutoRegressive eXogenous model (NARX), which is a very general memoryful causal channel. At some future time, likely using  $\epsilon$ -transducers extended to continuous variables as recently done for  $\epsilon$ -machines in Ref. [58], we will understand better the kinds of structure these channels can represent.

### B. Related Work: Classification

Beyond developing a theoretical framework for structured transformations, one that is sufficiently constructive to be of use in statistical inference, there are issues that concern how they give a new view, if any, of the organization of the space of structured processes itself.

Specifically, computational mechanics up to this point focused on processes and developed  $\epsilon$ -machines to describe them as stochastic sets.  $\epsilon$ -Machines are, most simply stated, compact representations of distributions over sequences. With the  $\epsilon$ -transducers introduced here, computational mechanics now has formalized stochastic mappings of these stochastic sets. And, to get to the point, with sets and mappings one finally has a framework capable of addressing the recoding equivalence notion and the geometry of the space of processes proposed in Ref. [59]. A key component of this will be a measure of distance between processes that uses a structural measure from the minimal optimal mapping ( $\epsilon$ -transducer) between them. This would offer a constructive, in the sense we use the word, approach to the view of process space originally introduced by Shannon [13, 60, 61].

This then leads to the historically prior question of structurally classifying processes—paralleling schemes developed in computation theory [24]. Indeed, our devel-



opment is much closer to input-output processes from the earliest days of dynamical systems and automata theory—which were concerned with exploring the range of behaviors of mechanical systems and the then-new digital computers.

Briefly,  $\epsilon$ -transducers are probabilistic endomorphisms of subshifts as studied in symbolic dynamics [25]. The (nonprobabilistic) endomorphisms there were developed to explore the equivalence of processes via conjugacies. Notably, this area grew out of efforts in the 1920s and 1930s by Hedlund, Morse, Thue, and others to define symbolic dynamical systems that were more analytically tractable than continuum-state systems [62]. Their efforts played a role in forming the logical foundations of mathematics and so eventually in the emergence of a theory of computation via Church, Gödel, Post, and Turing [63–66]. This led eventually to Moore’s abstractions of sequential machines and transducers [67] and to Huffman’s concept of a minimal implementation [68] and information lossless automata [69–71]. Today, though expositions are increasingly rare, finite-state transducers are covered by several texts on computation theory; see, for example, Ref. [72].

Once one allows for distributions over sequences, though, then one shifts from the overtly structural approach of symbolic dynamics and automata to Shannon’s information sources and communication channels [26] and a strong emphasis on stochastic process theory. As noted in the introduction, one principal difference is that here we considered channels with memory, while the latter in its elementary treatments considers memoryless channels or channels with very restricted forms of memory. Finite-state channels have been developed in limited way, though; for example, see Ref. [73, Ch. 7] and for very early efforts see Refs. [74] and [75]. There are also overlaps, as we attempted to show in the selected examples, with classifications developed in digital filter theory [76].

There are also differences in focus and questions. Whereas information theory [12, 26] studies *quantities* of information such as intrinsic randomness and informational correlation, computational mechanics [1] goes an additional step and attempts to quantify the information *itself*—the *computational structure* or memory within a system. This is achieved not by assuming a class of model directly, but by making a simple assumption about modeling itself: The only relevant information is that which contributes to prediction—the “difference that makes a difference” to the future [77]. Via the causal equivalence relation, this assumption leads directly to the *unique*, *maximally predictive*, and *minimally complex* model of our measurement data—the  $\epsilon$ -*machine*. Another way to express this is that  $\epsilon$ -transducers give a constructive way to explore the information theory of channels with and without memory.

## C. Applications

Our development of  $\epsilon$ -transducers was targeted to provide the foundation for several related problems—problems that we will address elsewhere, but will briefly describe here to emphasize general relevance and also to suggest future directions.

### 1. Inference versus Experimentation

If all data collected is produced by a measuring device, then any model formed from that data captures both the structure of the system *and sensor* in combination. Is there a natural separation of measuring instrument from system measured? We can now pose this precisely in terms of  $\epsilon$ -transducers: Are there optimal decompositions of a process’s  $\epsilon$ -machine into a possibly smaller  $\epsilon$ -machine (representing the hidden process) composed with a  $\epsilon$ -transducer (representing the measuring instrument)?

### 2. Information Flow within and between Systems

In nonlinear dynamics and in information theory there has been a long-lived interest in how information “flows” and how such flows relate to a system’s mechanical organization; see Refs. [78–81], to mention only a few. These employ specializations of Eq. (6)’s excess entropy, being various forms of conditional mutual information. The *transfer entropy* [82] and the earlier *directional information* [83, 84] are two such. The main issue concerns how one process affects another and so this is a domain in which  $\epsilon$ -transducers—as optimal models of the structured transformations between processes—can help clarify the issues.

In particular, there has been recent criticism of the use of these as measures of information flow and, specifically, their relation to the structural organization of the flows [85]. We can now do better, we believe, since  $\epsilon$ -transducers give a canonical presentation with which to describe and extract the structure of such mappings. And this, in turn, allows one to explicitly relate how causal-state structure supports or precludes information flows. We address this problem in a sequel [86].

### 3. Process Decomposition

Given a process, we can now analyze what internal components drive or are driven by other internal components. As one example, Is a subset of the measurement alphabet the “output” being driven by another subset that is “input”? The question hints at the solution that one can now provide: Produce the  $\epsilon$ -transducer for each bipartite input-output partitioning of the global  $\epsilon$ -machine alphabet, giving a set of candidate input-output models.

One can then invoke, based on a notion of first principles (such as parsimony) or prior knowledge, a way to choose the “best” input-output, driver-drivee decomposition.

#### 4. Perception-Action Cycles

Probably one of the most vexing contemporary theoretical and practical problems, one that occurs quite broadly, is how to describe long-term and emergent features of *dynamic learning* in which a system models its input, makes a decision based on what it has gleaned, and takes an action that affects the environment producing the inputs. In psychology and cognitive sciences this problem goes under the label of the *perception-action cycle*; in neuroscience, under *sensori-motor loop* [7, 8]. The problem transcends both traditional mathematical statistics and modern machine learning, as their stance is that the data is not affected by what is learned. And in this, it transcends the time-worn field of experiment design [87, 88] and the more recent machine learning problem of *active learning* [54]. Though related to computational mechanics via Ref. [41], the recent proposal [89] for *interactive learning* is promising, but is not grounded in a systematic approach to structure. It also transcends control theory, as the latter does not address dynamically building models, but rather emphasizes how to monitor and drive a given system into given states [90].

$\epsilon$ -Transducers suggest a way to model the transduction of sensory input to a model and from the model to a decision process that generates actions. Thus, the computational mechanics representation of the perception-action cycle is two cross-coupled  $\epsilon$ -transducers—one’s output is the other’s input and vice versa. Formulating the problem in this way promises progress in analyzing and in quantifying structures in the space of models and strategies.

Physical applications of  $\epsilon$ -transducers to analyze the information thermodynamics of feedback control in Maxwellian Demons can be seen in Szilard’s Engine [4] and the Mandal-Jarzynski ratchet [3, 5].

## XV. CONCLUSION

Previously, computational mechanics focused on extracting and analyzing the informational and structural properties of individual processes. The premise being that once a process’s  $\epsilon$ -machine had been obtained, it can be studied in lieu of other more cumbersome (or even inappropriate) process presentations. Since the  $\epsilon$ -machine is also unique and minimal for a process, its structure and quantities were treated as being those of the underlying system that generated the process. Strengthening this paradigm, virtually all of a process’s correlational, information, and structural quantities can now be calculated in closed form using new methods of  $\epsilon$ -machine spectral decomposition [91].

By way of explaining this paradigm, we opened with a review of stationary processes and their  $\epsilon$ -machines, turning to broaden the setting to joint input-output processes and communication channels. We then defined the (conditional) causal equivalence relation, which led immediately to transducer causal states and the  $\epsilon$ -transducer. A series of theorems then established their optimality. To illustrate the range of possible transformations we considered a systematic set of example channels that, in addition, provided an outline of a structural classification scheme. As an aide in this, we gave a graphical way to view structured transformations via causal-state wordmaps. With the framework developed, one sees that the same level of computational mechanics’ prior analysis of individual processes can now be brought to bear on understanding structural transformations between processes.

The foregoing, however, is simply the first in a series on the structural analysis of mappings between processes. The next will address the information-theoretic measures appropriate to joint input-output processes. We then will turn to an analysis that blends the present results on the causal architecture of structured transformations and the information-theoretic measures, showing how the internal mechanism expressed in the  $\epsilon$ -transducer supports information creation, loss, and manipulation during flow. From that point, the sequels will branch out to address channel composition, decomposition, and inversion.

Given the diversity of domains in which structured transformations (and their understanding) appear to play a role, there looks to be a wide range of applications. In addition to addressing several of these applications, Sec. XIV outlined several future research directions. The  $\epsilon$ -transducer development leads, for example, to a number of questions that can now be precisely posed and whose answers now seem in reach: How exactly do different measuring devices change the  $\epsilon$ -machine formed from measurements of a fixed system? What precisely is lost in the measurement process, and how well can we model a system using a given measuring device? When is it possible to see *past* a measuring device into a system, and how can we optimize our choice of measuring device in practice?

## ACKNOWLEDGMENTS

We thank Cina Aghamohammadi, Alec Boyd, David Darmon, Chris Ellison, Ryan James, John Mahoney, and Paul Riechers for helpful comments and the Santa Fe Institute for its hospitality during visits. JPC is an SFI External Faculty member. This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract numbers W911NF-12-1-0234, W911NF-13-1-0390, and W911NF-13-1-0340. NB was partially supported by NSF VIGRE grant DMS0636297.

## Appendix A: Equivalence of Two $\epsilon$ -Transducer Definitions

We show the equivalence of two different  $\epsilon$ -transducer definitions, that presented in the main paper and an earlier version requiring additional assumptions. Since the  $\epsilon$ -transducer is determined by its causal equivalence relation, we show that the respective equivalence relations are the same. The first is defined and discussed at length above and duplicated here for convenience.

**Definition 1.** *The causal equivalence relation  $\sim_\epsilon$  for channels is defined as follows:*

$$\begin{aligned} \overleftarrow{(x, y)} \sim_\epsilon \overleftarrow{(x, y)'} &\iff \\ \mathbb{P}(\overrightarrow{Y} | \overrightarrow{X}, \overleftarrow{(X, Y)} = \overleftarrow{(x, y)}) & \\ = \mathbb{P}(\overrightarrow{Y} | \overrightarrow{X}, \overleftarrow{(X, Y)} = \overleftarrow{(x, y)'}) & \end{aligned}$$

The second definition is an implicit equivalence relation consisting of an explicit equivalence relation, along with an additional unifilarity constraint that, of course, is quite strong [49, 50]. Here, we make both requirements explicit.

**Definition 2.** *The single-symbol unifilar equivalence relation  $\sim_\epsilon^1$  for channels is defined as follows:*

$$\begin{aligned} \overleftarrow{(x, y)} \sim_\epsilon^1 \overleftarrow{(x, y)'} &\iff \\ (i) \quad \mathbb{P}(Y_0 | X_0, \overleftarrow{(X, Y)}_0 = \overleftarrow{(x, y)}) & \\ = \mathbb{P}(Y_0 | X_0, \overleftarrow{(X, Y)}_0 = \overleftarrow{(x, y)'}) & \\ \text{and:} & \\ (ii) \quad \mathbb{P}(Y_1 | X_1, \overleftarrow{(X, Y)}_0 = \overleftarrow{(x, y)}, (X, Y)_0 = (a, b)) & \\ = \mathbb{P}(Y_1 | X_1, \overleftarrow{(X, Y)}_0 = \overleftarrow{(x, y)'}, (X, Y)_0 = (a, b)) & , \end{aligned}$$

for all  $a \in \mathcal{X}$  and  $b \in \mathcal{Y}$  such that:

$$\mathbb{P}((X, Y)_0 = (a, b) | \overleftarrow{(x, y)}) > 0$$

and:

$$\mathbb{P}((X, Y)_0 = (a, b) | \overleftarrow{(x, y)'}) > 0 .$$

The second requirement (ii) in the above definition requires that appending any joint symbol to two single-symbol-equivalent pasts will also result in a pair of pasts that are single-symbol-equivalent. This is unifilarity. The second part of the second requirement ensures that we are only considering *possible* joint symbols  $(a, b)$ —symbols that can follow  $\overleftarrow{(x, y)}$  or  $\overleftarrow{(x, y)'}$  with some nonzero probability.

**Proposition 7.** *The single-symbol unifilar equivalence relation is identical to the causal equivalence relation.*

*Proof.* Let  $\overleftarrow{(x, y)}$  and  $\overleftarrow{(x, y)'}$  be two pasts, equivalent under  $\sim_\epsilon^1$ . This provides our base case for induction:

$$\mathbb{P}(Y_0 | X_0, \overleftarrow{(x, y)}) = \mathbb{P}(Y_0 | X_0, \overleftarrow{(x, y)'}) . \quad (\text{A1})$$

Now, let's assume that  $\overleftarrow{(x, y)}$  and  $\overleftarrow{(x, y)'}$  are equivalent for length- $L - 1$  future morphs:

$$\mathbb{P}(Y_{0:L} | X_{0:L}, \overleftarrow{(x, y)}) = \mathbb{P}(Y_{0:L} | X_{0:L}, \overleftarrow{(x, y)'}) . \quad (\text{A2})$$

We need to show that  $\overleftarrow{(x, y)}$  and  $\overleftarrow{(x, y)'}$  are equivalent for length- $L$  future morphs by using the unifilarity constraint. Unifilarity requires that appending a joint symbol  $(a, b)$  to both  $\overleftarrow{(x, y)}$  and  $\overleftarrow{(x, y)'}$  results in two new pasts also equivalent to each other for length- $L - 1$  future morphs:

$$\begin{aligned} \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)}(a, b)) & \\ = \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)'}(a, b)) . & \end{aligned} \quad (\text{A3})$$

Since this must be true for *any* joint symbol, we replace  $(a, b)$  with  $(X, Y)_0$  in Eq. (A3), giving:

$$\begin{aligned} \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)}, (X, Y)_0) & \\ = \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)'}, (X, Y)_0) & \\ \iff & \\ \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)}, X_0, Y_0) & \\ = \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0, Y_0) . & \end{aligned} \quad (\text{A4})$$

To arrive at our result, we need to multiply the left side of Eq. (A4) by  $\mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)}, X_0)$  and the right side by  $\mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0)$ , which we can do when these quantities are equal. Since our channel is *causal*,  $X_{1:L+1}$  has no effect on  $Y_0$  when we condition on the infinite joint past and present input symbol. The two quantities,  $\mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)}, X_0)$  and  $\mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0)$ , therefore reduce to  $\mathbb{P}(Y_0 | \overleftarrow{(x, y)}, X_0)$  and  $\mathbb{P}(Y_0 | \overleftarrow{(x, y)'}, X_0)$ , respectively. But these are equal by the single-symbol unifilar equivalence relation—the base for induction. Multiplying each side of Eq. (A4) by these two terms (in their original form) gives:

$$\begin{aligned} \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)}, X_0, Y_0) & \\ \cdot \mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)}, X_0) & \\ = \mathbb{P}(Y_{1:L+1} | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0, Y_0) & \\ \cdot \mathbb{P}(Y_0 | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0) & \\ \iff & \\ \mathbb{P}(Y_{1:L+1}, Y_0 | X_{1:L+1}, \overleftarrow{(x, y)}, X_0) & \\ = \mathbb{P}(Y_{1:L+1}, Y_0 | X_{1:L+1}, \overleftarrow{(x, y)'}, X_0) & \\ \iff & \\ \mathbb{P}(Y_{0:L+1} | X_{0:L+1}, \overleftarrow{(x, y)}) = \mathbb{P}(Y_{0:L+1} | X_{0:L+1}, \overleftarrow{(x, y)'}) . & \end{aligned}$$

The two pasts are therefore equivalent for length- $L$  future morphs. By induction, the two pasts are equivalent for arbitrarily long future morphs.  $\square$

- [1] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(January):17–24, 2012.
- [2] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Lett.*, 45:712, 1980.
- [3] D. Mandal and C. Jarzynski. Work and information processing in a solvable model of Maxwell’s Demon. *Proc. Natl. Acad. Sci. USA*, 109(29):11641–11645, 2012.
- [4] A. B. Boyd and J. P. Crutchfield. Demon dynamics: Deterministic chaos, the Szilard map, and the intelligence of thermodynamic systems. 2015. SFI Working Paper 15-06-019; arxiv.org:1506.04327 [cond-mat.stat-mech].
- [5] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Identifying functional thermodynamics in autonomous Maxwellian ratchets. In preparation, 2015.
- [6] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. Bradford Book, New York, 1999.
- [7] V. Cutsuridis, A. Hussain, and J. G. Taylor. *Perception-Action Cycle*. Springer, New York, New York, 2011.
- [8] G. Gordon, D. Kaplan, D. Lankow, D. Little, J. Sherwin, B. Suter, et al. Toward an integrated approach to perception and action: Conference report and future directions. *Front. Syst. Neurosci.*, 5(20), 2011.
- [9] J. F. Padgett, D. Lee, and N. Collier. Economic production as chemistry. *Industrial and Corporate Change*, 12(4):843–877, 2003.
- [10] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63:105–108, 1989.
- [11] R. M. Gray. *Probability, Random Processes, and Ergodic Theory*. Springer-Verlag, New York, second edition, 2009.
- [12] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006.
- [13] R. W. Yeung. *Information Theory and Network Coding*. Springer, New York, 2008.
- [14] C. J. Ellison, J. R. Mahoney, and J. P. Crutchfield. Prediction, retrodiction, and the amount of information stored in the present. *J. Stat. Phys.*, 136(6):1005–1034, 2009.
- [15] J. P. Crutchfield, C. J. Ellison, J. R. Mahoney, and R. G. James. Synchronization and control in intrinsic and designed computation: An information-theoretic analysis of competing models of stochastic computation. *CHAOS*, 20(3):037105, 2010.
- [16] R. M. Gray and L. D. Davisson. The ergodic decomposition of stationary discrete random processes. *IEEE Trans. Info. Th.*, 20(5):625–636, 1974.
- [17] P. Billingsley. Statistical methods in Markov chains. *Ann. Math. Stat.*, 32:12, 1961.
- [18] J. P. Crutchfield and C. R. Shalizi. Thermodynamic depth of causal states: Objective complexity via minimal representations. *Phys. Rev. E*, 59(1):275–283, 1999.
- [19] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001.
- [20] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney. Time’s barbed arrow: Irreversibility, crypticity, and stored information. *Phys. Rev. Lett.*, 103(9):094101, 2009.
- [21] N. Travers and J. P. Crutchfield. Equivalence of history and generator  $\epsilon$ -machines. 2011. SFI Working Paper 11-11-051; arxiv.org:1111.4500 [math.PR].
- [22] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.
- [23] A. Paz. *Introduction to Probabilistic Automata*. Academic Press, New York, 1971.
- [24] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Reading, 1979.
- [25] D. Lind and B. Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, New York, 1995.
- [26] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.
- [27] Though see Ref. [73, see Ch. 7] and for early efforts Refs. [74] and [75].
- [28] R. M. Gray. *Entropy and Information Theory*. Springer-Verlag, New York, 1990.
- [29] A. V. Oppenheim and R. W. Schaffer. *Discrete-time signal processing*. Prentice-Hall signal processing series. Prentice Hall, 2010.
- [30] B. Mandelbrot. *Fractals: Form, Chance and Dimension*. W. H. Freeman and Company, New York, 1977.
- [31] L. R. Rabiner. A tutorial on hidden Markov models and selected applications. *IEEE Proc.*, 77:257, 1989.
- [32] R. J. Elliot, L. Aggoun, and J. B. Moore. *Hidden Markov Models: Estimation and Control*, volume 29 of *Applications of Mathematics*. Springer, New York, 1995.
- [33] R. G. James, J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Many roads to synchrony: Natural time scales and their algorithms. *Phys. Rev. E*, 89:042135, 2014.
- [34] J. P. Crutchfield and J. E. Hanson. Turbulent pattern bases for cellular automata. *Physica D*, 69:279 – 301, 1993.
- [35] C.-B. Li, H. Yang, and T. Komatsuzaki. Multiscale complex network of protein conformational fluctuations in single-molecule time series. *Proc. Natl. Acad. Sci. USA*, 105:536–541, 2008.
- [36] D. P. Varn, G. S. Canright, and J. P. Crutchfield.  $\epsilon$ -Machine spectral reconstruction theory: A direct method for inferring planar disorder and structure from X-ray diffraction studies. *Acta. Cryst. Sec. A*, 69(2):197–206, 2013.
- [37] C. C. Streliaoff and J. P. Crutchfield. Bayesian structural inference for hidden processes. *Phys. Rev. E*, 89:042119, 2014.
- [38] J. P. Crutchfield and B. S. McNamara. Equations of motion from a data series. *Complex Systems*, 1:417–452, 1987.
- [39] B. D. Johnson, J. P. Crutchfield, C. J. Ellison, and C. S. McTague. Enumerating finitary processes. 2012. SFI Working Paper 10-11-027; arxiv.org:1011.0036 [cs.FL].
- [40] C. R. Shalizi, K. L. Shalizi, and J. P. Crutchfield. Pattern discovery in time series, Part I: Theory, algorithm, analysis, and convergence. *J. Machine Learning Research*, submitted, 2002. Santa Fe Institute Working Paper 02-10-060; arXiv.org/abs/cs.LG/0210025.
- [41] S. Still, J. P. Crutchfield, and C. J. Ellison. Optimal causal inference: Estimating stored information and ap-

- proximating causal architecture. *CHAOS*, 20(3):037111, 2010.
- [42] J. P. Crutchfield. Reconstructing language hierarchies. In H. A. Atmanspracher and H. Scheingraber, editors, *Information Dynamics*, pages 45–60, New York, 1991. Plenum.
- [43] J. E. Hanson and J. P. Crutchfield. The attractor-basin portrait of a cellular automaton. *J. Stat. Phys.*, 66:1415–1462, 1992.
- [44] J. P. Crutchfield. Discovering coherent structures in nonlinear spatial systems. In A. Brandt, S. Ramberg, and M. Shlesinger, editors, *Nonlinear Ocean Waves*, pages 190–216, Singapore, 1992. World Scientific.
- [45] J. E. Hanson and J. P. Crutchfield. Computational mechanics of cellular automata: An example. *Physica D*, 103:169–189, 1997.
- [46] C. S. McTague and J. P. Crutchfield. Automated pattern discovery—An algorithm for constructing optimally synchronizing multi-regular language filters. *Theo. Comp. Sci.*, 359(1-3):306–328, 2006.
- [47] J. P. Crutchfield and O. Görnerup. Objects that make objects: The population dynamics of structural complexity. *J. Roy. Soc. Interface*, 3:345–349, 2006.
- [48] J. P. Crutchfield and S. Whalen. Structural drift: The population dynamics of sequential learning. *PLoS Computational Biology*, 8(6):e1002510, 2010.
- [49] J. P. Crutchfield. Optimal structural transformations—the  $\epsilon$ -transducer. *UC Berkeley Physics Research Report*, 1994.
- [50] C. R. Shalizi. *Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata*. PhD thesis, University of Wisconsin, Madison, Wisconsin, 2001.
- [51] J. J. Hopfield. Neural networks and physical systems with emergent collective behavior. *Proc. Natl. Acad. Sci.*, 79:2554, 1982.
- [52] J. Hertz, A. Krogh, and R. G. Palmer. *An Introduction to the Theory of Neural Networks*, volume 1 of *Lecture Notes, Studies in the Sciences of Complexity*. Addison-Wesley, Redwood City, California, 1991.
- [53] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, second edition, 2011.
- [54] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge, Cambridge, United Kingdom, 2003.
- [55] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. *Neural Computation*, 13:2409–2463, 2001.
- [56] C. R. Shalizi and J. P. Crutchfield. Information bottlenecks, causal states, and statistical relevance bases: How to represent relevant information in memoryless transduction. *Adv. Comp. Sys.*, 5(1):91–95, 2002.
- [57] D. Albers, J. C. Sprott, and J. P. Crutchfield. Persistent chaos in high dimensions. *Phys. Rev. E*, 74(5):057201, 2006.
- [58] S. Marzen and J. P. Crutchfield. Informational and causal architecture of discrete-time renewal processes. *Entropy*, to appear, 2015. SFI Working Paper 14-08-032; arxiv.org:1408.6876 [cond-mat.stat-mech].
- [59] J. P. Crutchfield. Information and its metric. In L. Lam and H. C. Morris, editors, *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, pages 119–130, New York, 1990. Springer-Verlag.
- [60] C. E. Shannon. The lattice theory of information. *IEEE Trans. Info. Th.*, 1:105–107, 1953.
- [61] H. Li and E. K. P. Chong. On a connection between information and group lattices. *Entropy*, 13:683–798, 2011.
- [62] M. Morse and G. A. Hedlund. Symbolic dynamics. *Am. J. Math.*, 60(4):815–866, 1938.
- [63] E. Post. Introduction to the general theory of elementary propositions. *Am. J. Math.*, 43:163–185, 1921.
- [64] K. Gödel. *On Formally Undecidable Propositions of Principia Mathematica and Related Systems*. Dover Publications, 1992.
- [65] A. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.*, 42, 43:230–265, 544–546, 1937.
- [66] A. Church. A note on the Entscheidungsproblem. *J. Symbolic Logic*, 1:40–41, 1936.
- [67] E. F. Moore. Gedanken-experiments on sequential machines. In C. Shannon and J. McCarthy, editors, *Automata Studies*, number 34 in *Annals of Mathematical Studies*, pages 129–153. Princeton University Press, Princeton, New Jersey, 1956.
- [68] D. Huffman. The synthesis of sequential switching circuits. *J. Franklin Inst.*, 257:161–190, 275–303, 1954.
- [69] D. Huffman. Information conservation and sequence transducers. In *Proc. Symp. Information Networks*, pages 291–307. Polytechnic Institute of Brooklyn, Brooklyn, New York, 1954.
- [70] D. Huffman. Canonical forms for information-lossless finite-state logical machines. *IRE Trans. Circ. Th.*, 6:41–59, 1959.
- [71] D. Huffman. Notes on information-lossless finite-state automata. *Il Nuovo Cimento*, 13(2 Supplement):397–405, 1959.
- [72] J. G. Brookshear. *Theory of computation: Formal languages, automata, and complexity*. Benjamin/Cummings, Redwood City, California, 1989.
- [73] R. B. Ash. *Information Theory*. John Wiley and Sons, New York, 1965.
- [74] D. Blackwell, L. Breiman, and A. J. Thomasian. Proof of Shannon’s transmission theorem for finite-state indecomposable channels. *Ann. Math. Stat.*, 29(4):1209–1220, 1958.
- [75] D. Blackwell. Exponential error bounds for finite state channels. In *Proc. Fourth Berkeley Symp. on Math. Statist. and Prob.*, volume 1, pages 57–63. Univ. of Calif. Press, 1961.
- [76] R. W. Hamming. *Digital Filters*. Dover Publications, New York, third edition, 1997.
- [77] G. Bateson. *Mind and Nature: A Necessary Unity*. E. P. Dutton, New York, 1979.
- [78] R. Shaw. Strange attractors, chaotic behavior, and information flow. *Z. Naturforsch.*, 36a:80, 1981.
- [79] R. Ahlswede and J. Körner. Appendix: On common information and related characteristics of correlated information sources. In R. Ahlswede, Baumer, N. Cai, H. Aydinian, V. Blinovskiy, C. Deppe, and H. Mashurian, editors, *General Theory of Information Transfer and Combinatorics*, volume 4123 of *Lecture Notes in Computer Science*, pages 664–677. Springer Berlin, 2006.
- [80] P. L. Williams and R. D. Beer. Generalized measures of information transfer. *arXiv:1102.1507*, abs/1102.1507, 2011.
- [81] L. Barnett, A. B. Barrett, and A. K. Seth. Granger causality and transfer entropy are equivalent for Gaus-

- sian variables. *Phys. Rev. Lett.*, 103, 2009.
- [82] T. Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85:461–464, 2000.
- [83] H. Marko. The bidirectional communication theory—A generalization of information theory. *IEEE Trans. Comm.*, COM-21:1345–1351, 1973.
- [84] J. L. Massey. Causality, feedback and directed information. In *Proc. 1990 Intl. Symp. on Info. Th. and its Applications, Waikiki, Hawaii, Nov. 27-30*, pages 1–6. 1990.
- [85] J. Sun and E. M. Bollt. Causation entropy identifies indirect influences, dominance of neighbors and anticipatory couplings. *Physica D*, 267:49–57, 2014.
- [86] N. Barnett and J. P. Crutchfield. Computational mechanics of input-output processes: Shannon information measures and decompositions. *in preparation*, 2014.
- [87] V. V. Fedorov. *Theory of Optimal Experiments*. Probability and Mathematical Statistics. Academic Press, New York, 1972.
- [88] A. Atkinson, B. Bogacka, and A. A. Zhigljavsky, editors. *Optimum Design 2000*. Nonconvex Optimization and Its Applications. Springer, New York, 2001.
- [89] S. Still. Information-theoretic approach to interactive learning. *EuroPhys. Lett.*, 85:28005, 2009.
- [90] K. J. Astrom and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, Princeton, New Jersey, 2008.
- [91] J. P. Crutchfield, P. Riechers, and C. J. Ellison. Exact complexity: Spectral decomposition of intrinsic computation. Santa Fe Institute Working Paper 13-09-028; arXiv:1309.3792 [cond- mat.stat-mech].