

# Spectral Simplicity of Apparent Complexity, Part I: The Nondiagonalizable Metadynamics of Prediction

Paul M. Riechers\* and James P. Crutchfield†

*Complexity Sciences Center  
Department of Physics  
University of California at Davis  
One Shields Avenue, Davis, CA 95616*

(Dated: January 12, 2018)

Virtually all questions that one can ask about the behavioral and structural complexity of a stochastic process reduce to a linear algebraic framing of a time evolution governed by an appropriate hidden-Markov process generator. Each type of question—correlation, predictability, predictive cost, observer synchronization, and the like—induces a distinct generator class. Answers are then functions of the class-appropriate transition dynamic. Unfortunately, these dynamics are generically nonnormal, nondiagonalizable, singular, and so on. Tractably analyzing these dynamics relies on adapting the recently introduced *meromorphic functional calculus*, which specifies the spectral decomposition of functions of nondiagonalizable linear operators, even when the function poles and zeros coincide with the operator’s spectrum. Along the way, we establish special properties of the spectral projection operators that demonstrate how they capture the organization of subprocesses within a complex system. Circumventing the spurious infinities of alternative calculi, this leads in the sequel, Part II, to the first closed-form expressions for complexity measures, couched either in terms of the Drazin inverse (negative-one power of a singular operator) or the eigenvalues and projection operators of the appropriate transition dynamic.

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

Keywords: hidden Markov model, entropy rate, excess entropy, predictable information, statistical complexity, projection operator, complex analysis, resolvent, Drazin inverse

For well over a century, science compared the randomness in physical systems via their temperatures or their thermodynamic entropies. These are measures of energy disorder. Using them, we say one or the other system is more random (hotter or more entropic) than the other. Curiously, even today, we do not know how to compare two physical systems in terms of how organized they are. Such comparisons are particularly important when the systems of interest do not have well defined energies—as found across mathematics and the sciences, from abstract dynamical systems to economic and social systems. This is what the endeavor of exploring complexity measures addresses—develop quantities that allow one to compare how nonlinear systems are structured, how they store and process information, and how they intrinsically compute. To date, complexity measures have been estimated empirically from experimental measurements, from large-scale simulations that generate synthetic data, or theoretically in the very few cases that are analytically tractable. We show that this arduous and limited state of affairs is no

longer necessary, if one can theoretically deduce or empirically estimate a statistical representation called a hidden Markov model. We provide analytic, closed-form expressions for almost all complexity measures of processes generated by hidden Markov models.

## I. INTRODUCTION

Complex systems—that is, many-body systems with strong interactions—are usually observed through low-resolution feature detectors. The consequence is that their hidden structure is, at best, only revealed over time. Since individual observations cannot capture the full resolution of each degree of freedom, let alone a sufficiently full set of them, measurement time series often appear stochastic and non-Markovian, exhibiting long-range correlations. Empirical challenges aside, restricting to the purely theoretical domain, even finite systems can appear quite complicated. Despite admitting finite descriptions, stochastic processes with sofic support, to take one example, exhibit infinite-range dependencies among the chain of random variables they generate [1]. While such infinite-correlation processes are legion in complex physical and biological systems, even approximately analyzing them is generally appreciated as difficult, if not impossible. Generically, even finite systems lead to uncountably

---

\* pmriechers@ucdavis.edu

† chaos@ucdavis.edu

infinite sets of predictive features [2]. These facts seem to put physical sciences’ most basic goal—prediction—out of reach.

We aim to show that this direct, but sobering conclusion is too bleak. Rather, there is a collection of constructive methods that address hidden structure and the challenges associated with predicting complex systems. This follows up on our recent introduction of a functional calculus that uncovered new relationships among supposedly different complexity measures [3] and that demonstrated the need for a generalized spectral theory to answer such questions [4]. Those efforts yielded elegant, closed-form solutions for complexity measures that, when compared, offered insight into the overall theory of complexity measures. Here, providing the necessary background for and greatly expanding those results, we show that different questions regarding correlation, predictability, and prediction each require their own analytical structures, expressed as various kinds of hidden transition dynamic. The resulting transition dynamic among hidden variables summarizes symmetry breaking, synchronization, and information processing, for example. Each of these metadynamics, though, is built up from the original given system.

The shift in perspective that allows the new level of tractability begins by recognizing that—beyond their ability to *generate* many sophisticated processes of interest—hidden Markov models can be treated as exact *mathematical objects* when analyzing the processes they generate. Crucially, and especially when addressing nonlinear processes, most questions that we ask imply a linear transition dynamic over *some* hidden state space. Speaking simply, something happens, then it evolves linearly in time, then we snapshot a selected characteristic. This broad type of sequential questioning *cascades*, in the sense that the influence of the initial preparation cascades through state space as time evolves, affecting the final measurement. Alternatively, other, complementary kinds of questioning involve *accumulating* such cascades. The linear algebra underlying either kind is highlighted in Table I in terms of an appropriate discrete-time transition operator  $T$  or a continuous-time generator  $G$  of time evolution.

In this way, deploying linear algebra to analyze complex systems turns on identifying an appropriate hidden state space. And, in turn, the latter depends on the genre of the question. Here, we focus on closed-form expressions for a process’ complexity measures. This determines what the internal system setup  $\langle \cdot |$  and the final detection  $|\cdot \rangle$  should be. We show that complexity questions fall into three subgenres and, for each of these, we identify the appropriate linear dynamic and closed-form expressions for several of the key questions in each genre.

Linear Algebra Underlying Complexity		
Question type	Discrete time	Continuous time
Cascading	$\langle \cdot   T^L   \cdot \rangle$	$\langle \cdot   e^{tG}   \cdot \rangle$
Accumulating	$\langle \cdot   (\sum_L T^L)   \cdot \rangle$	$\langle \cdot   (\int e^{tG} dt)   \cdot \rangle$

TABLE I. Having identified the hidden linear dynamic, either a discrete-time operator  $T$  or continuous-time operator  $G$ , quantitative questions tend to be either *cascading* or *accumulating* type. What changes between distinct questions are the dot products with the initial setup  $\langle \cdot |$  and the final observations  $|\cdot \rangle$ .

See Table II. The burden of the following is to explain the table in detail. We return to a much-elaborated version at the end.

Associating observables  $x \in \mathcal{A}$  with transitions between hidden states  $s \in \mathcal{S}$ , gives a hidden Markov model (HMM) with observation-labeled transition matrices  $\{T^{(x)} : T_{i,j}^{(x)} = \Pr(x, s_j | s_i)\}_{x \in \mathcal{A}}$ . They sum to the row-stochastic state-to-state transition matrix  $T = \sum_{x \in \mathcal{A}} T^{(x)}$ . (The continuous-time versions are similarly defined, which we do later on.) Adding measurement symbols  $x \in \mathcal{A}$  this way—to transitions—can be considered a model of measurement itself [5]. The efficacy of our choice will become clear.

It is important to note that HMMs, in continuous and discrete time, arise broadly in the sciences, from quantum mechanics [6, 7], statistical mechanics [8], and stochastic thermodynamics [9–11] to communication theory [12, 13], information processing [14–16], computer design [17], population and evolutionary dynamics [18, 19], and economics. Thus, HMMs appear in the most fundamental physics and in the most applied engineering and social sciences. The breadth suggests that the thoroughgoing HMM analysis developed here is worth the required effort to learn.

Since complex processes have highly structured, directional transition dynamics— $T$  or  $G$ —we encounter the full richness of matrix algebra in analyzing HMMs. We explain how analyzing complex systems *induces* a nondiagonalizable metadynamics, even if the original dynamic is diagonalizable in its underlying state-space. Normal and diagonalizable restrictions, so familiar in mathematical physics, simply fail us here.

The diversity of nondiagonalizable dynamics presents a technical challenge, though. A new calculus for functions of nondiagonalizable operators—e.g.,  $T^L$  or  $e^{tG}$ —becomes a necessity if one’s goal is an exact analysis of complex processes. Moreover, complexity measures naively and easily lead one to consider illegal operations. Taking the inverse of a singular operator is a particularly central, useful, and fraught example. Fortunately,

such illegal operations can be skirted since the complexity measures only extract the excess transient behavior of an infinitely complicated orbit space.

To explain how this arises—how certain modes of behavior, such as excess transients, are selected as relevant, while others are ignored—we apply the meromorphic functional calculus and new results for spectral projection operators recently derived in Ref. [4] to analyze complex processes generated by HMMs.

The following shows that this leads to a simplified spectral theory of weighted directed graphs, that even non-diagonalizable eigenspaces can be manipulated individually, and that, more specifically, the techniques can be applied to the challenges of prediction. The results developed here greatly extend and (finally) explain those announced in Ref. [3]. The latter introduced the basic methods and results by narrowly focusing on closed-form expressions for several measures of intrinsic computation, applying them to prototype complex systems.

The meromorphic functional calculus, summarized in detail later, concerns functions of nondiagonalizable operators when poles (or zeros) of the function of interest coincide with poles of the operator’s resolvent—poles that appear precisely at the eigenvalues of the transition dynamics. Pole–pole and pole–zero interactions transform the complex-analysis residues within the functional calculus. One notable result is that the negative-one power of a singular operator exists in the meromorphic functional calculus. We derive its form, note that it is the *Drazin inverse*, and show how widely useful and common it is.

For example, the following gives the first closed-form expressions for many complexity measures in wide use—many of which turn out to be expressed most concisely in terms of a Drazin inverse. Furthermore, spectral decomposition gives insight into subprocesses of a complex system in terms of the spectral projection operators of the appropriate transition dynamic.

In the following, we emphasize that when we observe processes generated by a source capable of even the simplest computations, much of the predictable structure lies beyond pairwise correlation. We clarify how different measures of complexity quantify and distinguish nuanced aspects of what is predictable and what is necessary for prediction. We then give closed-form solutions for this quantification, resulting in a new level of rigor, tractability, and insight.

To get started, sections §II through §III briefly review relevant background in stochastic processes, the HMMs that generate them, and complexity measures. Several classes of HMMs are discussed in §III. Mixed-state presentations (MSPs)—HMM generators of a process that also track distributions induced by observation—are re-

Questions and Their Linear Dynamics		
Genre	Measures	Hidden dynamic
Observation	Correlations $\gamma(L)$ Power spectra $P(\omega)$	HMM matrix $T$
Predictability	Myopic entropy $h_\mu(L)$ Excess entropy $\mathbf{E}, \mathcal{E}(\omega)$	HMM MSP matrix $W$
Prediction	Causal synchrony $C_\mu, \mathcal{H}^+(L)$ $\mathbf{S}, \mathcal{S}(\omega)$	$\epsilon$ -Machine MSP matrix $\mathcal{W}$
Generation	State synchrony $C(\mathcal{M}), \mathcal{H}(L), \mathbf{S}'$	Generator MSP matrix

TABLE II. Question genres (leftmost column) about process complexity listed with increasing sophistication. Each genre implies a different linear transition dynamic (rightmost column). *Observational* questions concern the superficial, given dynamic. *Predictability* questions are about the observation-induced dynamic over distributions; that is, over states used to generate the superficial dynamic. *Prediction* questions address the dynamic over distributions over a process’ causally-equivalent histories. *Generation* questions concern the dynamic over any nonunifilar presentation  $\mathcal{M}$  and observation-induced dynamics over its distributions.

viewed in §IV. They are key to complexity measures within an information-theoretic framing. Section §V then shows how each complexity measure reduces to the linear algebra of an appropriate HMM adapted to the question genre.

To make progress at this point, we summarize the meromorphic functional calculus in §VI. Several of its mathematical implications are discussed in relation to projection operators in §VII and a spectral weighted directed graph theory is presented in §VIII.

With this all set out, the sequel Part II finally derives the promised closed-form complexities of a process and outlines common simplifications for special cases. This leads to the discovery of the *symmetry collapse index*, which indicates the sophistication of finite computational structures hidden in infinite-Markov-order processes. Leveraging the functional calculus, Part II introduces a novel extension—the complexity measure frequency spectrum and shows how to calculate it in closed form. It provides a suite of examples to ground the theoretical developments and works through in-depth a pedagogical example.

## II. STRUCTURED PROCESSES AND THEIR COMPLEXITIES

We first describe a system of interest in terms of its observed behavior, following the approach of computational mechanics, as reviewed in Ref. [20]. Again, a *process* is the collection of behaviors that the system produces and their probabilities of occurring.

A process's behaviors are described via a bi-infinite chain of random variables, denoted by capital letters  $\dots X_{t-2} X_{t-1} X_t X_{t+1} X_{t+2} \dots$ . A realization is indicated by lowercase letters  $\dots x_{t-2} x_{t-1} x_t x_{t+1} x_{t+2} \dots$ . We assume values  $x_t$  belong to a discrete alphabet  $\mathcal{A}$ . We work with blocks  $X_{t:t'}$ , where the first index is inclusive and the second exclusive:  $X_{t:t'} = X_t \dots X_{t'-1}$ . Block realizations  $x_{t:t'}$  we often refer to as *words*  $w$ . At each time  $t$ , we can speak of the *past*  $X_{-\infty:t} = \dots X_{t-2} X_{t-1}$  and the *future*  $X_{t:\infty} = X_t X_{t+1} \dots$ .

A process's probabilistic specification is a density over these chains:  $\mathbb{P}(X_{-\infty:\infty})$ . Practically, we work with finite blocks and their probability distributions  $\Pr(X_{t:t'})$ . To simplify the development, we primarily analyze stationary, ergodic processes: those for which  $\Pr(X_{t:t+L}) = \Pr(X_{0:L})$  for all  $t \in \mathbb{Z}$ ,  $L \in \mathbb{Z}^+$ . In such cases, we only need to consider a process's length- $L$  *word distributions*  $\Pr(X_{0:L})$ .

### A. Directly observable organization

A common first step to understand how processes express themselves is to analyze correlations among observables. Pairwise correlation in a sequence of observables is often summarized by the *autocorrelation function*:

$$\gamma(L) = \langle \overline{X}_t X_{t+L} \rangle_t ,$$

where the bar above  $X_t$  denotes its complex conjugate, and the angled brackets denote an average over all times  $t \in \mathbb{Z}$ . Alternatively, structure in a stochastic process is often summarized by the *power spectral density*, also referred to more simply as the *power spectrum*:

$$P(\omega) = \lim_{N \rightarrow \infty} \frac{1}{N} \left\langle \left| \sum_{L=1}^N X_L e^{-i\omega L} \right|^2 \right\rangle ,$$

where  $\omega \in \mathbb{R}$  is the angular frequency [21]. Though a basic fact, it is not always sufficiently emphasized in applications that power spectra capture only pairwise correlation. Indeed, it is straightforward to show that the power spectrum  $P(\omega)$  is the windowed Fourier transform of the autocorrelation function  $\gamma(L)$ . That is, power spectra describe how pairwise correlations are distributed across frequencies. Power spectra are common in signal processing, both in technological settings and physical experiments [22]. As a physical example, diffraction patterns are the power spectra of a sequence of structure factors [23].

Other important measures of observable organization called Green–Kubo coefficients determine transport properties in near-equilibrium thermodynamic systems—

but are rather more application-specific [24, 25]. These coefficients reflect the idea that dissipation depends on correlation structure. They usually appear in the form of integrating the autocorrelation of *derivatives* of observables. A change of observables, however, turns this into an integration of a standard autocorrelation function. Green–Kubo transport coefficients then involve the limit  $\lim_{\omega \rightarrow 0} P(\omega)$  for the process of appropriate observables.

One theme in the following is that, though widely used, correlation functions and power spectra give an impoverished view of a process's structural complexity, since they only consider ensemble averages over pairwise events. Moreover, creating a list of higher-order correlations is an impractical way to summarize complexity, as seen in the *connected correlation functions* of statistical mechanics [26].

### B. Intrinsic predictability

Information measures, in contrast, can involve all orders of correlation and thus help to go beyond pairwise correlation in understanding, for example, how a process' past behavior affects predicting it at later times. Information theory, as developed for general complex processes [1], provides a suite of quantities that capture prediction properties using variants of Shannon's entropy  $H[\cdot]$  and mutual information  $I[\cdot; \cdot]$  [13] applied to sequences. Each measure answers a specific question about a process' predictability. For example:

- How much information is contained in the words generated? The *block entropy* [1]:

$$H(L) = - \sum_{w \in \mathcal{A}^L} \Pr(w) \log_2 \Pr(w) .$$

- How random is a process? Its *entropy rate* [12]:

$$h_\mu = \lim_{L \rightarrow \infty} H(L)/L .$$

For dynamical systems with a continuous phase-space  $\mathcal{B}$ , the *metric entropy* also known as *Kolmogorov–Sinai (KS) entropy* is the supremum of entropy rates induced by partitioning  $\mathcal{B}$  into different finite alphabets  $\mathcal{A}$ . [27].

- How is the irreducible randomness  $h_\mu$  approached? Via the *myopic entropy rates* [28]:

$$h_\mu(L) = H[X_0 | X_{1-L} \dots X_{-1}] .$$

- How much of the future can be predicted? Its *excess entropy*, which is the past–future mutual information [1, and references therein]:

$$\mathbf{E} = \mathbb{I}[X_{-\infty:0}; X_{0:\infty}] .$$

$\mathbf{E}$  has also been investigated in ergodic theory [29] and under the names *stored information* [30], *effective measure complexity* [31], and *predictive information* [32].

- How much information must be extracted to know its predictability and so see its intrinsic randomness  $h_\mu$ ? Its *transient information* [1]:

$$\mathbf{T} = \sum_{L=0}^{\infty} [\mathbf{E} + h_\mu L - \mathbb{H}(L)] .$$

The spectral approach, our subject, naturally leads to allied, but new information measures. To give a sense, later we introduce the *excess entropy spectrum*  $\mathcal{E}(\omega)$ . It completely, yet concisely, summarizes the structure of myopic entropy reduction, in a way similar to how the power spectrum completely describes autocorrelation. However, while the power spectrum summarizes only pairwise linear correlation, the excess entropy spectrum captures all orders of nonlinear dependency between random variables, making it an incisive probe of hidden structure.

Before leaving the measures related to predictability, we must also point out that they have important refinements—measures that lend a particularly useful, even functional, interpretation. These include the bound, ephemeral, elusive, and related informations [33, 34]. Though amenable to the spectral methods of the following, we leave their discussion for another venue. Fortunately, their spectral development is straightforward, but would take us beyond the minimum necessary presentation to make good on the overall discussion of spectral decomposition.

### C. Prediction overhead

Process predictability measures, as just enumerated, certainly say much about a process’ intrinsic information processing. They leave open, though, the question of the structural complexity associated with *implementing* prediction. This challenge entails a complementary set of measures that directly address the inherent complexity of actually *predicting* what is *predictable*. For that matter, how cryptic is a process?

Computational mechanics describes minimal-memory maximal prediction—using the minimal memory nec-

essary to predict everything that is predictable about the future—via a process’ hidden, effective or causal states and transitions, as summarized by the process’s  $\epsilon$ -machine [20]. A *causal state*  $\sigma \in \mathcal{S}^+$  is an equivalence class of histories  $x_{-\infty:0}$  that all yield the same probability distribution over observable futures  $X_{0:\infty}$ . Therefore, knowing a process’s current causal state—that  $\mathcal{S}_0^+ = \sigma$ , say—is sufficient for maximal prediction.

The computational mechanics framework can also be related to several more recent attempts at describing effective levels of complex systems. For example, if individual histories are taken to be the microstates of a stochastic process, then causal states are the minimal high-level description of a stochastic process that satisfies the *informational closure* criterion of Ref. [35].

Computational mechanics provides an additional suite of quantities that capture the overhead of prediction, again using variants of Shannon’s entropy and mutual information applied to the  $\epsilon$ -machine. Each also answers a specific question about an observer’s burden of prediction. For example:

- How much historical information must be stored for maximal prediction? The Shannon information in the causal states or *statistical complexity* [36]:

$$C_\mu = \mathbb{H}[\mathcal{S}_0^+] .$$

- How unpredictable is a causal state upon observing a process for duration  $L$ ? The *myopic causal-state uncertainty* [1]:

$$\mathcal{H}^+(L) = \mathbb{H}[\mathcal{S}_0^+ | X_{-L} \dots X_{-1}] .$$

- How much information must an observer extract to synchronize to—that is, to know with certainty—the causal state? The optimal predictor’s *synchronization information* [1]:

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

Paralleling the purely informational suite of the previous section, we later introduce the *optimal synchronization spectrum*  $\mathcal{S}(\omega)$ . It completely and concisely summarizes the frequency distribution of state-uncertainty reduction, similar to how the power spectrum  $P(\omega)$  completely describes autocorrelation and the excess entropy spectrum  $\mathcal{E}(\omega)$  the myopic entropy reduction. Helpfully, the above optimal prediction measures can be found from the optimal synchronization spectrum.

The structural complexities monitor an observer’s burden in optimally predicting a process. And so, they have practical relevance when an intelligent artificial or biolog-

ical agent must take advantage of a structured stochastic environment—e.g., a Maxwellian Demon taking advantage of correlated environmental fluctuations [37], prey avoiding easy prediction, or profiting from stock market volatility, come to mind.

Prediction has many natural generalizations. For example, since maximal prediction often requires infinite resources, sub-maximal prediction (i.e., predicting with lower fidelity) is of practical interest. Fortunately, there are principled ways to investigate the tradeoffs between predictive accuracy and computational burden [2, 38–40]. As another example, maximal prediction in the presence of noisy or irregular observations can be investigated with a properly generalized framework; see Ref. [41]. Blending the existing tools, resource-limited prediction under such observational constraints can also be investigated. There are also many applications where prediction is relevant to the task at hand, but is not necessarily the ultimate objective; this of course has a long history, and Ref. [42] has recently tried to formalize this effort. In all of these settings, information measures similar to those listed above are key to understanding and quantifying the tradeoffs arising in prediction.

Having highlighted the difference between prediction and predictability, we can appreciate that some processes hide more internal information—are more cryptic—than others. It turns out, this can be quantified. The *crypticity*  $\chi = C_\mu - \mathbf{E}$  is the difference between the process’s stored information  $C_\mu$  and the mutual information  $\mathbf{E}$  shared between past and future observables [43]. Operationally, crypticity contrasts predictable information content  $\mathbf{E}$  with an observer’s minimal stored-memory overhead  $C_\mu$  required to *make* predictions. To predict what is predictable, therefore, an optimal predictor must account for a process’s crypticity.

#### D. Generative complexities

How does a physical system produce its output process? This depends on many details. Some systems employ vast internal mechanistic redundancy, while others under constraints have optimized internal resources down to a minimally necessary generative structure. Different pressures give rise to different kinds of optimality. For example, minimal state-entropy generators turn out to be distinct from minimal state-set generators [44–46]. The challenge then is to develop ways to monitor differences in generative mechanism [47].

Any generative model [1, 48]  $\mathcal{M}$  with state-set  $\mathcal{S}$  has a statistical state complexity (state entropy):  $C(\mathcal{M}) = \mathbf{H}[\mathcal{S}]$ . Consider the corresponding *myopic state-*

*uncertainty* given  $L$  sequential observations:

$$\mathcal{H}(L) = \mathbf{H}[\mathcal{S}_0 | X_{-L:0}] ,$$

And so:

$$\mathcal{H}(0) = C(\mathcal{M}) .$$

We also have the asymptotic uncertainty  $\mathcal{H} \equiv \lim_{L \rightarrow \infty} \mathcal{H}(L)$ . Related, there is the *excess synchronization information*:

$$\mathbf{S}' = \sum_{L=0}^{\infty} [\mathcal{H}(L) - \mathcal{H}] .$$

Such quantities are relevant even when an observer never fully synchronizes to a generative state; i.e., even when  $\mathcal{H} > 0$ . Finite-state  $\epsilon$ -machines always synchronize [49, 50] and so their  $\mathcal{H}$  vanishes.

Since many different mechanisms can generate a given process, we need useful bounds on the statistical state complexity of possible process generators. For example, the minimal *generative complexity*  $C_g = \min_{\{\mathcal{S}\}} C(\mathcal{M})$  is the minimal state-information a physical system must store to generate its future [46]. The predictability and the statistical complexities bound each other:

$$\mathbf{E} \leq C_g \leq C_\mu .$$

That is, the *predictable* future information  $\mathbf{E}$  is less than or equal to the information  $C_g$  necessary to *produce* the future which, in turn, is less than or equal to the information  $C_\mu$  necessary to *predict* the future [1, 44–47]. Such relationships have been explored even for quantum generators of (classical) stochastic processes [51, and references therein].

### III. HIDDEN MARKOV MODELS

Up to this point, the development focused on introducing and interpreting various information and complexity measures. It was not constructive in that there was no specification of how to calculate these quantities for a given process. To do so requires models or, in the vernacular, a *presentation* of a process. Fortunately, a common mathematical representation describes a wide class of process generators: the edge-labeled *hidden Markov models* (HMMs), also known as a Mealy HMMs [48, 52] [53]. Using these as our preferred presentations, we will first classify them and then describe how to calculate the information measures of the processes they generate.

**Definition 1.** A finite-state, edge-labeled hidden Markov model  $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \{T^{(x)}\}_{x \in \mathcal{A}}, \eta_0\}$  consists of:

- A finite set of hidden states  $\mathcal{S} = \{s_1, \dots, s_M\}$ .  $\mathcal{S}_t$  is the random variable for the hidden state at time  $t$ .
- A finite output alphabet  $\mathcal{A}$ .
- A set of  $M \times M$  symbol-labeled transition matrices  $\{T^{(x)}\}_{x \in \mathcal{A}}$ , where  $T_{i,j}^{(x)} = \Pr(x, s_j | s_i)$  is the probability of transitioning from state  $s_i$  to state  $s_j$  and emitting symbol  $x$ . The corresponding overall state-to-state transition matrix is the row-stochastic matrix  $T = \sum_{x \in \mathcal{A}} T^{(x)}$ .
- An initial distribution over hidden states:  $\eta_0 = (\Pr(\mathcal{S}_0 = s_1), \Pr(\mathcal{S}_0 = s_2), \dots, \Pr(\mathcal{S}_0 = s_M))$ .

The dynamics of such finite-state models are governed by transition matrices amenable to the linear algebra of vector spaces. As a result, bra-ket notation is useful. Bras  $\langle \cdot |$  are row vectors and kets  $|\cdot\rangle$  are column vectors. One benefit of the notation is immediately recognizing mathematical object type. For example, on the one hand, any expression that forms a closed bra-ket pair—either  $\langle \cdot | \cdot \rangle$  or  $\langle \cdot | \cdot | \cdot \rangle$ —is a scalar quantity and commutes as a unit with anything. On the other hand, when useful, an expression of the ket-bra form  $|\cdot\rangle \langle \cdot|$  can be interpreted as a matrix.

$T$ 's row-stochasticity means that each of its rows sum to unity. Introducing  $|\mathbf{1}\rangle$  as the column vector of all 1s, this can be restated as:

$$T |\mathbf{1}\rangle = |\mathbf{1}\rangle .$$

This is readily recognized as an eigenequation:  $T |\lambda\rangle = \lambda |\lambda\rangle$ . That is, the all-ones vector  $|\mathbf{1}\rangle$  is always a right eigenvector of  $T$  associated with the eigenvalue  $\lambda$  of unity.

When the internal Markov transition matrix  $T$  is irreducible, the Perron-Frobenius theorem guarantees that there is a unique asymptotic state distribution  $\pi$  determined by:

$$\langle \pi | T = \langle \pi | ,$$

with the further condition that  $\pi$  is normalized in probability:  $\langle \pi | \mathbf{1}\rangle = 1$ . This again is recognized as an eigenequation: the asymptotic distribution  $\pi$  over the hidden states is  $T$ 's left eigenvector associated with the eigenvalue of unity.

To describe a stationary process, as done often in the following, the initial hidden-state distribution  $\eta_0$  is set to the asymptotic one:  $\eta_0 = \pi$ . The resulting process generated is then stationary. Choosing an alternative  $\eta_0$  is useful in many contexts. For example, starting with  $\eta_t$  the state distribution at the next time is simply  $\langle \eta_{t+1} | T = \langle \eta_t |$ . However, starting with such alternatives typically yields a nonstationary process.

An HMM  $\mathcal{M}$  describes a process' behaviors as a formal language  $\mathcal{L} \subseteq \bigcup_{\ell=1}^{\infty} \mathcal{A}^{\ell}$  of allowed realizations. More-

over,  $\mathcal{M}$  succinctly describes a process's word distribution  $\Pr(w)$  over all words  $w \in \mathcal{L}$ . (Appropriately,  $\mathcal{M}$  also assigns zero probability to words outside of the process' language:  $\Pr(w) = 0$  for all  $w \in \mathcal{L}^c$ ,  $\mathcal{L}$ 's complement.) Specifically, the stationary probability of observing a particular length- $L$  word  $w = x_0 x_1 \dots x_{L-1}$  is given by:

$$\Pr(w) = \langle \pi | T^{(w)} | \mathbf{1}\rangle , \quad (1)$$

where  $T^{(w)} \equiv T^{(x_0)} T^{(x_1)} \dots T^{(x_{L-1})}$ .

More generally, given a nonstationary state distribution  $\eta$ , the probability of seeing  $w$  is:

$$\Pr(X_{t:t+L} = w | \mathcal{S}_t \sim \eta) = \langle \eta | T^{(w)} | \mathbf{1}\rangle , \quad (2)$$

where  $\mathcal{S}_t \sim \eta$  means that the random variable  $\mathcal{S}_t$  is distributed as  $\eta$  [13]. And, the state distribution having seen word  $w$  starting in  $\eta$  is:

$$\Pr(\mathcal{S}_{t+L} | \mathcal{S}_t \sim \eta, X_{t:t+L} = w) = \frac{\langle \eta | T^{(w)} | \mathbf{1}\rangle}{\langle \eta | T^{(w)} | \mathbf{1}\rangle} . \quad (3)$$

These conditional distributions are used often since, for example, most observations induce a nonstationary distribution over hidden states. Tracking such observation-induced distributions is the role of a related model class—the mixed-state presentation, introduced shortly. To get there, we must first introduce several, prerequisite HMM classes. See Fig. 1. A simple example of the general HMM just discussed is shown in Fig. 1a.

## A. Unifilar HMMs

An important class of HMMs consists of those that are unifilar. *Unifilarity* guarantees that, given a start state and a sequence of observations, there is a unique path through the internal states [54]. This, in turn, allows one to directly translate between properties of the internal Markov chain and properties of the observed behavior generated along the sequence of edges traversed. The states of unifilar HMMs are maximally predictive [55].

In contrast, general—that is, nonunifilar—HMMs have an exponentially growing number of possible state paths as a function of observed word length. Thus, nonunifilar process presentations break most all quantitative connections between internal dynamics and observations, rendering them markedly less useful process presentations. While they can be used to generate realizations of a given process, they cannot be used outright to predict a process. Unifilarity is required.

**Definition 2.** A finite-state, edge-labeled, unifilar HMM (*uHMM*) [56] is a finite-state, edge-labeled HMM with the following property:

- **Unifilarity:** For each state  $s \in \mathcal{S}$  and each symbol  $x \in \mathcal{A}$  there is at most one outgoing edge from state  $s$  that emits symbol  $x$ .

An example is shown in Fig. 1b.

### B. Minimal unifilar HMMs

Minimal models are not only convenient to use, but very often allow for determining essential informational properties, such as a process' memory  $C_\mu$ . A process' minimal state-entropy uHMM is the same as its minimal-state uHMM. And, the latter turns out to be the process'  $\epsilon$ -machine in computational mechanics [20]. Computational mechanics shows how to calculate a process'  $\epsilon$ -machine from the process' conditional word distributions. Specifically,  $\epsilon$ -machine states, the process' *causal states*  $\sigma \in \mathcal{S}^+$ , are equivalence classes of histories that yield the same predictions for the future. Explicitly, two histories  $\overleftarrow{x}$  and  $\overleftarrow{x}'$  map to the same causal state  $\epsilon(\overleftarrow{x}) = \epsilon(\overleftarrow{x}') = \sigma$  if and only if  $\Pr(\overrightarrow{X}|\overleftarrow{x}) = \Pr(\overrightarrow{X}|\overleftarrow{x}')$ . Thus, each causal state comes with a prediction of the future  $\Pr(\overrightarrow{X}|\sigma)$ —its *future morph*. In short, a process'  $\epsilon$ -machine is its minimal size, maximally-predictive predictor.

Converting a given uHMM to its corresponding  $\epsilon$ -machine employs probabilistic variants of well-known state-minimization algorithms in automata theory [57]. One can also verify that a given uHMM is minimal by checking that all its states are probabilistically distinct [49, 50].

**Definition 3.** A uHMM's states are probabilistically distinct if for each pair of distinct states  $s_k, s_j \in \mathcal{S}$  there exists some finite word  $w = x_0x_1 \dots x_{L-1}$  such that:

$$\Pr(\overrightarrow{X} = w | \mathcal{S} = s_k) \neq \Pr(\overrightarrow{X} = w | \mathcal{S} = s_j) .$$

If this is the case, then the process' uHMM is its  $\epsilon$ -machine.

An example is shown in Fig. 1c.

### C. Finitary stochastic process hierarchy

The finite-state presentations in these classes form a hierarchy in terms of the processes they can finitely generate [44]:  $\text{Processes}(\epsilon\text{-machines}) = \text{Processes}(\text{uHMMs}) \subset \text{Processes}(\text{HMMs})$ . That is, finite HMMs generate a strictly larger class of stochastic processes than finite uHMMs. The class of processes generated by finite uHMMs, though, is the same as generated by finite  $\epsilon$ -machines.

### D. Continuous-time HMMs

Though we concentrate on discrete-time processes, many of the process classifications, properties, and calculational methods carry over easily to continuous time. In this setting transition *rates* are more appropriate than transition probabilities. Continuous-time HMMs can often be obtained as a discrete-time limit  $\Delta t \rightarrow 0$  of an edge-labeled HMM whose edges operate for a time  $\Delta t$ . The most natural continuous-time HMM presentation, though, has a continuous-time generator  $G$  of time evolution over hidden states, with observables emitted as deterministic *functions of an internal Markov chain*:  $f : \mathcal{S} \rightarrow \mathcal{A}$ .

Respecting the continuous-time analog of probability conservation, each row of  $G$  sums to zero. Over a finite time interval  $t$ , marginalizing over all possible observations, the row-stochastic state-to-state transition dynamic is:

$$T_{t_0 \rightarrow t_0+t} = e^{tG} .$$

Any nontrivial continuous-time process generated by such a continuous-time HMM has an uncountably infinite number of possible realizations within a finite time interval, and most of these have vanishing probability. However, probabilities regarding what state the system is in at any finite set of times can be easily calculated. For this purpose we introduce the continuous-time *observation matrices*:

$$\Gamma_x = \sum_{\rho \in \mathcal{R}} \delta_{x,f(\rho)} |\delta_\rho\rangle \langle \delta_\rho| ,$$

where  $\delta_{x,f(\rho)}$  is a Kronecker delta,  $|\delta_\rho\rangle$  the column vector of all 0s except for a 1 at the position for state  $\rho$ , and  $\langle \delta_\rho|$  its transpose  $(|\delta_\rho\rangle)^\top$ . These ‘‘projectors’’ sum to the identity:  $\sum_{x \in \mathcal{A}} \Gamma_x = I$ .

An example is shown in Fig. 1d.

### IV. MIXED-STATE PRESENTATIONS

A given process can be generated by nonunifilar, unifilar, and  $\epsilon$ -machine HMM presentations. A process'  $\epsilon$ -machine is unique. However, within either the unifilar or nonunifilar HMM classes, there can be an unbounded number of presentations that generate the process.

This flexibility suggests that we can create a HMM process generator that, through its embellished structure, answers more refined questions than the information generation ( $h_\mu$ ) and memory ( $C_\mu$ ) calculated from the  $\epsilon$ -machine. To this end, we introduce the *mixed-state presentation* (MSP). An MSP tracks important supple-






FIG. 1. Finite HMM classes and example processes they generate, depicted by their state-transition diagrams: For any setting of the transition probabilities  $p, q \in (0, 1)$  and transition rates  $a, b, c \in (0, \infty)$ , each HMM generates an observable stochastic process over its alphabet  $\mathcal{A} \subset \{0, 1, 2\}$ —the latent states themselves are not directly observable from the output process and so are “hidden”. (a) Simple nonunifilar source: Two transitions leaving from the same state generate the same output symbol. (b) Nonminimal unifilar HMM. (c)  $\epsilon$ -Machine: Minimal unifilar HMM for the stochastic process generated. (d) Generator of a continuous-time stochastic process.

mentary information *in* the hidden states and, through well-crafted dynamics, *over* the hidden states. In particular, an MSP generates a process while tracking the observation-induced distribution over the states of an alternative process generator. Here, we review only that subset of mixed-state theory required by the following.

Consider a HMM presentation  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \{T^{(x)}\}_{x \in \mathcal{A}}, \pi)$  of some process in statistical equilibrium. A mixed state  $\eta$  can be any state distribution over  $\mathcal{S}$  and so we work with the simplex  $\Delta_{\mathcal{S}}$  of state distributions. However,  $\Delta_{\mathcal{S}}$  is uncountable and so contains far more mixed states than needed to calculate many complexity measures.

How to efficiently monitor the way in which an observer comes to know the HMM state as it sees successive symbols from the process? This is the problem of *observer-state synchronization*. To analyze the evolution of the observer’s knowledge through a sequence of observation-induced mixed states, we use the set  $\mathcal{R}_{\pi}$  of mixed states that are induced by all allowed words  $w \in \mathcal{L}$  from initial mixed state  $\eta_0 = \pi$ :

$$\mathcal{R}_{\pi} = \bigcup_{w \in \mathcal{L}} \left\{ \frac{\langle \pi | T^{(w)} \rangle}{\langle \pi | T^{(w)} | \mathbf{1} \rangle} \right\} \cup \{ \langle \pi | \} .$$

The cardinality of  $\mathcal{R}_{\pi}$  is finite when there are only a finite number of distinct probability distributions over  $\mathcal{M}$ ’s states that can be induced by observed sequences, if starting from the stationary distribution  $\pi$ .

If  $w$  is the first (in lexicographic order) word that induces a particular distribution over  $\mathcal{S}$ , then we denote this distribution as  $\eta^w$ ; a shorthand for Eq. (3). For example, if the two words 010 and 110110 both induce the same distribution  $\eta$  over  $\mathcal{S}$  and no word shorter than 010 induces that distribution, then the mixed state is denoted

$\eta^{010}$ . It corresponds to the distribution:

$$\langle \eta^{010} | = \frac{\langle \pi | T^{(0)} T^{(1)} T^{(0)} \rangle}{\langle \pi | T^{(0)} T^{(1)} T^{(0)} | \mathbf{1} \rangle} .$$

Since a given observed symbol induces a unique updated distribution from a previous distribution, the dynamic over mixed states is unifilar. Transition probabilities among mixed states can be obtained via Eqs. (2) and (3). So, if:

$$\langle \eta | T^{(x)} | \mathbf{1} \rangle > 0$$

and:

$$\langle \eta' | = \frac{\langle \eta | T^{(x)} \rangle}{\langle \eta | T^{(x)} | \mathbf{1} \rangle} ,$$

then:

$$\begin{aligned} \Pr(\eta', x | \eta) &= \Pr(x | \eta) \\ &= \langle \eta | T^{(x)} | \mathbf{1} \rangle . \end{aligned}$$

These transition probabilities over the mixed states in  $\mathcal{R}_{\pi}$  are the matrix elements for the observation-labeled transition matrices  $\{W^{(x)}\}_{x \in \mathcal{A}}$  of  $\mathcal{M}$ ’s *synchronizing MSP* ( $\mathcal{S}$ -MSP):

$$\mathcal{S}\text{-MSP}(\mathcal{M}) = (\mathcal{R}_{\pi}, \mathcal{A}, \{W^{(x)}\}_{x \in \mathcal{A}}, \delta_{\pi}) ,$$

where  $\delta_{\pi}$  is the distribution over  $\mathcal{R}_{\pi}$  peaked at the unique start-(mixed)-state  $\pi$ . The row-stochastic net mixed-state-to-state transition matrix of  $\mathcal{S}$ -MSP( $\mathcal{M}$ ) is  $W = \sum_{x \in \mathcal{A}} W^{(x)}$ . If irreducible, then there is a unique stationary probability distribution  $\langle \pi_W |$  over  $\mathcal{S}$ -MSP( $\mathcal{M}$ )’s states obtained by solving  $\langle \pi_W | = \langle \pi_W | W$ . We use  $\mathcal{R}_t$  to

denote the random variable for the MSP’s state at time  $t$ .

Figure 2 illustrates how a  $\mathcal{S}$ -MSP relates to the HMM whose state distributions it tracks. Both the original HMM and its MSP generate the same process. However, the MSP’s dynamic over mixed-states also tracks how, through the mixed states induced by successive observations, an observer comes to know the original HMM’s current state. Importantly, this MSP’s dynamic is nondiagonalizable, while the original HMM’s dynamic is diagonalizable. The appearance of nondiagonalizability when analyzing particular properties, in this case observer synchronization, is generic. Constructively working with this unavoidable fact motivates much of the following. Crucially, we find that the burden of predicting a stochastic process is fundamentally dependent on the nondiagonalizable characteristics of its MSP.

That feedforward network structures lead to nondiagonalizability and that it is important to network functionality was also observed in a quite different setting—oscillator synchronization on complex networks with directional coupling dynamics [58]. Briefly comparing settings reveals the origins of the commonality. In each, nondiagonalizability arises from an irreducible interdependence among elements—an interdependence that can be harnessed for hierarchical control. In Ref. [58], nondiagonalizable network structures allow upstream oscillators to influence downstream oscillators and this enables optimal synchronization among all oscillators. In contrast, while our setting does not concern synchronizing oscillator nodes to each other, it does analyze how an observer’s belief state synchronizes to the true state of the system under study. During this kind of synchronization, past states of knowledge feed into future states of knowledge. In short, nondiagonalizability corresponds to intrinsically interdependent updates in the evolution of knowledge.

More generally, we work with distributions over  $\Delta_{\mathcal{S}}$ . We must consider a mixed-state dynamic that starts from a general (nonpeaked) distribution over  $\Delta_{\mathcal{S}}$ . This may be counterintuitive, since a distribution over distributions should correspond to a single distribution. However, general MSP theory with a general starting distribution over  $\Delta_{\mathcal{S}}$  allows us to consider a weighted average of behaviors originating from different histories. And, this is distinct from considering the behavior originating from a weighted average of histories. This more general MSP formalism arises in the closed-form solutions for more sophisticated complexity measures, such as the bound information. This anticipates tools needed in a sequel.

With this brief overview of mixed states, we can now turn to use them. Section § V shows that tracking dis-

tributions over the states of another generator makes the MSP an ideal algebraic object for closed-form complexity expressions involving conditional entropies—measures that require conditional probabilities. Sections § II B and § II C showed that many of the complexity measures for predictability and predictive burden are indeed framed as conditional entropies. And so, MSPs are central to their closed-form expressions.

Historically, mixed states were already implicit in Ref. [59], introduced in their modern form by Ref. [44, 45], and have been used recently; e.g., in Refs. [60, 61]. Most of these efforts, however, used mixed-states in the specific context of the *synchronizing MSP* ( $\mathcal{S}$ -MSP). A greatly extended development of mixed-state dynamics appears in Ref. [41].

The overall strategy, though, is easy to explain. Different information-theoretic questions require different mixed-state dynamics, each of which is a unifilar presentation. Employing the mathematical methods developed here, we find that the desired closed-form solutions are often simple functions of the transition dynamic of an appropriate MSP. Specifically, the spectral properties of the relevant MSP control the form of information-theoretic quantities.

Finally, we note that similar linear-algebraic constructions—whose hidden states track relevant information—that are nevertheless *not* MSPs are just as important for answering different sets of questions about a process. Since these constructions are not directly about predictability and prediction, we report on them elsewhere.

## V. IDENTIFYING THE HIDDEN LINEAR DYNAMIC

We are now in a position to identify the hidden linear dynamic appropriate to many of the questions that arise in complex systems—their observation, predictability, prediction, and generation, as outlined in Table II. In part, this section addresses a very practical need for specific calculations. In part, it also lays the foundations for further generalizations, to be discussed at the end. Identifying the linear dynamic means identifying the linear operator  $A$  such that a question of interest can be reformulated as either being of the cascading form  $\langle \cdot | A^n | \cdot \rangle$  or as an accumulation of such cascading events via  $\langle \cdot | (\sum_n A^n) | \cdot \rangle$ ; recall Table I. Helpfully, many well-known questions of complexity can be mapped to these archetypal forms. And so, we now proceed to uncover the hidden linear dynamics of the cascading questions approximately in the order they were introduced in § II.




FIG. 2. (a) Example HMM and (b) its  $\mathcal{S}$ -MSP: Both of these HMMs generate the (2-1)-GP-(2) Process considered in Part II. However,  $\mathcal{S}$ -MSP also tracks observation-induced states of knowledge about the example HMM’s distribution over internal states A, B, C, and D. The double-circle state in the MSP state-transition diagram denotes  $\mathcal{S}$ -MSP’s start-state. Green states are transient, whereas blue states are recurrent. The MSP’s mixed-state-to-state dynamic is nondiagonalizable—which is generic—even though the example HMM’s state-to-state dynamic is diagonalizable.

### A. Simple complexity from any presentation

For observable correlation, any HMM transition operator will do as the linear dynamic. We simply observe, let time (or space) evolve forward, and observe again. Let’s be concrete.

Recall the familiar autocorrelation function. For a discrete-domain process it is [62]:

$$\gamma(L) = \langle \bar{X}_t X_{t+L} \rangle_t ,$$

where  $L \in \mathbb{Z}$  and the bar denotes the complex conjugate. The autocorrelation function is symmetric about  $L = 0$ , so we can focus on  $L \geq 0$ . For  $L = 0$ , we simply have:

$$\begin{aligned} \langle \bar{X}_t X_t \rangle_t &= \sum_{x \in \mathcal{A}} |x|^2 \Pr(X_t = x) \\ &= \sum_{x \in \mathcal{A}} |x|^2 \langle \pi | T^{(x)} | \mathbf{1} \rangle . \end{aligned}$$

For  $L > 0$ , we have:

$$\begin{aligned} \gamma(L) &= \langle \bar{X}_t X_{t+L} \rangle_t \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \Pr(X_t = x, X_{t+L} = x') \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \Pr(x \underbrace{* \cdots *}_{L-1 \text{ *s}} x') \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \sum_{w \in \mathcal{A}^{L-1}} \Pr(x w x') . \end{aligned}$$

Each ‘\*’ above is a wildcard symbol denoting indifference to the particular symbol observed in its place. That is,

the \*s denote marginalizing over the intervening random variables. We develop the consequence of this, explicitly calculating [63] and finding:

$$\begin{aligned} \gamma(L) &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \sum_{w \in \mathcal{A}^{L-1}} \langle \pi | T^{(x)} T^{(w)} T^{(x')} | \mathbf{1} \rangle \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \langle \pi | T^{(x)} \left( \sum_{w \in \mathcal{A}^{L-1}} T^{(w)} \right) T^{(x')} | \mathbf{1} \rangle \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \langle \pi | T^{(x)} \underbrace{\left( \prod_{i=1}^{L-1} \sum_{x_i \in \mathcal{A}} T^{(x_i)} \right)}_{=T} T^{(x')} | \mathbf{1} \rangle \\ &= \sum_{x \in \mathcal{A}} \sum_{x' \in \mathcal{A}} \bar{x} x' \langle \pi | T^{(x)} T^{L-1} T^{(x')} | \mathbf{1} \rangle \\ &= \langle \pi | \left( \sum_{x \in \mathcal{A}} \bar{x} T^{(x)} \right) T^{L-1} \left( \sum_{x' \in \mathcal{A}} x' T^{(x')} \right) | \mathbf{1} \rangle . \end{aligned}$$

The result is the autocorrelation in cascading form  $\langle \cdot | T^L | \cdot \rangle$ , which can be made particularly transparent by subsuming time-independent factors on the left and right into the bras and kets. Let’s introduce the new row vector:

$$\langle \pi \bar{\mathcal{A}} | = \langle \pi | \left( \sum_{x \in \mathcal{A}} \bar{x} T^{(x)} \right)$$

and column vector:

$$| \mathcal{A} \mathbf{1} \rangle = \left( \sum_{x \in \mathcal{A}} x T^{(x)} \right) | \mathbf{1} \rangle .$$

Then, the autocorrelation function for nonzero integer  $L$

is simply:

$$\gamma(L) = \langle \pi \bar{\mathcal{A}} | T^{|L|-1} | \mathcal{A} \mathbf{1} \rangle . \quad (4)$$

Clearly, the autocorrelation function is a direct, albeit filtered, signature of iterates of the transition dynamic of any process presentation.

This result can easily be translated to the continuous-time setting. If the process is represented as a function of a Markov chain and we make the translation that:

$$\langle \pi \bar{\mathcal{A}} | = \langle \pi | \left( \sum_{x \in \mathcal{A}} \bar{x} \Gamma_x \right) \quad \text{and} \quad | \mathcal{A} \mathbf{1} \rangle = \left( \sum_{x \in \mathcal{A}} x \Gamma_x \right) | \mathbf{1} \rangle ,$$

then the autocorrelation function for any  $\tau \in \mathbb{R}$  is simply:

$$\gamma(\tau) = \langle \pi \bar{\mathcal{A}} | e^{|\tau|G} | \mathcal{A} \mathbf{1} \rangle , \quad (5)$$

where  $G$  is determined from  $T$  following §III D. Again, the autocorrelation function is a direct fingerprint of the transition dynamic over the hidden states.

The power spectrum is a modulated accumulation of the autocorrelation function. With some algebra, one can show that it is:

$$P(\omega) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{L=-N}^N (N - |L|) \gamma(L) e^{-i\omega L} .$$

Reference [62] shows that for discrete-domain processes the continuous part of the power spectrum is simply:

$$P_c(\omega) = \langle |x|^2 \rangle + 2 \operatorname{Re} \langle \pi \bar{\mathcal{A}} | (e^{i\omega I} - T)^{-1} | \mathcal{A} \mathbf{1} \rangle , \quad (6)$$

where  $\operatorname{Re}(\cdot)$  denotes the real part of its argument and  $I$  is the identity matrix. Similarly, for continuous-domain processes one has:

$$P_c(\omega) = 2 \operatorname{Re} \langle \pi \bar{\mathcal{A}} | (i\omega I - G)^{-1} | \mathcal{A} \mathbf{1} \rangle . \quad (7)$$

Although useful, these signatures of pairwise correlation are only first-order complexity measures. Common measures of complexity that include higher orders of correlation can also be written in the simple cascading and accumulating forms, but require a more careful choice of representation.

## B. Predictability from a presentation MSP

For example, any HMM presentation allows us to calculate using Eq. (1) a process's block entropy:

$$\mathbb{H}(L) = \mathbb{H}[X_0 X_1 \dots X_{L-1}] ,$$

but at a computational cost  $\mathcal{O}(|\mathcal{S}|^3 L |\mathcal{A}|^L)$  exponential in  $L$ , due to the exponentially growing number of words in  $\mathcal{L} \cap \mathcal{A}^L$ . Consequently, using a general HMM one can neither directly nor efficiently calculate many key complexity measures, including a process's entropy rate and excess entropy.

These limitations motivate using more specialized HMM classes. To take one example, it has been known for some time that a process' entropy rate  $h_\mu$  can be calculated directly from any of its unifilar presentations [54]. Another is that we can calculate the excess entropy directly from a process's uHMM forward  $\mathcal{S}^+$  and reverse  $\mathcal{S}^-$  states [60, 61]:  $\mathbf{E} = \mathbb{I}[\bar{X}; \bar{X}] = \mathbb{I}[\mathcal{S}^+; \mathcal{S}^-]$ .

However, efficient computation of *myopic* entropy rates  $h_\mu(L)$  remained elusive for some time, and we only recently found their closed-form expression [3]. The myopic entropy rates are important because they represent the apparent entropy rate of a process if it is modeled as a finite Markov order- $(L-1)$  process—a very common approximation. Crucially, the difference  $h_\mu(L) - h_\mu$  from the process' true entropy rate is the *surplus entropy rate* incurred by using an order- $L-1$  Markov approximation. Similarly, these surplus entropy rates lead directly to not only an apparent loss of predictability, but errors in inferred physical properties. These include overestimates of dissipation associated with the surplus entropy rate assigned to a physical thermodynamic system [37].

Unifilarity, it turns out, is not enough to calculate a process'  $h_\mu(L)$  directly. Rather, the  $\mathcal{S}$ -MSP of any process presentation is what is required. Let's now develop the closed-form expression for the myopic entropy rates, following Ref. [41].

The length- $L$  myopic entropy rate is the expected uncertainty in the  $L^{\text{th}}$  random variable  $X_{L-1}$ , given the preceding  $L-1$  random variables  $X_{0:L-1}$ :

$$\begin{aligned} h_\mu(L) &\equiv \mathbb{H}(L) - \mathbb{H}(L-1) \\ &= \mathbb{H}[X_{0:L} | \eta_0 = \pi] - \mathbb{H}[X_{0:L-1} | \eta_0 = \pi] \\ &= \mathbb{H}[X_{L-1}, X_{0:L-1} | \eta_0 = \pi] - \mathbb{H}[X_{0:L-1} | \eta_0 = \pi] \\ &= \mathbb{H}[X_{L-1} | X_{0:L-1}, \eta_0 = \pi] , \end{aligned} \quad (8)$$

where, in the second line, we explicitly give the condition  $\eta_0 = \pi$  specifying our ignorance of the initial state. That is, without making any observations we can only assume that the initial distribution  $\eta_0$  over  $\mathcal{M}$ 's states is the expected asymptotic distribution  $\pi$ . For a mixing ergodic process, for example, even if another distribution  $\eta_{-N} = \alpha$  was known in distant past, we still have  $\langle \eta_0 | = \langle \eta_{-N} | T^N \rightarrow \langle \pi |$ , as  $N \rightarrow \infty$ .

Assuming an initial probability distribution over  $\mathcal{M}$ 's states, a given observation sequence induces a particular sequence of updated state distributions. That is,

the  $\mathcal{S}$ -MSP( $\mathcal{M}$ ) is unifilar regardless of whether  $\mathcal{M}$  is unifilar or not. Or, in other words, given the  $\mathcal{S}$ -MSP's unique start state— $\mathcal{R}_0 = \pi$ —and a particular realization  $X_{0:L-1} = w^{L-1}$  of the last  $L-1$  random variables, we end up at the particular mixed state  $\mathcal{R}_{L-1} = \eta_{w^{L-1}} \in \mathcal{R}_\pi$ . Moreover, the entropy of the next observation is uniquely determined by  $\mathcal{M}$ 's state distribution, suggesting that Eq. (8) becomes:

$$\mathbb{H}[X_{L-1}|X_{0:L-1}, \eta_0 = \pi] = \mathbb{H}[X_{L-1}|\mathcal{R}_{L-1}, \mathcal{R}_0 = \pi] ,$$

as proven elsewhere [41]. Intuitively, conditioning on all of the past observation random variables is equivalent to conditioning on the random variable for the state distribution induced by particular observation sequences.

We can now recast Eq. (8) in terms of the  $\mathcal{S}$ -MSP, finding:

$$\begin{aligned} h_\mu(L) &= \mathbb{H}[X_{L-1}|\mathcal{R}_{L-1}, \mathcal{R}_0 = \pi] \\ &= \sum_{\eta \in \mathcal{R}_\pi} \Pr(\mathcal{R}_{L-1} = \eta | \mathcal{R}_0 = \pi) \mathbb{H}[X_{L-1}|\mathcal{R}_{L-1} = \eta] \\ &= \sum_{\eta \in \mathcal{R}_\pi} \langle \delta_\pi | W^{L-1} | \delta_\eta \rangle \\ &\quad \times - \sum_{x \in \mathcal{A}} \langle \delta_\eta | W^{(x)} | \mathbf{1} \rangle \log_2 \langle \delta_\eta | W^{(x)} | \mathbf{1} \rangle \\ &= \langle \delta_\pi | W^{L-1} | \mathbb{H}(W^{\mathcal{A}}) \rangle . \end{aligned}$$

Here:

$$|\mathbb{H}(W^{\mathcal{A}})\rangle \equiv - \sum_{\eta \in \mathcal{R}_\pi} |\delta_\eta\rangle \sum_{x \in \mathcal{A}} \langle \delta_\eta | W^{(x)} | \mathbf{1} \rangle \log_2 \langle \delta_\eta | W^{(x)} | \mathbf{1} \rangle$$

is simply the column vector whose  $i^{\text{th}}$  entry is the entropy of transitioning from the  $i^{\text{th}}$  state of  $\mathcal{S}$ -MSP. Critically,  $|\mathbb{H}(W^{\mathcal{A}})\rangle$  is independent of  $L$ .

Notice that taking the logarithm of the sum of the entries of the row vector  $\langle \delta_\eta | W^{(x)} \rangle$  via  $\langle \delta_\eta | W^{(x)} | \mathbf{1} \rangle$  is only permissible since  $\mathcal{S}$ -MSP's unifilarity guarantees that  $W^{(x)}$  has at most one nonzero entry per row. (We also use the familiar convention that  $0 \log_2 0 = 0$  [13].)

The result is a particularly compact and efficient expression for the length- $L$  myopic entropy rates:

$$h_\mu(L) = \langle \delta_\pi | W^{L-1} | \mathbb{H}(W^{\mathcal{A}}) \rangle . \quad (9)$$

Thus, all that is required is computing powers  $W^{L-1}$  of the MSP transition dynamic. The computational cost  $\mathcal{O}(L|\mathcal{R}_\pi|^3)$  is now only linear in  $L$ . Moreover,  $W$  is very sparse, especially so with a small alphabet  $\mathcal{A}$ . And, this means that the computational cost can be reduced even further via numerical optimization.

With  $h_\mu(L)$  in hand, the hierarchy of complexity measures that derive from it immediately follow, including

the entropy rate  $h_\mu$ , the excess entropy  $\mathbf{E}$ , and the transient information  $\mathbf{T}$  [1]. Specifically, we have:

$$\begin{aligned} h_\mu &= \lim_{L \rightarrow \infty} h_\mu(L) , \\ \mathbf{E} &= \sum_{L=1}^{\infty} [h_\mu(L) - h_\mu] , \text{ and} \\ \mathbf{T} &= \sum_{L=1}^{\infty} L [h_\mu(L) - h_\mu] . \end{aligned}$$

The sequel, Part II, discusses these in more detail, introducing their closed-form expressions. To prepare for this, we must first review the meromorphic functional calculus, which is needed for working with the above operators.

### C. Continuous time?

We saw that correlation measures are easily extended to the continuous-time domain via continuous-time HMMs. Entropy rates (since they are rates) and state entropy (since it depends only on the instantaneous distribution) also carry over rather easily to continuous time. Indeed, the former is well studied for chaotic systems [64] and the latter is exemplified by the thermodynamic entropy. Yet, other information-theoretic measures of information transduction are awkward when directly translated to continuous time. At least one approach has been taken recently towards understanding their structure [65–67], but more work is necessary.

### D. Synchronization from generator MSP

If a process' state-space is known, then the  $\mathcal{S}$ -MSP of the generating model allows one to track the observation-induced distributions over its states. This naturally leads to closed-form solutions to informational questions about how an observer comes to know, or how it *synchronizes* to, the system's states.

To monitor how an observer's knowledge of a process' internal state changes with increasing measurements we use the *myopic state uncertainty*  $\mathcal{H}(L) = \mathbb{H}[\mathcal{S}_0|X_{-L:0}]$  [1]. Expressing it in terms of the  $\mathcal{S}$ -MSP, one finds [41]:

$$\begin{aligned} \mathcal{H}(L) &= - \sum_{w \in \mathcal{A}^L} \Pr(w) \sum_{s \in \mathcal{S}} \Pr(s|w) \log_2 \Pr(s|w) \\ &= \sum_{\eta \in \mathcal{R}_\pi} \Pr(\mathcal{R}_L = \eta | \mathcal{R}_0 = \pi) \mathbb{H}[\eta] . \end{aligned}$$

Here,  $\mathbb{H}[\eta]$  is the presentation-state uncertainty specified

by the mixed state  $\eta$ :

$$\mathbf{H}[\eta] \equiv - \sum_{s \in \mathcal{S}} \langle \eta | \delta_s \rangle \log_2 \langle \eta | \delta_s \rangle , \quad (10)$$

where  $|\delta_s\rangle$  is the length- $|\mathcal{S}|$  column vector of all zeros except for a 1 at the appropriate index of the presentation-state  $s$ .

Continuing, we re-express  $\mathcal{H}(L)$  in terms of powers of the  $\mathcal{S}$ -MSP transition dynamic:

$$\begin{aligned} \mathcal{H}(L) &= \sum_{\eta \in \mathcal{R}_\pi} \langle \delta_\pi | W^L | \delta_\eta \rangle \mathbf{H}[\eta] \\ &= \langle \delta_\pi | W^L | \mathbf{H}[\eta] \rangle . \end{aligned} \quad (11)$$

Here, we defined:

$$|\mathbf{H}[\eta]\rangle \equiv \sum_{\eta \in \mathcal{R}_\pi} |\delta_\eta\rangle \mathbf{H}[\eta] ,$$

which is the  $L$ -independent length- $|\mathcal{R}_\pi|$  column vector whose entries are the appropriately indexed entropies of each mixed state.

The forms of Eqs. (9) and (11) demonstrate that  $h_\mu(L+1)$  and  $\mathcal{H}(L)$  differ only in the type of information being extracted after being evolved by the operator: observable entropy  $\mathbf{H}(W^A)$  or state entropy  $\mathbf{H}[\eta]$ , as implicated by their respective kets. Each of these entropies decreases as the distributions induced by longer observation sequences converge to the synchronized distribution. If synchronization is achieved, the distributions become  $\delta$ -functions on a single state and the associated entropies vanish.

Paralleling  $h_\mu(L)$ , there is a complementary hierarchy of complexity measures that are built up from functions of  $\mathcal{H}(L)$ . These include the *asymptotic state uncertainty*  $\mathcal{H}$  and *excess synchronization information*  $\mathbf{S}'$ , to mention only two:

$$\begin{aligned} \mathcal{H} &= \lim_{L \rightarrow \infty} \mathcal{H}(L) \text{ and} \\ \mathbf{S}' &= \sum_{L=0}^{\infty} [\mathcal{H}(L) - \mathcal{H}] . \end{aligned}$$

Compared to the  $h_\mu(L)$  family of measures,  $\mathcal{H}$  and  $\mathbf{S}'$  mirror the roles of  $h_\mu$  and  $\mathbf{E}$ , respectively.

The *model state-complexity*:

$$\begin{aligned} C(\mathcal{M}) &= \mathcal{H}(0) \\ &= \langle \delta_\pi | \mathbf{H}[\eta] \rangle \end{aligned}$$

also has an analog in the  $h_\mu(L)$  hierarchy—the process'

*alphabet complexity*:

$$\begin{aligned} \mathbf{H}[X_0] &= h_\mu(1) \\ &= \langle \delta_\pi | \mathbf{H}(W^A) \rangle . \end{aligned}$$

### E. Maximal prediction from $\epsilon$ -machine MSP

We just reviewed the linear underpinnings of synchronizing to *any* model of a process. However, the myopic state uncertainty of the  $\epsilon$ -machine has a distinguished role in determining the synchronization cost for maximally predicting a process, regardless of the presentation that generated it. Using the  $\epsilon$ -machine's  $\mathcal{S}$ -MSP, the  $\epsilon$ -machine myopic state uncertainty can be written in direct parallel to the myopic state uncertainty of any model:

$$\begin{aligned} \mathcal{H}^+(L) &= - \sum_{w \in \mathcal{A}^L} \Pr(w) \sum_{\sigma \in \mathcal{S}^+} \Pr(\sigma|w) \log_2 \Pr(\sigma|w) \\ &= \sum_{\eta \in \mathcal{R}_\pi^+} \Pr(\mathcal{R}_L = \eta | \mathcal{R}_0 = \pi) \mathbf{H}[\eta] \\ &= \langle \delta_\pi | \mathcal{W}^L | \mathbf{H}[\eta] \rangle . \end{aligned}$$

The script  $\mathcal{W}$  emphasizes that we are now specifically working with the state-to-state transition dynamic of the  $\epsilon$ -machine's MSP.

Paralleling  $\mathcal{H}(L)$ , an obvious hierarchy of complexity measures is built from functions of  $\mathcal{H}^+(L)$ . For example, the  $\epsilon$ -machine's state-complexity is the *statistical complexity*  $C_\mu = \mathcal{H}^+(0)$ . The information that must be obtained to synchronize to the causal state in order to maximally predict—the *causal synchronization information*—is given in terms of the  $\epsilon$ -machine's  $\mathcal{S}$ -MSP by  $\mathbf{S} = \sum_{L=0}^{\infty} \mathcal{H}^+(L)$ .

An important difference when using  $\epsilon$ -machine presentations is that they have zero asymptotic state uncertainty:

$$\mathcal{H}^+ = 0 .$$

Therefore,  $\mathbf{S} = \mathbf{S}'(\epsilon\text{-machine})$ . Moreover, we conjecture that  $\mathbf{S} = \min_{\mathcal{M}} \sum_{L=0}^{\infty} \mathcal{H}(L)$  for any presentation  $\mathcal{M}$  that generates the process, even if  $C_\mu \geq C_g$ .

### F. Beyond the MSP

Many of the complexity measures use a mixed-state presentation as the appropriate linear dynamic, with particular focus on the  $\mathcal{S}$ -MSP. However, we want to emphasize that this is more a reflection of questions that have become common. It does not indicate the general answer

that one expects in the broader approach to finding the hidden linear dynamic. Here, we give a brief overview for how other linear dynamics can appear for different types of complexity question. These have been uncovered recently and will be reported on in more detail in sequels.

First, we found the reverse-time *mixed-functional presentation* (MFP) of any forward-time generator. The MFP tracks the reverse-time dynamic over linear functionals  $|\eta\rangle$  of state distributions induced by reverse-time observations:

$$|\eta\rangle \in \mathcal{R}_1 = \left\{ \frac{T^{(w)}|\mathbf{1}\rangle}{\langle \pi | T^{(w)} | \mathbf{1} \rangle} \right\}_w .$$

The MFP allows direct calculation of the convergence of the *preparation uncertainty*  $\mathfrak{H}(L) \equiv \mathbb{H}(\mathcal{S}_0 | X_{0:L})$  via powers of the linear MFP transition dynamic. The preparation uncertainty in turn gives a new perspective on the transient information since:

$$\mathbf{T} = \sum_{L=0}^{\infty} (\mathbb{H}(\mathcal{S}_0^+ | X_{0:L}) - \chi)$$

can be interpreted as the predictive advantage of hindsight. Related, the *myopic process crypticity*  $\chi(L) = \mathfrak{H}^+(L) - \mathcal{H}^+(L)$  had been previously introduced [43]. Since  $\lim_{L \rightarrow \infty} \mathcal{H}^+(L) = \mathcal{H}^+ = 0$ , the asymptotic crypticity is  $\chi = \mathfrak{H}^+ + \mathcal{H}^+ = \mathfrak{H}^+$ . And, this reveals a refined partitioning underlying the sum:

$$\sum_{L=0}^{\infty} (\chi - \chi(L)) = \mathbf{S} - \mathbf{T} \geq 0 .$$

Crypticity  $\chi = \mathbb{H}(\mathcal{S}_0^+ | X_{0:\infty})$  itself is positive only if the process' *cryptic order*:

$$k = \min\{\ell \in \{0, 1, \dots\} : \mathbb{H}(\mathcal{S}_0^+ | X_{-\ell:\infty}) = 0\} ,$$

is positive. The cryptic order is always less than or equal to its better-known cousin, the *Markov order*  $R$ :

$$R = \min\{\ell \in \{0, 1, \dots\} : \mathbb{H}(\mathcal{S}_0^+ | X_{-\ell:0}) = 0\} ,$$

since conditioning can never increase entropy. In the case of cryptic order, we condition on future observations  $X_{0:\infty}$ .

The forward-time *cryptic operator presentation* gives the forward-time observation-induced dynamic over the operators:

$$\mathcal{O} \in \left\{ \frac{|s^-\rangle \langle \eta^w|}{\langle \eta^w | s^- \rangle} : s^- \in \mathcal{S}^-, \langle \eta^w | \in \mathcal{R}_\pi, \langle \eta^w | s^- \rangle > 0 \right\} .$$

Since the reverse causal state  $\mathcal{S}_0^-$  at time 0 is a linear combination of forward causal states [68, 69], this presentation allows new calculations of the convergence to crypticity that implicate  $\Pr(\mathcal{S}_0^+ | X_{-L:\infty})$ .

In fact, the cryptic operator presentation is a special case of the more general myopic bidirectional dynamic over operators :

$$\mathcal{O} \in \left\{ \frac{|\eta^{w'}\rangle \langle \eta^w|}{\langle \eta^w | \eta^{w'} \rangle} : \langle \eta^w | \in \mathcal{R}_\pi, |\eta^{w'}\rangle \in \mathcal{R}_1, \langle \eta^w | \eta^{w'} \rangle > 0 \right\}$$

induced by new observations of either the future or the past. This is key to understanding the interplay between forgetfulness and shortsightedness:  $\Pr(\mathcal{S}_0 | X_{-M:0}, X_{0:N})$ .

The list of these extensions continues. Detailed bounds on entropy-rate convergence are obtained from the transition dynamic of the so-called *possibility machine*, beyond the asymptotic result obtained in Ref. [50]. And, the importance of post-synchronized monitoring, as quantified by the information lost due to negligence over a duration  $\ell$ :

$$b_\mu(\ell) = \mathbb{I}(X_{0:\ell}; X_{\ell:\infty} | X_{-\infty:0}) ,$$

can be determined using yet another type of modified MSP.

These examples all find an exact solution via a theory parallel to that outlined in the following, but applied to the linear dynamic appropriate for the corresponding complexity question. Furthermore, they highlight the opportunity, enabled by the full meromorphic functional calculus [4], to ask and answer more nuanced and, thus, more probing questions about structure, predictability, and prediction.

## G. The end?

It would seem that we achieved our goal. We identified the appropriate transition dynamic for common complexity questions and, by some standard, gave formulae for their exact solution. In point of fact, the effort so far has all been in preparation. Although we set the framework up appropriately for linear analysis, closed-form expressions for the complexity measures still await the mathematical developments of the following sections. At the same time, at the level of qualitative understanding and scientific interpretation, so far we failed to answer the simple question:

- What range of possible behaviors do these complexity measures exhibit?

and the natural follow-up question:

- What mechanisms produce qualitatively different informational signatures?

The following section reviews the recently developed functional calculus that allows us to actually decompose arbitrary functions of the nondiagonalizable hidden dynamic to give conclusive answers to these fundamental questions [4]. We then analyze the range of possible behaviors and identify the internal mechanisms that give rise to qualitatively different contributions to complexity.

The investment in this and the succeeding sections allow Part II to express new closed-form solutions for many complexity measures beyond those achieved to date. In addition to obvious calculational advantages, this also gives new insights into possible behaviors of the complexity measures and, moreover, their unexpected similarities with each other. In many ways, the results shed new light on what we were (implicitly) probing with already-familiar complexity measures. Constructively, this suggests extending complexity magnitudes to complexity functions that succinctly capture the organization to all orders of correlation. Just as our intuition for pairwise correlation grows out of power spectra, so too these extensions unveil the workings of both a process' predictability and the burden of prediction for an observer.

## VI. SPECTRAL THEORY BEYOND THE SPECTRAL THEOREM

Here, we briefly review the spectral decomposition theory from Ref. [4] needed for working with nondiagonalizable linear operators. As will become clear, it goes significantly beyond the spectral theorem for normal operators. Although linear operator theory (especially as developed in the mathematical literature of functional analysis) already addresses nonnormal operators, it had not delivered the comparable machinery for a tractable spectral decomposition of nonnormal and nondiagonalizable operators. Reference [4] explored this topic and derived new relations that enable the practical analysis of nonnormal and nondiagonalizable systems—in which independent subprocesses (irreducible subspaces) can be directly manipulated.

### A. Spectral primer

We restrict our attention to operators that have at most a countably infinite spectrum. Such operators share many features with finite-dimensional square matrices. And so, we review several elementary but essential facts that are used extensively in the following.

Recall that if  $A$  is a finite-dimensional square matrix, then  $A$ 's spectrum is simply its set of eigenvalues:

$$\Lambda_A = \{ \lambda \in \mathbb{C} : \det(\lambda I - A) = 0 \} ,$$

where  $\det(\cdot)$  is the determinant of its argument.

For reference later, recall that the *algebraic multiplicity*  $a_\lambda$  of eigenvalue  $\lambda$  is the power of the term  $(z - \lambda)$  in the characteristic polynomial  $\det(zI - A)$ . In contrast, the *geometric multiplicity*  $g_\lambda$  is the dimension of the kernel of the transformation  $A - \lambda I$  or the number of linearly independent eigenvectors for the eigenvalue. Moreover,  $g_\lambda$  is the number of Jordan blocks associated with  $\lambda$ . The algebraic and geometric multiplicities are all equal when the matrix is diagonalizable.

Since there can be multiple subspaces associated with a single eigenvalue, corresponding to different Jordan blocks in the Jordan canonical form, it is structurally important to introduce the index of the eigenvalue to describe the size of its largest-dimension associated subspace.

**Definition 4.** *The index  $\nu_\lambda$  of eigenvalue  $\lambda$  is the size of the largest Jordan block associated with  $\lambda$ .*

The index gives information beyond what the algebraic and geometric multiplicities themselves reveal. Nevertheless, for  $\lambda \in \Lambda_A$ , it is always true that  $\nu_\lambda - 1 \leq a_\lambda - g_\lambda \leq a_\lambda - 1$ . In the diagonalizable case,  $a_\lambda = g_\lambda$  and  $\nu_\lambda = 1$  for all  $\lambda \in \Lambda_A$ .

The *resolvent*:

$$R(z; A) \equiv (zI - A)^{-1} ,$$

defined with the help of the continuous complex variable  $z \in \mathbb{C}$ , captures all of the spectral information about  $A$  through the poles of the resolvent's matrix elements. In fact, the resolvent contains more than just the spectrum: the order of each pole gives the index of the corresponding eigenvalue.

Each eigenvalue  $\lambda$  of  $A$  has an associated *spectral projection operator*  $A_\lambda$ , which is the residue of the resolvent as  $z \rightarrow \lambda$ :

$$A_\lambda \equiv \frac{1}{2\pi i} \oint_{C_\lambda} R(z; A) dz , \quad (12)$$

where  $C_\lambda$  is a counterclockwise contour in the complex plane around eigenvalue  $\lambda$ . The residue of the matrix can be calculated elementwise.

The projection operators are orthonormal:

$$A_\lambda A_\zeta = \delta_{\lambda, \zeta} A_\lambda , \quad (13)$$



and sum to the identity:

$$I = \sum_{\lambda \in \Lambda_A} A_\lambda . \quad (14)$$

For cases where  $\nu_\lambda = 1$ , we found that the projection operator associated with  $\lambda$  can be calculated as [4]:

$$A_\lambda = \prod_{\substack{\zeta \in \Lambda_A \\ \zeta \neq \lambda}} \left( \frac{A - \zeta I}{\lambda - \zeta} \right)^{\nu_\zeta} . \quad (15)$$

Not all projection operators of a nondiagonalizable operator can be found directly from Eq. (15), since some have index larger than one. However, if there is only one eigenvalue that has index larger than one—the *almost diagonalizable* case treated in Part II—then Eq. (15), together with the fact that the projection operators must sum to the identity, *does* give a full solution to the set of projection operators. Next, we consider the general case, with no restriction on  $\nu_\lambda$ .

### B. Eigenprojectors via left and right generalized eigenvectors

In general, as we now discuss, an operator's eigenprojectors can be obtained from all left and right eigenvectors and generalized eigenvectors associated with the eigenvalue. Let  $(\Lambda_A) = (\lambda_1, \lambda_2, \dots, \lambda_n)$  be the  $n$ -tuple of eigenvalues in which each eigenvalue  $\lambda \in \Lambda_A$  is listed  $g_\lambda$  times. So  $\sum_{k=1}^n \delta_{\lambda_k, \lambda} = g_\lambda$ , and  $n = \sum_{\lambda \in \Lambda_A} g_\lambda$  is the total number of Jordan blocks in the Jordan canonical form. Each  $\lambda_k \in (\Lambda_A)$  corresponds to a particular Jordan block of size  $m_k$ . The index  $\nu_\lambda$  of  $\lambda$  is thus:

$$\nu_\lambda = \max\{\delta_{\lambda, \lambda_k} m_k\}_{k=1}^n .$$

There is a corresponding  $n$ -tuple of  $m_k$ -tuples of linearly-independent *generalized right-eigenvectors*:

$$\left( (|\lambda_1^{(m)}\rangle)_{m=1}^{m_1}, (|\lambda_2^{(m)}\rangle)_{m=1}^{m_2}, \dots, (|\lambda_n^{(m)}\rangle)_{m=1}^{m_n} \right) ,$$

where:

$$(|\lambda_k^{(m)}\rangle)_{m=1}^{m_k} \equiv \left( |\lambda_k^{(1)}\rangle, |\lambda_k^{(2)}\rangle, \dots, |\lambda_k^{(m_k)}\rangle \right)$$

and a corresponding  $n$ -tuple of  $m_k$ -tuples of linearly-independent *generalized left-eigenvectors*:

$$\left( \langle\langle\lambda_1^{(m)}|\rangle\rangle_{m=1}^{m_1}, \langle\langle\lambda_2^{(m)}|\rangle\rangle_{m=1}^{m_2}, \dots, \langle\langle\lambda_n^{(m)}|\rangle\rangle_{m=1}^{m_n} \right) ,$$

where:

$$\langle\langle\lambda_k^{(m)}|\rangle\rangle_{m=1}^{m_k} \equiv \left( \langle\lambda_k^{(1)}|, \langle\lambda_k^{(2)}|, \dots, \langle\lambda_k^{(m_k)}| \right)$$

such that:

$$(A - \lambda_k I) |\lambda_k^{(m+1)}\rangle = |\lambda_k^{(m)}\rangle \quad (16)$$

and:

$$\langle\lambda_k^{(m+1)}| (A - \lambda_k I) = \langle\lambda_k^{(m)}| , \quad (17)$$

for  $0 \leq m \leq m_k - 1$ , where  $|\lambda_j^{(0)}\rangle = \vec{0}$  and  $\langle\lambda_j^{(0)}| = \vec{0}$ . Specifically,  $|\lambda_k^{(1)}\rangle$  and  $\langle\lambda_k^{(1)}|$  are conventional right and left eigenvectors, respectively.

Most directly, the generalized right and left eigenvectors can be found as the nontrivial solutions to:

$$(A - \lambda_k I)^m |\lambda_k^{(m)}\rangle = \vec{0}$$

and:

$$\langle\lambda_k^{(m)}| (A - \lambda_k I)^m = \vec{0} ,$$

respectively. Imposing appropriate normalization, we find that:

$$\langle\lambda_j^{(m)}| \lambda_k^{(n)}\rangle = \delta_{j,k} \delta_{m+n, m_k+1} . \quad (18)$$

Crucially, right and left eigenvectors are no longer simply related by complex-conjugate transposition and right eigenvectors are not necessarily orthogonal to each other. Rather, left eigenvectors and generalized eigenvectors form a dual basis to the right eigenvectors and generalized eigenvectors. Somewhat surprisingly, the *most* generalized left eigenvector  $\langle\lambda_k^{(m_k)}|$  associated with  $\lambda_k$  is dual to the *least* generalized right eigenvector  $|\lambda_k^{(1)}\rangle$  associated with  $\lambda_k$ :

$$\langle\lambda_k^{(m_k)}| \lambda_k^{(1)}\rangle = 1 .$$

Explicitly, we find that the spectral projection operators for a nondiagonalizable matrix can be written as:

$$A_\lambda = \sum_{k=1}^n \sum_{m=1}^{m_k} \delta_{\lambda, \lambda_k} |\lambda_k^{(m)}\rangle \langle\lambda_k^{(m_k+1-m)}| . \quad (19)$$

### C. Companion operators and resolvent decomposition

It is useful to introduce the generalized set of *companion operators*:

$$A_{\lambda,m} = A_{\lambda}(A - \lambda I)^m, \quad (20)$$

for  $\lambda \in \Lambda_A$  and  $m \in \{0, 1, 2, \dots\}$ . These operators satisfy the following semigroup relation:

$$A_{\lambda,m} A_{\lambda,n} = \delta_{\lambda,\zeta} A_{\lambda,m+n}. \quad (21)$$

$A_{\lambda,m}$  reduces to the eigenprojector for  $m = 0$ :

$$A_{\lambda,0} = A_{\lambda}, \quad (22)$$

and it exactly reduces to the zero-matrix for  $m \geq \nu_{\lambda}$ :

$$A_{\lambda,m} = \mathbf{0}. \quad (23)$$

Crucially, we can rewrite the resolvent as a weighted sum of the companion matrices  $\{A_{\lambda,m}\}$ , with complex coefficients that have poles at each eigenvalue  $\lambda$  up to the eigenvalue's index  $\nu_{\lambda}$ :

$$R(z; A) = \sum_{\lambda \in \Lambda_A} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(z - \lambda)^{m+1}} A_{\lambda,m}. \quad (24)$$

Ultimately these results allow us to easily evaluate arbitrary functions of nondiagonalizable operators, to which we now turn. (Reference [4] gives more background.)

### D. Functions of nondiagonalizable operators

The *meromorphic functional calculus* [4] gives meaning to arbitrary functions  $f(\cdot)$  of any linear operator  $A$ . Its starting point is the Cauchy-integral-like formula:

$$f(A) = \sum_{\lambda \in \Lambda_A} \frac{1}{2\pi i} \oint_{C_{\lambda}} f(z) R(z; A) dz, \quad (25)$$

where  $C_{\lambda}$  denotes a sufficiently small counterclockwise contour around  $\lambda$  in the complex plane such that no singularity of the integrand besides the possible pole at  $z = \lambda$  is enclosed by the contour.

Invoking Eq. (24) yields the desired formulation:

$$f(A) = \sum_{\lambda \in \Lambda_A} \sum_{m=0}^{\nu_{\lambda}-1} A_{\lambda,m} \frac{1}{2\pi i} \oint_{C_{\lambda}} \frac{f(z)}{(z - \lambda)^{m+1}} dz. \quad (26)$$

Hence, with the eigenprojectors  $\{A_{\lambda}\}_{\lambda \in \Lambda_A}$  in hand, eval-

uating an arbitrary function of the nondiagonalizable operator  $A$  comes down to the evaluation of several residues.

Typically, evaluating Eq. (26) requires less work than one might expect when looking at the equation in its full generality. For example, whenever  $f(z)$  is holomorphic (i.e., well behaved) at  $z = \lambda$ , the residue simplifies to:

$$\frac{1}{2\pi i} \oint_{C_{\lambda}} \frac{f(z)}{(z - \lambda)^{m+1}} dz = \frac{1}{m!} f^{(m)}(\lambda),$$

where  $f^{(m)}(\lambda)$  is the  $m^{\text{th}}$  derivative of  $f(z)$  evaluated at  $z = \lambda$ . However, if  $f(z)$  has a pole or zero at  $z = \lambda$ , then it substantially changes the complex contour integration. In the simplest case, when  $A$  is diagonalizable and  $f(z)$  is holomorphic at  $\Lambda_A$ , the matrix-valued function reduces to the simple form:

$$f(A) = \sum_{\lambda \in \Lambda_A} f(\lambda) A_{\lambda}.$$

Moreover, if  $\lambda$  is nondegenerate, then:

$$A_{\lambda} = \frac{|\lambda\rangle\langle\lambda|}{\langle\lambda|\lambda\rangle},$$

although  $\langle\lambda|$  here should be interpreted as the solution to the left eigenequation  $\langle\lambda|A = \lambda\langle\lambda|$  and, in general,  $\langle\lambda| \neq (|\lambda\rangle)^{\dagger}$ .

The meromorphic functional calculus agrees with the Taylor-series approach whenever the series converges and agrees with the holomorphic functional calculus of Ref. [70] whenever  $f(z)$  is holomorphic at  $\Lambda_A$ . However, when both these functional calculi fail, the meromorphic functional calculus extends the domain of  $f(A)$  in a way that is key to the following analysis. We show, for example, that within the meromorphic functional calculus, the negative-one power of a singular operator is the Drazin inverse. The Drazin inverse effectively inverts everything that is invertible. Notably, it appears ubiquitously in the new-found solutions to many complexity measures.

### E. Evaluating residues

How does one use Eq. (26)? It says that the spectral decomposition of  $f(A)$  reduces to the evaluation of several residues, where:

$$\text{Res}(g(z), z \rightarrow \lambda) = \frac{1}{2\pi i} \oint_{C_{\lambda}} g(z) dz.$$

So, to make progress with Eq. (26), we must evaluate functional-dependent residues of the form  $\text{Res}(f(z)/(z - \lambda)^{m+1}, z \rightarrow \lambda)$ . This is basic complex analysis. Recall that the residue of a complex-valued

function  $g(z)$  around its isolated pole  $\lambda$  of order  $n + 1$  can be calculated from:

$$\text{Res}(g(z), z \rightarrow \lambda) = \frac{1}{n!} \lim_{z \rightarrow \lambda} \frac{d^n}{dz^n} [(z - \lambda)^{n+1} g(z)] .$$

### F. Decomposing $A^L$

Equation (26) allows us to explicitly derive the spectral decomposition of powers of an operator. For  $f(A) = A^L \rightarrow f(z) = z^L$ ,  $z = 0$  can be either a zero or a pole of  $f(z)$ , depending on the value of  $L$ . In either case, an eigenvalue of  $\lambda = 0$  will distinguish itself in the residue calculation of  $A^L$  via its unique ability to change the order of the pole (or zero) at  $z = 0$ .

For example, at this special value of  $\lambda$  and for integer  $L > 0$ ,  $\lambda = 0$  induces poles that *cancel* with the zeros of  $f(z) = z^L$ , since  $z^L$  has a zero at  $z = 0$  of order  $L$ . For integer  $L < 0$ , an eigenvalue of  $\lambda = 0$  *increases* the order of the  $z = 0$  pole of  $f(z) = z^L$ . For all other eigenvalues, the residues will be as expected.

Hence, for any  $L \in \mathbb{C}$ :

$$A^L = \left[ \sum_{\substack{\lambda \in \Lambda_A \\ \lambda \neq 0}} \sum_{m=0}^{\nu_\lambda - 1} \binom{L}{m} \lambda^{L-m} A_{\lambda,m} \right] + [0 \in \Lambda_A] \sum_{m=0}^{\nu_0 - 1} \delta_{L,m} A_0 A^m , \quad (27)$$

where  $\binom{L}{m}$  is the generalized binomial coefficient:

$$\binom{L}{m} = \frac{1}{m!} \prod_{n=1}^m (L - n + 1) , \quad (28)$$

with  $\binom{L}{0} = 1$  and where  $[0 \in \Lambda_A]$  is the Iverson bracket. The latter takes value 1 if 0 is an eigenvalue of  $A$  and value 0 if not. Equation (27) applies to any linear operator with only isolated singularities in its resolvent.

If  $L$  is a nonnegative integer such that  $L \geq \nu_\lambda - 1$  for all  $\lambda \in \Lambda_A$ , then:

$$A^L = \sum_{\substack{\lambda \in \Lambda_A \\ \lambda \neq 0}} \sum_{m=0}^{\nu_\lambda - 1} \binom{L}{m} \lambda^{L-m} A_{\lambda,m} , \quad (29)$$

where  $\binom{L}{m}$  is now reduced to the traditional binomial coefficient  $L!/m!(L-m)!$ .

The form of Eq. (27), together with our earlier operator expressions for complexity measures that take on a cascading form, directly leads to the first fully general closed-form expressions for correlation, myopic entropy rates, and remaining state uncertainty, among others,

for the broad class of processes that can be generated by HMMs. This will be made explicit in Part II, where the consequences will also be unraveled.

### G. Drazin inverse

The negative-one power  $A^{-1}$  of a linear operator is in general *not* the same as its inverse  $\text{inv}(A)$ , since the latter need not exist. However, the negative-one power of a linear operator is always defined via Eq. (27):

$$A^{-1} = \sum_{\lambda \in \Lambda_A \setminus \{0\}} \sum_{m=0}^{\nu_\lambda - 1} (-1)^m \lambda^{-1-m} A_{\lambda,m} . \quad (30)$$

Notably, when the operator is singular, we find that:

$$AA^{-1} = I - A_0 .$$

This is the *Drazin inverse*  $A^{\mathcal{D}}$  of  $A$ , also known as the  $\{1^{\nu_0}, 2, 5\}$ -inverse [71]. (Note that it is *not* the same as the Moore–Penrose pseudo-inverse.) Although the Drazin inverse is usually defined axiomatically to satisfy certain criteria. In contrast, Ref. [4] naturally *derived* it as the negative one power of a singular operator in the meromorphic functional calculus.

Whenever  $A$  is invertible, however,  $A^{-1} = \text{inv}(A)$ . That said, we should not confuse this coincidence with equivalence. More to the point, there is no reason other than historical accidents of notation that the negative-one power should in general be equivalent to the inverse—especially if an operator is not invertible. To avoid confusing  $A^{-1}$  with  $\text{inv}(A)$ , we use the notation  $A^{\mathcal{D}}$  for the Drazin inverse of  $A$ . Still,  $A^{\mathcal{D}} = \text{inv}(A)$  whenever  $0 \notin \Lambda_A$ .

Although Eq. (30) is a constructive way to build the Drazin inverse, it suggests more work than is actually necessary. We derived several simple constructions for it that require only the original operator and the eigenvalue-0 projector. For example, Ref. [4] found that, for any  $c \in \mathbb{C} \setminus \{0\}$ :

$$A^{\mathcal{D}} = (I - A_0)(A + cA_0)^{-1} . \quad (31)$$

Later, we will also need the decomposition of  $(I - W)^{\mathcal{D}}$ , as it enters into many closed-form complexity expressions related to accumulated transients—the past–future mutual information among them. Reference [4] showed that:

$$(I - T)^{\mathcal{D}} = [I - (T - T_1)]^{-1} - T_1 \quad (32)$$

for any stochastic matrix  $T$ , where  $T_1$  is the projection operator associated with  $\lambda = 1$ . If  $T$  is the state-transition matrix of an ergodic process, then the RHS of

Eq. (32) becomes especially simple to evaluate since then  $T_1 = |\mathbf{1}\rangle\langle\pi|$ .

Somewhat tangentially, this connects to the *fundamental matrix*  $Z = (I - T + T_1)^{-1}$  used by Ref. [72] in its analysis of Markovian dynamics. More immediately, Eq. (32) plays a prominent role when deriving excess entropy and synchronization information. The explicit spectral decomposition is also useful:

$$(I - T)^{\mathcal{D}} = \sum_{\lambda \in \Lambda_T \setminus \{1\}} \sum_{m=0}^{\nu_\lambda - 1} \frac{1}{(1 - \lambda)^{m+1}} T_{\lambda, m}. \quad (33)$$

## VII. PROJECTION OPERATORS FOR STOCHASTIC DYNAMICS

The preceding employed the notation that  $A$  is a general linear operator. In the following, we reserve  $T$  for the operator of a stochastic transition dynamic, as in the state-to-state transition dynamic of an HMM:  $T = \sum_{x \in \mathcal{A}} T^{(x)}$ . If the state space is finite and has a stationary distribution, then  $T$  has a representation that is a nonnegative row-stochastic—all rows sum to unity—transition matrix.

We are now in a position to summarize several useful properties for the projection operators of any row-stochastic matrix  $T$ . Naturally, if one uses column-stochastic instead of row-stochastic matrices, all results can be translated by simply taking the transpose of every line in the derivations. (Recall that  $(ABC)^{\top} = C^{\top} B^{\top} A^{\top}$ .)

The fact that all elements of the transition matrix are real-valued guarantees that, for each  $\lambda \in \Lambda_T$ , its complex conjugate  $\bar{\lambda}$  is also in  $\Lambda_T$ . Moreover, the spectral projection operator associated with the complex conjugate of  $\lambda$  is  $T_{\bar{\lambda}}$ 's complex conjugate:

$$T_{\bar{\lambda}} = \overline{T_{\lambda}}.$$

This also implies that  $T_{\lambda}$  is real if  $\lambda$  is real.

If the dynamic induced by  $T$  has a stationary distribution over the state space, then  $T$ 's spectral radius is unity and all its eigenvalues lie on or within the unit circle in the complex plane. The maximal eigenvalues have unity magnitude and  $1 \in \Lambda_T$ . Moreover, an extension of the Perron–Frobenius theorem guarantees that eigenvalues on the unit circle have algebraic multiplicity equal to their geometric multiplicity. And, so,  $\nu_{\zeta} = 1$  for all  $\zeta \in \{\lambda \in \Lambda_T : |\lambda| = 1\}$ .

$T$ 's index-one eigenvalue  $\lambda = 1$  is associated with stationarity of the hidden Markov model.  $T$ 's other eigenvalues on the unit circle are roots of unity and correspond to deterministic periodicities within the process.




FIG. 3. (a) Weighted directed graph (digraph) of the feedback matrix  $A$  of a cyclic cluster structure that contributes eigenvalues  $\Lambda_A = \left\{ \left( \prod_{i=1}^N \alpha_i \right)^{1/N} e^{in2\pi/N} \right\}_{n=0}^{N-1}$  with algebraic multiplicities  $a_\lambda = 1$  for all  $\lambda \in \Lambda_A$ . (b) Weighted digraph of the feedback matrix  $A$  of a doubly cyclic cluster structure that contributes eigenvalues  $\Lambda_A = \{0\} \cup \left\{ \alpha \left[ \left( \prod_{i=1}^N \beta_i \right) + \left( \prod_{i=1}^N \gamma_i \right) \right]^{\frac{1}{N+1}} e^{\frac{in2\pi}{N+1}} \right\}_{n=0}^N$  with algebraic multiplicities  $a_0 = N - 1$  and  $a_\lambda = 1$  for  $\lambda \neq 0$ . (This eigenvalue “rule” depends on having the same number of  $\beta$ -transitions as  $\gamma$ -transitions.) The 0-eigenvalue only has geometric multiplicity of  $g_0 = 1$ , so the structure is nondiagonalizable for  $N > 2$ . Nevertheless, the generalized eigenvectors are easy to construct. The spectral analysis of the cluster structure in (b) suggests more general rules that can be gleaned from reading-off eigenvalues from digraph clusters; e.g., if a *chain* of  $\alpha$ 's appears in the bisecting path.

### A. Row sums

If  $T$  is row-stochastic, then by definition:

$$T|\mathbf{1}\rangle = |\mathbf{1}\rangle.$$

Hence, via the general eigenprojector construction Eq. (19) and the general orthogonality condition Eq. (18), we find that:

$$T_{\lambda}|\mathbf{1}\rangle = \delta_{\lambda,1}|\mathbf{1}\rangle. \quad (34)$$

This shows that  $T$ 's projection operator  $T_1$  is row-stochastic, whereas each row of every other projection operator must sum to zero. This can also be viewed as a consequence of conservation of probability for dynamics over Markov models.

### B. Expected stationary distribution

If unity is the only eigenvalue of  $\Lambda_T$  on the unit circle, then the process has no deterministic periodicities. In this case, every initial condition leads to a stationary asymptotic distribution. The expected stationary distribution

FIG. 4. Construction of  $W$ -eigenprojectors  $W_\lambda$  from low-level  $A$ -projectors and  $C$ -projectors, when  $W = \begin{bmatrix} A & B \\ \mathbf{0} & C \end{bmatrix}$ . (Recall that  $(\lambda I - A)^{-1}$  and  $(\lambda I - C)^{-1}$  can be constructed from the lower-level projectors.) For simplicity, we assume that the algebraic multiplicity  $a_\lambda = 1$  in each of these cases.

bution  $\pi_\alpha$  from any initial distribution  $\alpha$  is:

$$\begin{aligned} \langle \pi_\alpha | &= \lim_{L \rightarrow \infty} \langle \alpha | T^L \\ &= \langle \alpha | T_1 . \end{aligned} \quad (35)$$

An attractive feature of Eq. (35) is that it holds even for nonergodic processes—those with multiple stationary components.

When the stochastic process is ergodic (one stationary component), then  $a_1 = 1$  and there is only one stationary distribution  $\pi$ . The  $T_1$  projection operator becomes:

$$T_1 = |\mathbf{1}\rangle \langle \pi| , \quad (36)$$

even if there are deterministic periodicities. Deterministic periodicities imply that different initial conditions may still induce different asymptotic oscillations, according to  $\{T_\lambda : |\lambda| = 1\}$ . In the case of ergodic processes without deterministic periodicities, every initial condition relaxes to the same steady-state distribution over the hidden states:  $\langle \pi_\alpha | = \langle \alpha | T_1 = \langle \pi |$  regardless of  $\alpha$ , so long as  $\alpha$  is a properly normalized probability distribution.

As suggested in Ref. [4], the new results above extend spectral theory to arbitrary functions of nondiagonalizable operators in a way that contributes to a *spectral weighted digraph theory* beyond the purview of spectral

graph theory proper [73]. Moreover, this enables new analyses. In particular, the spectra of undirected graphs and their graph Laplacian matrices have been studied extensively and continue to be. However, those efforts have been extensive due in part to the spectral theory for normal operators applying directly to both undirected graphs and their Laplacians. Digraph spectra have also been studied [74], but to a much lesser extent. Again, this is due in part to the spectral theorem not typically applying, rendering this case much more complicated. Thus, the spectral theory of nonnormal and nondiagonalizable operators offers new opportunities. This not only hints at the importance of extracting eigenvalues from directed graph motifs, but also begins to show how eigenvectors and eigenprojectors can be built up iteratively from directed graph clusters.

## VIII. SPECTRA BY INSPECTION

The next sections show how spectra and eigenprojectors can be intuited, computed, and applied in the analysis of complex systems. These techniques often make the problem at hand analytically tractable, and they will be used in the examples of Part II to give exact expressions for complexity measures.

	Discrete time	Continuous time
Derivatives of cascading $\uparrow$		
Cascading	$\langle \cdot   A^L   \cdot \rangle$	$\langle \cdot   e^{tG}   \cdot \rangle$
Integrals of cascading $\downarrow$		
Accumulated transients	$\langle \cdot   (\sum_L (A - A_1)^L)   \cdot \rangle$	$\langle \cdot   (\int (e^{tG} - G_0) dt)   \cdot \rangle$
modulated accumulation	$\langle \cdot   (\sum_L (zA)^L)   \cdot \rangle$	$\langle \cdot   (\int (ze^{tG})^t dt)   \cdot \rangle$

TABLE III. Once we identify the hidden linear dynamic behind our questions, most are either of the *cascading* or *accumulating* type. Moreover, if a complexity measure accumulates transients, the Drazin inverse is likely to appear. Interspersed accumulation can be a helpful theoretical tool, since all derivatives and integrals of cascading type can be calculated, if we know the modified accumulation with  $z \in \mathbb{C}$ . With  $z \in \mathbb{C}$ , modulated accumulation involves an operator-valued  $z$ -transform. However with  $z = e^{i\omega}$  and  $\omega \in \mathbb{R}$ , modulated accumulation involves an operator-valued Fourier-transform.

### A. Eigenvalues from graph structure

Consider a directed graph structure with cascading dependencies: one cluster of nodes feeds back only to itself according to matrix  $A$  and feeds forward to another cluster of nodes according to matrix  $B$ , which is not necessarily a square matrix. The second cluster feeds back only to itself according to matrix  $C$ . The latter node cluster might also feed forward to another cluster, but such considerations can be applied iteratively.

The simple situation just described is summarized, with proper index permutation, by a block matrix of the form:  $W = \begin{bmatrix} A & B \\ \mathbf{0} & C \end{bmatrix}$ . In this case, it is easy to see that:

$$\det(W - \lambda I) = \begin{vmatrix} A - \lambda I & B \\ \mathbf{0} & C - \lambda I \end{vmatrix} \quad (37)$$

$$= |A - \lambda I| |C - \lambda I|. \quad (38)$$

And so,  $\Lambda_W = \Lambda_A \cup \Lambda_C$ . This simplification presents an opportunity to read off eigenvalues from clustered graph structures that often appear in practice, especially for transient graph structures associated with synchronization, as with transient mixed-state transitions in MSPs.

Cyclic cluster structures (say, of length  $N$  and edge-weights  $\alpha_1$  through  $\alpha_N$ ) yield especially simple spectra:

$$\Lambda_A = \left\{ \left( \prod_{i=1}^N \alpha_i \right)^{1/N} e^{in2\pi/N} \right\}_{n=0}^{N-1}. \quad (39)$$

That is, the eigenvalues are simply the  $N^{\text{th}}$  roots of the product of all of the edge-weights. See Fig. 3a.

Similar rules for reading off spectra from other cluster structures exist. Although we cannot list them exhaustively here, we give another simple but useful rule in Fig. 3b. It also indicates the ubiquity of nondiagonalizability in weighted digraph structures. This second rule is suggestive of further generalizations where spectra can be read off from common digraph motifs.

### B. Eigenprojectors from graph structure

We just outlined how clustered directed graph structures yield simplified joint spectra. Is there a corresponding simplification of the spectral projection operators? In fact, there is and it leads to an iterative construction of “higher-level” projectors from “lower-level” clustered components. In contrast to the joint spectrum though, that completely ignores the feedforward matrix  $B$ , the emergent projectors do require  $B$  to pull the associated eigencontributions into the generalized setting. Figure 4 summarizes the results for the simple case of nondegenerate eigenvalues. The general case is constructed similarly.

The preceding results imply a number of algorithms, both for analytic and numerical calculations. Most directly, this points to the fact that eigenanalysis can be partitioned into a series of simpler problems that are later combined to a final solution. However, in addition to more efficient serial computation, there are opportunities for numerical parallelization of the algorithms to compute the eigenprojectors, whether they are computed directly, say from Eq. (15), or from right and left eigenvectors and generalized eigenvectors. These opportunities for further optimization are perhaps rare considering how extremely well developed the field of numerical linear algebra already is. That said, the automation now possible will be key to applying our analysis methods to real systems with immense data produced from very high-dimensional state spaces.

## IX. CONCLUSION

Surprisingly, many questions we ask about a structured stochastic nonlinear process imply a linear dynamic over a preferred hidden state space. These questions often concern predictability and prediction. To make predictions about the real world, though, it is not sufficient to have a model of the world. Additionally, the predictor must synchronize their model to the real-world data

Genre	Implied linear transition dynamic	Example Questions		
		Cascading	Accumulated transients	Modulated accumulation
OVERT OBSERVATIONAL	Transition matrix $T$ of any HMM	Correlations, $\gamma(L)$ : $\langle \pi \bar{\mathcal{A}}   T^{ L -1}   \mathcal{A} \mathbf{1} \rangle$	Green–Kubo transport coefficients	Power spectra, $P(\omega)$ : $2\Re \langle \pi \bar{\mathcal{A}}   (e^{i\omega} I - T)^{-1}   \mathcal{A} \mathbf{1} \rangle$
PREDICTABILITY	Transition matrix $W$ of MSP of any HMM	Myopic entropy rate, $h_\mu(L)$ : $\langle \delta_\pi   W^{L-1}   H(W^{\mathcal{A}}) \rangle$	Excess entropy, $\mathbf{E}$ : $\langle \delta_\pi   (I - W)^{\mathcal{D}}   H(W^{\mathcal{A}}) \rangle$	$E(z)$ : $\langle \delta_\pi   (zI - W)^{-1}   H(W^{\mathcal{A}}) \rangle$
OPTIMAL PREDICTION	Transition matrix $\mathcal{W}$ of MSP of $\epsilon$ -machine	Causal state uncertainty, $\mathcal{H}^+(L)$ : $\langle \delta_\pi   \mathcal{W}^L   H[\eta] \rangle$	Synchronization info, $\mathbf{S}$ : $\langle \delta_\pi   (I - \mathcal{W})^{\mathcal{D}}   H[\eta] \rangle$	$S(z)$ : $\langle \delta_\pi   (zI - \mathcal{W})^{-1}   H[\eta] \rangle$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

TABLE IV. Genres of complexity questions given in order of increasing sophistication; summary of Part I and a preview of Part II. Each implies a different linear transition dynamic. Closed-form formulae are given for several complexity measures, showing the similarity among them down the same column. Formulae in the same row have matching bra-ket pairs. The similarity within the column corresponds to similarity in the time-evolution implied by the question type. The similarity within the row corresponds to similarity in question genre.

that has been observed up to the present time. This metadynamic of synchronization—the transition structure among belief states—is intrinsically linear, but is typically nondiagonalizable.

We presented results for observed processes generated by HMMs. However, the results easily apply to other state-based models, including *observable operator models* (OOMs) [75] and *generalized HMMs* (GHMMs) [45]. In each case, the observation-induced synchronizing metadynamic is still an HMM. It will also be useful to adapt our methods to open quantum models, where a density matrix evolves via environmental influence and a protocol for partial measurements (POVMs) induces a synchronizing (meta)dynamics.

Recall organizational Tables I and II from the Introduction. After all the intervening detail, let’s consider a more nuanced formulation. We saw that once we frame questions in terms of the hidden linear transition dynamic, complexity measures are usually either of the cascading or accumulation type. Scalar complexity measures often accumulate only the interesting transient structure that rides on top of the asymptotics. Skimming off the asymptotics led to the Drazin inverse. Modified accumulation turned complexity scalars into complexity functions. Tables III and IV summarize the results. Notably, Table IV gives closed-form formulae for many complexity measures that previously were only expressed as infinite sums over functions of probabilities.

Let us remind ourselves: Why, in this analysis, were nondiagonalizable dynamics noteworthy? They are since the *metadynamics* of diagonalizable dynamics are generi-

cally nondiagonalizable. And, this is typically due to the 0-eigenvalue subspace that is responsible for the initial, ephemeral epoch of symmetry collapse. The metadynamics of transitioning between belief states demonstrated this explicitly. However, other metadynamics beyond those focused on prediction are also generically nondiagonalizable. For example, in the analysis of quantum compression, crypticity, and other aspects of hidden structure, the relevant linear dynamic is not the MSP. Instead, it is a nondiagonalizable structure that can be fruitfully analyzed with the same generalized spectral theory of nonnormal operators [4].

Using the appropriate dynamic for common complexity questions and the meromorphic functional calculus to overcome nondiagonalizability, the sequel (Part II) goes on to develop closed-form expressions for complexity measures as simple functions of the corresponding transition dynamic of the implied HMM.

## ACKNOWLEDGMENTS

JPC thanks the Santa Fe Institute for its hospitality. The authors thank Chris Ellison, Ryan James, John Mahoney, Alec Boyd, and Dowman Varn for helpful discussions. 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-13-1-0340, and W911NF-13-1-0390.

[1] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy conver-

gence. *CHAOS*, 13(1):25–54, 2003. 1, 4, 5, 6, 13

- [2] S. E. Marzen and J. P. Crutchfield. Nearly maximally predictive features and their dimensions. *Phys. Rev. E*, 95(5):051301(R), 2017. 2, 6
- [3] J. P. Crutchfield, C. J. Ellison, and P. M. Riechers. Exact complexity: The spectral decomposition of intrinsic computation. *Phys. Lett. A*, 380(9):998–1002, 2016. 2, 3, 12
- [4] P. M. Riechers and J. P. Crutchfield. Beyond the spectral theorem: Decomposing arbitrary functions of nondiagonalizable operators. arxiv.org:1607.06526. 2, 3, 15, 16, 17, 18, 19, 21, 23
- [5] While we follow Shannon [12] in this, it differs from the more widely used state-labeled HMMs. 2
- [6] C. Moore and J. P. Crutchfield. Quantum automata and quantum grammars. *Theoret. Comp. Sci.*, 237:1-2:275–306, 2000. 2
- [7] L. A. Clark, W. Huang, T. M. Barlow, and A. Beige. Hidden quantum markov models and open quantum systems with instantaneous feedback. *New. J. Phys.*, 14:143–151, 2015. 2
- [8] O. Penrose. *Foundations of statistical mechanics; a deductive treatment*. Pergamon Press, Oxford, 1970. 2
- [9] U. Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.*, 75:126001, 2012. 2
- [10] R. Klages, W. Just, and C. Jarzynski, editors. *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*. Wiley, New York, 2013.
- [11] J. Bechhoefer. Hidden Markov models for stochastic thermodynamics. *New. J. Phys.*, 17:075003, 2015. 2, 25
- [12] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948. 2, 4, 24
- [13] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006. 2, 4, 7, 13, 25
- [14] L. R. Rabiner and B. H. Juang. An introduction to hidden Markov models. *IEEE ASSP Magazine*, January, 1986. 2
- [15] R. A. Roberts and C. T. Mullis. *Digital Signal Processing*. Addison-Wesley, Reading, Massachusetts, 1987.
- [16] L. R. Rabiner. A tutorial on hidden Markov models and selected applications. *IEEE Proc.*, 77:257, 1989. 2
- [17] M. O. Rabin. Probabilistic automata. *Info. Control*, 6:230–245, 1963. 2
- [18] W. J. Ewens. *Mathematical Population Genetics*. Springer, New York, second edition, 2004. 2
- [19] M. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belnap Press, New York, 2006. 2
- [20] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(January):17–24, 2012. 3, 5, 8
- [21] P. Stoica and R. L. Moses. *Spectral Analysis of Signals*. Pearson Prentice Hall, Upper Saddle River, New Jersey, 2005. 4
- [22] R. W. Hamming. *Digital Filterns*. Dover Publications, New York, third edition, 1997. 4
- [23] M M. Woolfson. *An Introduction to X-ray Crystallography*. Cambridge University Press, Cambridge, United Kingdom, 1997. 4
- [24] M. S. Green. Markoff random processes and the statistical mechanics of time-dependent phenomena. II. Irreversible processes in fluids. *J. Chem. Physics*, 22(3):398–413, 1954. 4
- [25] R. Zwanzig. Time-correlation functions and transport coefficients in statistical mechanics. *Ann. Rev. Phys. Chemistry*, 16(1):67–102, 1965. 4
- [26] J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. *The Theory of Critical Phenomena*. Oxford University Press, Oxford, 1992. 4
- [27] A. N. Kolmogorov. Entropy per unit time as a metric invariant of automorphisms. *Dokl. Akad. Nauk. SSSR*, 124:754, 1959. (Russian) Math. Rev. vol. 21, no. 2035b. 4
- [28] J. P. Crutchfield and N. H. Packard. Symbolic dynamics of noisy chaos. *Physica*, 7D(1-3):201 – 223, 1983. 4
- [29] A. del Junco and M. Rahe. Finitary codings and weak bernoulli partitions. *Proc. AMS*, 75:259, 1979. 5
- [30] R. Shaw. *The Dripping Faucet as a Model Chaotic System*. Aerial Press, Santa Cruz, California, 1984. 5
- [31] P. Grassberger. Toward a quantitative theory of self-generated complexity. *Intl. J. Theo. Phys.*, 25:907, 1986. 5
- [32] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. *Neural Comput.*, 13(11):2409–2463, November 2001. 5
- [33] R. G. James, C. J. Ellison, and J. P. Crutchfield. Anatomy of a bit: Information in a time series observation. *CHAOS*, 21(3):037109, 2011. 5
- [34] P. M. Ara, R. G. James, and J. P. Crutchfield. The elusive present: Hidden past and future dependence and why we build models. *Phys. Rev. E*, 93(2):022143, 2016. 5
- [35] O. Pfante, N. Bertschinger, E. Olbrich, N. Ay, and J. Jost. Comparison between different methods of level identification. *Adv. Complex Sys.*, 17(02):1450007, 2014. 5
- [36] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989. 5
- [37] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Leveraging environmental correlations: The thermodynamics of requisite variety. *J. Stat. Phys.*, 167(6):1555–1585, 2016. 6, 12
- [38] S. Still, J. P. Crutchfield, and C. J. Ellison. Optimal causal inference: Estimating stored information and approximating causal architecture. *CHAOS*, 20(3):037111, 2010. 6
- [39] F. Creutzig, A. Globerson, and N. Tishby. Past-future information bottleneck in dynamical systems. *Phys. Rev. E*, 79(4):041925, 2009.
- [40] S. E. Marzen and J. P. Crutchfield. Predictive rate-distortion for infinite-order Markov processes. *J. Stat. Phys.*, 163(6):1312–1338, 2016. 6
- [41] C. J. Ellison and J. P. Crutchfield. States of states of uncertainty. in preparation. 6, 10, 12, 13
- [42] D. H. Wolpert, J. A. Grochow, E. Libby, and S. DeDeo. Optimal high-level descriptions of dynamical systems.



*arXiv:1409.7403*. 6

- [43] J. R. Mahoney, C. J. Ellison, R. G. James, and J. P. Crutchfield. How hidden are hidden processes? a primer on crypticity and entropy convergence. *CHAOS*, 21(3):037112, 2011. 6, 15
- [44] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994. 6, 8, 10
- [45] D. R. Upper. *Theory and Algorithms for Hidden Markov Models and Generalized Hidden Markov Models*. PhD thesis, University of California, Berkeley, 1997. Published by University Microfilms Intl, Ann Arbor, Michigan. 10, 23, 25
- [46] W. Lohr and N. Ay. Non-sufficient memories that are sufficient for prediction. In *Complex Sciences*, volume 4 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 265–276. 2009. 6
- [47] J. Ruebeck, R. G. James, John R. Mahoney, and J. P. Crutchfield. Prediction and generation of binary Markov processes: Can a finite-state fox catch a markov mouse? *CHAOS*, in press, 2017. arxiv.org:1708.00113. 6
- [48] 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. 6
- [49] N. Travers and J. P. Crutchfield. Exact synchronization for finite-state sources. *J. Stat. Phys.*, 145(5):1181–1201, 2011. 6, 8
- [50] N. Travers and J. P. Crutchfield. Asymptotic synchronization for finite-state sources. *J. Stat. Phys.*, 145(5):1202–1223, 2011. 6, 8, 15
- [51] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. Occam’s quantum strop: Synchronizing and compressing classical cryptic processes via a quantum channel. *Scientific Reports*, 6:20495, 2016. 6
- [52] B. Vanluyten, J. C. Willems, and B. De Moor. Equivalence of state representations for hidden Markov models. *Systems & Control Letters*, 57(5):410 – 419, 2008. 6, 25
- [53] Contrast this with the class-equivalent state-labeled HMMs, also known as Moore HMMs [11, 45, 52, 76]. In automata theory, a finite-state HMM is called a *probabilistic nondeterministic finite automaton* [77]. Information theory [13] refers to them as finite-state *information sources*. And, stochastic process theory defines them as *functions of a Markov chain* [54, 59, 78, 79]. 6
- [54] R. B. Ash. *Information Theory*. Dover Books on Advanced Mathematics. Dover Publications, 1965. 7, 12, 25
- [55] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001. 7
- [56] Automata theory would refer to a uHMM as a *probabilistic deterministic finite automaton* [77]. The awkward terminology does not recommend itself. 7
- [57] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Reading, 1979. 8
- [58] T. Nishikawa and A. E. Motter. Synchronization is optimal in nondiagonalizable networks. *Phys. Rev. E*, 73:065106, 2006. 10
- [59] D. Blackwell. The entropy of functions of finite-state Markov chains. volume 28, pages 13–20, Publishing House of the Czechoslovak Academy of Sciences, Prague, 1957. 10, 25
- [60] 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. 10, 12
- [61] 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. 10, 12
- [62] P. M. Riechers and J. P. Crutchfield. Power spectra of stochastic processes from transition matrices of hidden Markov models. in preparation. 11, 12
- [63] Averaging over  $t$  invokes unconditioned word probabilities that must be calculated using the stationary probability  $\pi$  over the recurrent states. Effectively this ignores any transient nonstationarity that may exist in a process, since only the recurrent part of the HMM presentation plays a role in the autocorrelation function. One practical lesson is that if  $T$  has transient states, they might as well be trimmed prior to such a calculation. 11
- [64] P. Gaspard and X. Wang. Noise, chaos, and  $(\varepsilon, \tau)$ -entropy per unit time. *Physics Reports*, 235(6):291–343, 1993. 13
- [65] S. Marzen, M. R. DeWeese, and J. P. Crutchfield. Time resolution dependence of information measures for spiking neurons: Scaling and universality. *Front. Comput. Neurosci.*, 9:109, 2015. 13
- [66] S. Marzen and J. P. Crutchfield. Informational and causal architecture of continuous-time renewal processes. *J. Stat. Phys.*, 168(a):109–127, 2017.
- [67] S. Marzen and J. P. Crutchfield. Structure and randomness of continuous-time discrete-event processes. *J. Stat. Physics*, 169(2):303–315, 2016. 13
- [68] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes. *J. Phys. A: Math. Theo.*, 42:362002, 2009. 15
- [69] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes: Linear combinations of causal states. arxiv.org:0906.5099 [cond-mat]. 15
- [70] N. Dunford. Spectral theory I. Convergence to projections. *Trans. Am. Math. Soc.*, 54(2):pp. 185–217, 1943. 18
- [71] A. Ben-Israel and T.N.E. Greville. *Generalized Inverses: Theory and Applications*. CMS Books in Mathematics. Springer, 2003. 19
- [72] J. G. Kemeny and J. L. Snell. *Finite Markov chains*, volume 356. Springer, New York, 1960. 20
- [73] D. M. Cvetkovic, M. Doob, and H. Sachs. *Spectra of Graphs: Theory and Applications*. Wiley, New York, New York, third edition, 1998. 21

- [74] R. A. Brualdi. Spectra of digraphs. *Linear Algebra and its Applications*, 432(9):2181–2213, 2010. 21
- [75] H. Jaeger. Observable operator models for discrete stochastic time series. *Neural computation*, 12(6):1371–1398, 2000. 23
- [76] V. Balasubramanian. Equivalence and reduction of hidden Markov models. Technical Report AITR-1370, 1993. 25
- [77] M. Sipser. *Introduction to the Theory of Computation*. Cengage Learning, 2012. 25
- [78] J. J. Birch. Approximations for the entropy for functions of Markov chains. *Ann. Math. Statist.*, 33(2):930–938, 1962. 25
- [79] N. F. Travers. Exponential bounds for convergence of entropy rate approximations in hidden markov models satisfying a path-mergeability condition. *Stochastic Processes and their Applications*, 124(12):4149–4170, 2014. 25