

RESEARCH ARTICLE | AUGUST 30 2021

Divergent predictive states: The statistical complexity dimension of stationary, ergodic hidden Markov processes

Alexandra M. Jurgens; James P. Crutchfield ✉



Chaos 31, 083114 (2021)

<https://doi.org/10.1063/5.0050460>



View
Online



Export
Citation

CrossMark

Articles You May Be Interested In

Information symmetries in irreversible processes

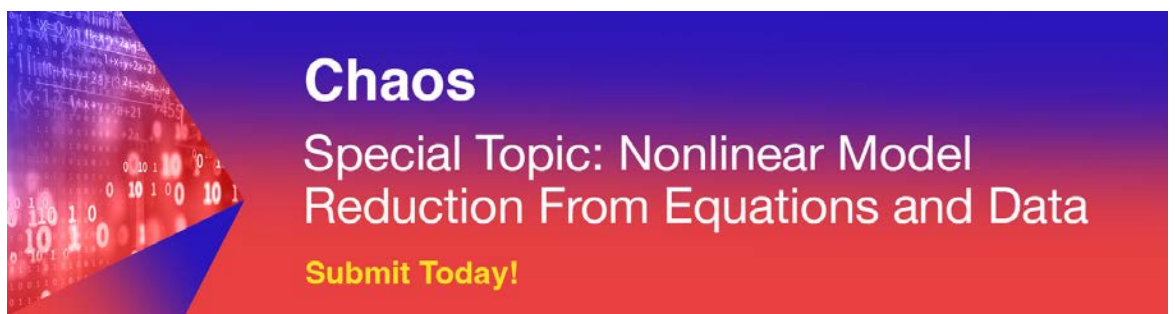
Chaos (September 2011)

Synchronization and control in intrinsic and designed computation: An information-theoretic analysis of competing models of stochastic computation

Chaos (September 2010)

The potential of silver nanoparticles as heterogeneous catalysts for generation of biodiesel from vegetable oils

AIP Conference Proceedings (June 2020)



Chaos
Special Topic: Nonlinear Model
Reduction From Equations and Data
Submit Today!

Divergent predictive states: The statistical complexity dimension of stationary, ergodic hidden Markov processes

Cite as: Chaos 31, 083114 (2021); doi: 10.1063/5.0050460

Submitted: 15 March 2021 · Accepted: 4 July 2021 ·

Published Online: 30 August 2021



View Online



Export Citation



CrossMark

Alexandra M. Jurgens^{a)} and James P. Crutchfield^{b)} 

AFFILIATIONS

Complexity Sciences Center, Physics and Astronomy Department, University of California at Davis, One Shields Avenue, Davis, California 95616, USA

^{a)}amjurgens@ucdavis.edu

^{b)}Author to whom correspondence should be addressed: chaos@ucdavis.edu

ABSTRACT

Even simply defined, finite-state generators produce stochastic processes that require tracking an uncountable infinity of probabilistic features for optimal prediction. For processes generated by hidden Markov chains, the consequences are dramatic. Their predictive models are generically infinite state. Until recently, one could determine neither their intrinsic randomness nor structural complexity. The prequel to this work introduced methods to accurately calculate the Shannon entropy rate (randomness) and to constructively determine their minimal (though, infinite) set of predictive features. Leveraging this, we address the complementary challenge of determining how structured hidden Markov processes are by calculating their statistical complexity dimension—the information dimension of the minimal set of predictive features. This tracks the divergence rate of the minimal memory resources required to optimally predict a broad class of truly complex processes.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0050460>

In the late 1950s, the American mathematician David Blackwell discovered a surprising and counterintuitive phenomenon—one that explains why modeling complex systems is so challenging. Blackwell realized that even simple generators of stochastic processes require infinite models to be optimally predicted. During the same period, the Soviet mathematicians Andrei Kolmogorov and his student Jasha Sinai adapted Claude Shannon's mathematical theory of communication to quantify the complexity of nonlinear physical dynamics. They introduced a key invariant of system randomness—the entropy rate—to measure the rate at which chaotic dynamical systems and stochastic processes produce information. The tension between these two threads serves to raise a natural and complementary question, How much memory does a system require to produce its information? Introduced in the 1980s as a complement to the entropy rate, the statistical complexity measures this structural resource. It quantifies how much of a system's past behavior is transmitted to its future through the present. Taken at face value, though, Blackwell's infinite complexity makes determining both entropy rate and statistical complexity very challenging. We show that the difficulties can be overcome and present practical algorithms for their

estimation. Somewhat surprisingly, the solution involves concepts from widely different areas—hidden Markov models, iterated function systems, information and dimension theories, and the spectrum of Lyapunov characteristic exponents.

I. INTRODUCTION

A delicate symbiotic relationship lives at the heart of highly complex systems—the intricate patterns these systems generate arise through an interplay between determinism and stochasticity. Despite progress identifying and measuring degrees of randomness and unpredictability, basic questions about this state of affairs remain. Specifically, how do we quantify “structure”? Can we detect system's emergent patterns and quantify their organization?

Clearly, posing these questions and developing the tools to answer them required, over the recent decades, integrating Turing's computation theory,^{1–3} Shannon's information theory,⁴ and Kolmogorov's dynamical systems theory.^{5–9} Together, they highlighted the central role that *information*—its generation, transmission, and

storage—plays in the organization and functioning of complex systems. Drawing from this convergence, *computational mechanics*¹⁰ introduced a definition of the structural complexity of stochastic processes—the *statistical complexity*, which measures the number and distribution of optimally predictive features.

Answers to the randomness-structure dichotomy have been carefully detailed and successfully implemented for stochastic processes that can be optimally predicted with *countably many* predictive features.^{11–13} However, there are complex systems arising in many engineering, physical, and biological systems^{14–18} that require an infinite number of predictive features. Somewhat soberingly, these truly complex systems are implicated in many natural phenomena, from the geophysics of earthquakes¹⁹ and physiological measurements of neural avalanches²⁰ to semantics in natural language²¹ and cascading failures in power transmission grids.²²

As first established by Blackwell in the 1950s,¹⁴ calculating the Shannon entropy rate of even apparently simple processes—such as those generated by a discrete time, N -state hidden Markov chain (HMC)—requires tracking an uncountably infinite set of *belief distributions* over the HMC's states. The following establishes that optimally predicting this class of processes requires using these belief distributions as predictive features. These feature sets, living on the $(N - 1)$ -simplex, are complex: generically highly ramified and self-similar, they are the support of similarly complicated measures.²³ This renders prediction of even simply defined processes very challenging. Probing their structure is even more difficult. Nevertheless, the broad popularity and application of HMCs—not only in the study of complex systems,¹⁰ but also in coding theory,²⁴ stochastic processes,²⁵ stochastic thermodynamics,²⁶ speech recognition,²⁷ computational biology,^{28,29} epidemiology,³⁰ and finance³¹—give testimony to the ubiquity of truly complex systems, in both theory and nature.

Our recent work³² addressed this state of affairs, showing how to produce infinite sets of predictive features and accurately calculate the entropy rate for processes generated by HMCs. This gave a new and efficient tool for consistently describing the randomness of truly complex systems. However, those results did not address the complementary side of the complex-system paradox—the structural aspect of the interplay between structure and randomness. Troublingly, for processes with uncountably infinite sets of predictive features, the statistical complexity diverges, substantially circumscribing its usefulness as a metric of system organization. As we will show, quantifying the structure of these truly complex systems requires a new approach and new tools and methods.

Historically, the need for such a measure of divergent information storage and Blackwell's discovery were perhaps anticipated by Shannon's definition in the 1940s of the *dimension rate*,⁴

$$\lambda = \lim_{\delta \rightarrow 0} \lim_{\epsilon \rightarrow 0} \lim_{T \rightarrow \infty} \frac{N(\epsilon, \delta, T)}{T \log \epsilon},$$

where $N(\epsilon, \delta, T)$ is the smallest number of elements that may be chosen such that all elements of a trajectory ensemble generated over time T , apart from a set of measure δ , are within the distance ϵ of at least one chosen trajectory. This is the minimal “number of dimensions” required to specify a member of a trajectory (or message) ensemble. Unfortunately, Shannon devotes barely a paragraph to the concept, leaving it largely unmotivated and uninterpreted.

Nonetheless, we take inspiration from this to develop the *statistical complexity dimension*, the asymptotic growth rate of the statistical complexity for an uncountably infinite set of predictive features.

We conjecture that the statistical complexity dimension is the same dimension rate proposed by Shannon. However, the following goes beyond Shannon's brief mention to provide constructive and accurate methods for determining this important system invariant. Technically, the statistical complexity dimension is defined as the information dimension of the (self-similar set of) predictive states. Several distinct steps are involved. Determining this information dimension requires establishing ergodicity, calculating the Lyapunov spectrum of an HMC's mixed-state presentation, and applying a suitably modified version of the Lyapunov-information dimension conjecture from dynamical systems that connects the spectrum to the dimension.

To highlight the usefulness of these informational quantities, that otherwise appear rather abstracted from natural systems, it should be noted that the following and its predecessor³² were preceded by two companion articles that applied the theoretical results here to two, rather different, physical domains. The first analyzed the origin of randomness and structural complexity engendered by quantum measurement.³³ The second solved a long-standing problem in exactly determining the thermodynamic functioning of Maxwellian demons, also known as information engines;¹⁸ that is, the predecessor and the present development (along with a sequel to be announced in the conclusion) lay out the mathematical and algorithmic tools required to successfully analyze structure and randomness in these applied problems. Taken together, we believe that the new approach will find an even wider use than in these application areas.

In the following, we introduce a practical and calculable measure of structural complexity analogous to Shannon's dimension rate in the form of *statistical complexity dimension* d_μ . Section II recalls the necessary background in stochastic processes, hidden Markov chains, and information theory. Section III recounts mixed states and their dynamic—the mixed-state presentation—as well as the connection to iterated function systems (IFSs) previously demonstrated. The main result follows in Sec. IV, where the Lyapunov-information dimension conjecture is reviewed and updated to our needs and the statistical complexity dimension d_μ is introduced. Finally, in Sec. VI, d_μ of a three-state parameterized HMC is calculated across a wide region of the parameter space, demonstrating the insights afforded by and computational efficiency of our methods.

II. HIDDEN MARKOV PROCESSES

Our main objects of study are stochastic processes and the mechanisms that generate them—hidden Markov chains, mixed states, and the ϵ -machine. We touch on several of their important properties, including stationarity, ergodicity, randomness, and memory. The reader familiar with the previous work in this series³² may skip to Sec. II D.

A. Processes

A *stochastic process* \mathcal{P} is a probability measure over a bi-infinite chain $\dots X_{t-2} X_{t-1} X_t X_{t+1} X_{t+2} \dots$ of random variables,

each X_t denoted by a capital letter. A particular *realization* $\dots x_{t-2} x_{t-1} x_t x_{t+1} x_{t+2} \dots$ is denoted via lowercase. 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: $X_{t:t'} = X_t \dots X_{t'-1}$. \mathcal{P} 's measure is defined via the collection of distributions over blocks: $\{\Pr(X_{t:t'}) : t < t', t, t' \in \mathbb{Z}\}$.

To simplify the development, we restrict to stationary, ergodic processes: those for which $\Pr(X_{t:t+\ell}) = \Pr(X_{0:\ell})$ for all $t \in \mathbb{Z}$, $\ell \in \mathbb{Z}^+$ and for which individual realizations obey all of those statistics. In such cases, we only need to consider process's length- ℓ *word distributions* $\Pr(X_{0:\ell})$.

A *Markov process* is the one that exhibits memory over a single time step: $\Pr(X_t|X_{-\infty:t}) = \Pr(X_t|X_{t-1})$. A *hidden Markov process* is the output of a memoryless channel³⁴ whose input is a Markov process.²⁵ The central problem that concerns us arises since the output process generically is not Markov.

B. Presentations

Directly working with processes—nominally, infinite sets of infinite sequences and their probabilities—is cumbersome. Therefore, we turn to consider finitely specified mechanistic models that generate them.

Definition 1. A *finite-state edge-labeled hidden Markov chain (HMC)* consists of

1. a finite set of states $\mathcal{S} = \{\sigma_1, \dots, \sigma_N\}$,
2. a finite alphabet \mathcal{A} of k symbols $x \in \mathcal{A}$, and
3. a set of N by N symbol-labeled transition matrices $T^{(x)}$, $x \in \mathcal{A}$: $T_{ij}^{(x)} = \Pr(\sigma_j, x|\sigma_i)$.

The associated state-to-state transitions are described by the row-stochastic matrix $T = \sum_{x \in \mathcal{A}} T^{(x)}$. The *internal-state Markov chain* is given by (\mathcal{S}, T) . The asymptotic, stationary state distribution is $\pi = \{\Pr(\sigma), \sigma \in \mathcal{S}\}$ and, as a vector, is given by T 's left eigenvector normalized in probability: $\pi = \pi T$.

The stochastic process \mathcal{P} generated by an HMC is the set of emitted-symbol sequences and their probabilities that arise from specifying an initial distribution over states \mathcal{S} and following all allowed state-to-state transitions. \mathcal{P} is stationary if the initial distribution is π .

A given stochastic process can be generated by any number of HMCs. These alternative mechanisms are called process's *presentations*. We now introduce a structural property of HMCs that has important consequences in determining a process's randomness and structure.

Definition 2. A *unifilar HMC (uHMC)* is an HMC such that for each state $\sigma_i \in \mathcal{S}$ and each symbol $x \in \mathcal{A}$, there is at most one outgoing edge from state σ_i labeled with symbol x .

One consequence is that uHMC's states are *predictive* in the sense that each is a function of the prior emitted sequence—the past $x_{-\infty:t} = \dots x_{t-2} x_{t-1} x_t$. Consider an infinitely long past that, in the present, has arrived at state σ_t . For uHMCs, it is not required that this infinitely long past arrives at a unique state, but it is the case that any state arrived at by this past must have the same past-conditioned distribution of future sequences $\Pr(X_{\infty:t}|x_{-\infty:t})$. We call this deterministic relationship between the past and the future a *prediction*.

In comparison, a generative (nonunifilar) HMC may return a set of states with varying conditional future distributions upon seeing this infinite past. All that is required for accurate generation is that, if this were to be repeated many times, averaging over these conditional future distributions returns the unique conditional future distribution $\Pr(X_{\infty:t}|x_{-\infty:t})$ given by the predictive uHMC state.

Although there can be many presentations for a process \mathcal{P} , there is a canonical presentation that is unique: a process's ϵ -*machine*.¹⁰

Definition 3. An ϵ -*machine* is a uHMC with probabilistically distinct states: For each pair of distinct states $\sigma_i, \sigma_j \in \mathcal{S}$, there exists a finite word $w = x_{0:\ell-1}$ such that

$$\Pr(X_{0:\ell} = w|\mathcal{S}_0 = \sigma_i) \neq \Pr(X_{0:\ell} = w|\mathcal{S}_0 = \sigma_j).$$

When probabilistically distinct states are enforced for the uHMC, one obtains the ϵ -machine for which the predictive state itself is unique. This is in contrast to non-minimal uHMCs and nonunifilar HMCs, where an infinitely long past is not guaranteed to lead to a unique state.

A process's ϵ -machine is its optimally predictive, minimal presentation in the sense that the set \mathcal{S} of predictive states is minimal compared to all its other unifilar presentations. That said, \mathcal{S} may be finite, countably infinite, or uncountably infinite. Since they capture a process's structure and are not merely predictive, ϵ -machine's states are called *causal states*.

C. Process intrinsic randomness: HMC entropy rate

A process's intrinsic randomness is the information in the present measurement, discounted by having observed the preceding infinitely long history. It is measured by Shannon's source entropy rate.⁴

Definition 4. A process's *entropy rate* h_μ is the asymptotic average Shannon entropy per symbol,¹¹

$$h_\mu = \lim_{\ell \rightarrow \infty} H[X_{0:\ell}]/\ell, \tag{1}$$

where $H[X_{0:\ell}]$ is the Shannon entropy of block $X_{0:\ell}$,

$$H[X_{0:\ell}] = - \sum_{x_{0:\ell} \in \mathcal{A}^\ell} \Pr(x_{0:\ell}) \log_2 \Pr(x_{0:\ell}). \tag{2}$$

Given a finite-state unifilar presentation M_μ of a process \mathcal{P} , we may directly calculate the entropy rate from the uHMC's transition matrices,⁴

$$h_\mu = - \sum_{\sigma \in \mathcal{S}} \pi_\sigma \sum_{x \in \mathcal{A}} \Pr(x|\sigma) \log_2 \Pr(x|\sigma). \tag{3}$$

In general, though, for processes generated by nonunifilar HMCs, there is no such closed-form expression for the entropy rate.¹⁴ For these processes, the closed-form expression [Eq. (3)] applied to the HMC states and transition matrices substantially misestimates the generated process's entropy rate.

Addressing this nonunifilar case was the focus of our previous development.³² We showed that the entropy rate of a general HMC may be determined using its *mixed states*, reviewed shortly in Sec. III. Tracking HMC's mixed states allows one to find the

entropy rate of the generated process and so the latter's intrinsic randomness.

D. Process intrinsic structure

A process's memory is determined using its ϵ -machine. Depending on the specific need, this may be measured either in terms of the number of causal states $|\mathcal{S}|$ or the amount of historical Shannon entropy they store—that is, the *statistical complexity* C_μ .

Definition 5. A process's statistical complexity is the Shannon entropy stored in its ϵ -machine's causal states,

$$C_\mu = H[\Pr(\mathcal{S})] = - \sum_{\sigma \in \mathcal{S}} \pi_\sigma \log_2 \pi_\sigma. \tag{4}$$

A process's ϵ -machine is its smallest uHMC presentation, in the sense that both $|\mathcal{S}|$ and C_μ are uniquely minimized by a process's ϵ -machine, compared to all other unifilar presentations. Due to this, the ϵ -machine's state entropy $H[\Pr(\mathcal{S})]$ is a unique measure of a process's structural complexity.

A challenge arises similar to that encountered with determining a process's entropy rate via its nonunifilar HMC presentations: there is no closed-form expression for the generated process's C_μ . We now turn to give a constructive answer to this challenge. The preceding presentation types—MC, HMC, uHMC, and ϵ -machine—give a useful path to understanding how process's different presentations help or hinder determining process properties. The strategy in the following turns on yet another presentation type. Here on in, with nothing else said, reference to an HMC means the general case—a nonunifilar HMC.

III. OBSERVER-PROCESS SYNCHRONIZATION

Previously, we introduced mixed-state presentations of HMCs and established their equivalence to random dynamical systems known as *iterated function systems* (IFSs).³² We now briefly review this construction. (Readers familiar with the previous results may skip to Sec. IV.)

Assume that an observer has a finite HMC presentation M for a process \mathcal{P} . Consider the *observer-process synchronization problem* in which the observer attempts to determine, at each moment, M 's state from observed data.³⁵ Note that this is a dynamical problem—with every new symbol emitted, the state of M may change.

Since M is a *hidden* Markov chain, the observer cannot directly detect the state. With only knowledge of M 's structure, the observer's best guess is that the states occur according to M 's internal-state stationary distribution π . This distribution describes the proportion of time spent in each state as observation length approaches infinity; therefore, by making this our initial guess, we assume M has been generating \mathcal{P} long enough that we may ignore its initial condition. The observer then refines this guess by monitoring the output data $x_0 x_1 x_2 \dots$ that M generates and comparing this to their knowledge of M 's structure. If and when the observer knows with certainty in which state the process is, they have *synchronized* to the process.




FIG. 1. A hidden Markov chain with two states $\mathcal{S} = \{\sigma_1, \sigma_2\}$ and two emitted symbols $\mathcal{A} = \{\square, \triangle\}$. It is nonunifilar and parameterized with $s \in [1, \infty)$. It becomes unifilar in the limit of $s \rightarrow \infty$.

A. Mixed-state presentation

We formalize the observer-synchronization problem by introducing a new presentation for a process \mathcal{P} , known as the mixed-state presentation.

1. Mixed-state set

Assume we have an $N + 1$ -state HMC presentation M with k symbols $x \in \mathcal{A}$. For a length- ℓ word w generated by M , let $\eta(w) = \Pr(\mathcal{S}|w)$ be the observer's *belief distribution* as to the process's current state after observing w ,

$$\eta(w) \equiv \Pr(\mathcal{S}_\ell | X_{0:\ell} = w, \mathcal{S}_0 \sim \pi). \tag{5}$$

$X \sim \eta$ means that random variable X is distributed according to η .

The belief distributions $\eta(w)$ that an HMC can visit define its set of *mixed states*,

$$\mathcal{R} = \{\eta(w) : w \in \mathcal{A}^+, \Pr(w) > 0\},$$

where \mathcal{A}^+ indicates the set of all words with positive length. Generically, the mixed-state set \mathcal{R} for an N -state HMC is infinite, even for finite N .¹⁴ Figure 2 shows a case where the HMC generates the Cantor set as its mixed-state set.

When observing an N -state machine, the vector $\langle \eta(w) |$ lives in the $(N-1)$ -simplex Δ^{N-1} , the set such that

$$\{\eta \in \mathbb{R}^N : \langle \eta | \mathbf{1} \rangle = 1, \langle \eta | \delta_i \rangle \geq 0, i = 1, \dots, N\},$$

where $\langle \delta_i | = [0 \ 0 \ \dots \ 1 \ \dots \ 0]$ and $|\mathbf{1}\rangle = [1 \ 1 \ \dots \ 1]^T$. We use this notation for components of the mixed-state vector η to avoid confusion with temporal indexing.

2. Mixed-state dynamic

The set of mixed states itself does not constitute a presentation. We must augment it with the state dynamic, which gives the transition probabilities between mixed states. Let the initial mixed state be the invariant probability π over M 's states; therefore, $\langle \eta_0 | = \langle \pi |$. In the context of the mixed-state dynamic, mixed-state subscripts denote time. The probability of transitioning from $\langle \eta(w) |$ to $\langle \eta(wx) |$ on observing symbol x follows from Eq. (5) immediately; we have

$$\Pr(\eta(wx) | \eta(w)) = \Pr(x | \mathcal{S}_\ell \sim \eta(w)).$$

This defines the mixed-state transition dynamic \mathcal{W} .

Given M 's symbol-labeled transition matrices, we can specify the mixed-state dynamic \mathcal{W} in a closed form. First, the probability




FIG. 2. Constructing the mixed-state presentation of the two-state nonunifilar HMC shown in (a): The invariant state-distribution $\pi = (1/2, 1/2)$. It becomes the first mixed state η_0 used in (b) to calculate the next set of mixed states—those having seen past words of length one. (c) In this example, infinitely many mixed states are generated and \mathcal{R} is the middle- $1/s$ Cantor set. Color is used to indicate the relative closeness of each mixed state to the original states in (a). From η_0 , one would need to see a word of infinite Δ s to reach σ_0 . In (d), the mixed states for $s = 3$ are pictured on the 1-simplex—the unit interval from $\eta = (0, 1)$ to $\eta = (1, 0)$. In this representation, the relationship to the Cantor set is visually clear.

Downloaded from http://pubs.aip.org/cha/article-pdf/doi/10.1063/5.0050460/14634261/083114_1_online.pdf

of generating symbol x when in a mixed state η is

$$\Pr(x|\eta) = \langle \eta | T^{(x)} | \mathbf{1} \rangle, \tag{6}$$

where $T^{(x)}$ is the symbol-labeled transition matrix associated with the symbol x . Now, given a mixed state at time t , we may calculate the probability of seeing each $x \in \mathcal{A}$. Upon seeing symbol x , the current mixed state $\langle \eta_t |$ is updated according to

$$\langle \eta_{t+1,x} | = \frac{\langle \eta_t | T^{(x)} | \mathbf{1} \rangle}{\langle \eta_t | T^{(x)} | \mathbf{1} \rangle}. \tag{7}$$

Equation (7) tells us that, by construction, the MSP is unifilar since each possible output symbol uniquely determines the next

(mixed) state. Taken together, Eqs. (6) and (7) define the mixed-state transition dynamic \mathcal{W} as

$$\begin{aligned} \Pr(\eta_{t+1}, x | \eta_t) &= \Pr(x | \eta_t) \\ &= \langle \eta_t | T^{(x)} | \mathbf{1} \rangle \end{aligned}$$

for all $\eta \in \mathcal{R}, x \in \mathcal{A}$.

Together, the mixed states and their dynamic define an HMC that is unifilar by construction. This is a process's *mixed-state presentation* (MSP) $\mathcal{U}(\mathcal{P}) = \{\mathcal{R}, \mathcal{W}\}$.

B. Constructing the mixed-state presentation

To explicitly find the MSP $\mathcal{U} = \{\mathcal{R}, \mathcal{W}\}$ for a given HMC M , we apply *mixed-state construction*:

1. Set $\mathcal{U} = \{\mathcal{R} = \emptyset, \mathcal{W} = \emptyset\}$.
2. Calculate M 's invariant state distribution: $\pi = \pi T$.
3. Take η_0 to be $\langle \pi \rangle$ and add it to \mathcal{R} .
4. For each current mixed state $\eta_t \in \mathcal{R}$, use Eq. (6) to calculate $\Pr(x|\eta_t)$ for each $x \in \mathcal{A}$.
5. For $\eta_t \in \mathcal{R}$, use Eq. (7) to find the updated mixed state $\eta_{t+1,x}$ for each $x \in \mathcal{A}$ with $\Pr(x|\eta_t) > 0$.
6. Add η_t 's transitions to \mathcal{W} and each $\eta_{t+1,x}$ to \mathcal{R} .
7. For each new η_{t+1} , repeat steps 4–6 until no new mixed states are produced.

This algorithm need not terminate, as shown in Fig. 2, which depicts the MSP construction for the HMC in Fig. 1. However, it can terminate for HMCs described by finite-state ϵ -machines. When working with HMCs for which the algorithm does not terminate, one must impose a limit on the number of generated mixed-states, effectively setting a level of approximation for \mathcal{R} .

One may ask, given that mixed-state construction returns a unifilar HMC of the underlying process, is the MSP the same as the ϵ -machine? It is not guaranteed to be so, as indeed is the case in Fig. 2. While the *Cantor machine* there generates an uncountably infinite set of mixed states, the underlying process's ϵ -machine is a single-state fair coin—which has a one-state HMC as its ϵ -machine. We see this by noting that the symbol-branching probabilities depicted in Fig. 2(c) are identical for every generated mixed state. This may seem a cause for concern. Perhaps we overestimated the necessary state size for the process in Fig. 2 by an infinite factor. However, this case is both rare and easy to check for, as discussed in Appendix B. By applying a simple check on the uniqueness of mixed states as they are constructed, we can confirm if the MSP is the underlying process's ϵ -machine. Unless otherwise noted, we take this to hold.

Although the most common purpose of applying mixed-state construction is to *unifilarize* an HMC, we may find the MSP of uHMCs as well. The MSPs of unifilar presentations are interesting and contain additional information beyond the initial unifilar presentation. For example, they typically contain transient causal states, and these are employed to calculate many complexity measures that track convergence statistics.³⁶

The following narrows the focus to mixed-state presentations of nonunifilar HMCs, which typically have an infinite mixed-state set \mathcal{R} . By way of emphasizing a principle result—and Blackwell's original point¹⁴—applying mixed-state construction even to nominally simple, finite-state, finite-alphabet nonunifilar HMCs results in an explosion of mixed states. Figure 3 gives three examples of MSPs with fractal mixed-state sets \mathcal{R} , each generated by a three-state nonunifilar HMC.

C. MSP as an IFS

Specifying MSP construction in this way reveals that generating mixed states is a type of random dynamical system known as a *place-dependent iterated function system* (IFS).³² For finite k and space Δ , place-dependent IFSs are characterized by a set of *mapping*

functions,

$$\{f^{(x)} : \Delta \rightarrow \Delta \mid x \in \{0, 1 \dots k\}\},$$

and associated *probability functions*,

$$\{p^{(x)} : \Delta \rightarrow [0, 1] \mid x \in \{0, 1 \dots k\}\}.$$

A place-dependent IFS generates a stochastic process over $\eta \in \Delta^N$ as follows: Given an initial position $\eta_0 \in \Delta^N$, the probability distribution $\{p^{(x)}(\eta_0) : x = 1, \dots, k\}$ is sampled. According to the sample x , apply $f^{(x)}$ to map η_0 to the next position $\eta_1 = f^{(x)}(\eta_0)$. Resample x from the distribution and continue, generating $\eta_0, \eta_1, \eta_2, \dots$.

An N -state HMC's associated mixed-state presentation defines a place-dependent IFS over the $N - 1$ simplex, with each symbol-labeled transition matrix $T^{(x)}$ defining a mapping function [Eq. (7)] and associated probability function [Eq. (6)]. Our previous results showed that an IFS defined by an ergodic HMC generates an ergodic mixed-state process and has a unique attractor—the set of mixed states \mathcal{R} .³² Additionally, this attractor has a unique, attracting, invariant measure known as the *Blackwell measure* $\mu_B(\mathcal{R})$.

IV. STRUCTURE OF INFINITE-STATE PROCESSES

Our prior development showed how to use the MSP to find a process's intrinsic randomness—in the form of Shannon entropy rate h_μ .³² Our goal here is to complement the measure of randomness with a measure of structure.

Recall that for processes generated by finite unifilar HMCs, a unique minimal presentation—the ϵ -machine—exists. Due to the latter's minimality, for such processes, we can then define the statistical complexity C_μ using Eq. (4). In contrast, general (nonunifilar) HMCs have no such canonical minimal presentation. Therefore, the Shannon entropy of their presentation states is ambiguous. However, Sec. III demonstrated how to find the mixed-state presentation $\mathcal{U}(M)$ of a nonunifilar HMC M , where operation unifilarizes it. A heavy cost is levied—explosion of the mixed-state set. Nonetheless, using the MSP to develop a measure of structure for processes generated by HMCs is naturally appealing since the MSP is unifilar and unique.³² Indeed, in most cases, the MSP is the ϵ -machine of the underlying process and so minimal (see Appendix B). When not, we can minimize the MSP by merging equivalent mixed states as we construct them.

However, the naïve approach of simply measuring structure with statistical complexity C_μ introduces a problem: the statistical complexity diverges for an HMC with an uncountably infinite state set \mathcal{R} . In general, MSPs of HMCs are uncountably infinite state, precluding distinguishing them via C_μ . This being said, it is clear visually from Fig. 3 that HMCs with uncountably infinite-state spaces still have significant and distinct structures. We wish to find a way to measure and distinguish such a structure. For this, we take inspiration from Shannon's dimension rate⁴ and call on a familiar tool.

A fractal dimension measures the rate at which a chosen size metric of a set diverges with the scale at which the set is observed.^{37–41} A fractal dimension is also useful to probe the “size” of objects when cardinality is not informative. For example, the mixed-state presentation, generically, has an uncountable infinity of causal states. That observation is far too coarse, though, to distinguish the clearly

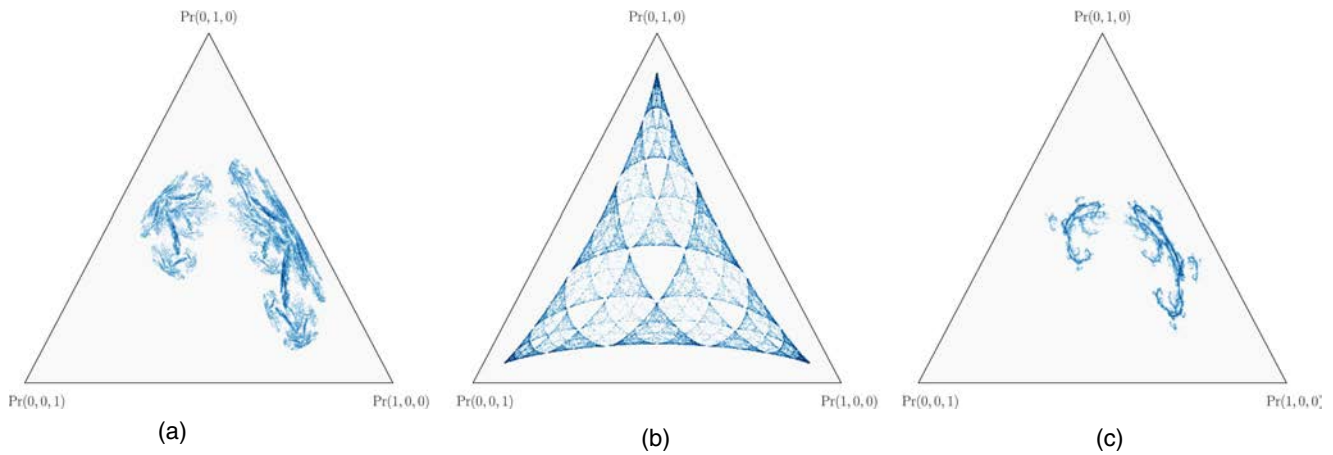


FIG. 3. Simple HMCs generate MSPs with a wide variety of structures, many fractal in nature. Each subplot displays 10^4 mixed states of a different, highly nonuniform three-state hidden Markov chain. The HMCs themselves are specified in Appendix A. (a) Mixed state attractor for a three-state “alpha” machine. (b) Mixed-state attractor for a three-state machine from Eq. (16), with $\alpha = 0.6$ and $x = 0.1$. (c) Mixed-state attractor for a three-state “beta” machine.

distinct mixed-state sets \mathcal{R} in Fig. 3. Each is uncountably infinite, but the \mathcal{R} ’s geometries differ. Determining their fractal and other dimensions will allow us to distinguish them and allow us to introduce additional insights into the original process’s intrinsic information processing.

A. Dimensions

Consider the mixed-state set \mathcal{R} on the simplex for an N -state HMC M that generates a process \mathcal{P} . We consider two types of dimension for \mathcal{R} : the Minkowski–Bouligand or box-counting dimension, often simply called the fractal dimension, and the information dimension.

To calculate the first, coarse-grain the N -simplex with evenly spaced subsimplex cells of side length ϵ . Let $\mathcal{F}(\epsilon)$ be the set of cells that encompass at least one mixed state. Then, \mathcal{R} ’s *box-counting dimension* is

$$d_0(\mathcal{R}) = - \lim_{\epsilon \rightarrow 0} \frac{\log |\mathcal{F}(\epsilon)|}{\log \epsilon}, \tag{8}$$

where $|C|$ is the size of set C .

The information dimension tracks how the Blackwell measure $\mu_B(\mathcal{R})$ scales with ϵ . Let each cell in $\mathcal{F}(\epsilon)$ be a state and approximate the dynamic over $\mathcal{U}(M)$ by grouping all transitions to and from states encompassed by the same cell. This results in a finite-state Markov chain that generates an approximation of the original mixed-state process and has a stationary distribution $\mu(\mathcal{F}(\epsilon))$. Then, $\mu_B(\mathcal{R})$ ’s *information dimension* is

$$d_1(\mu_B(\mathcal{R})) = \lim_{\epsilon \rightarrow 0} \frac{H_\mu[\mathcal{F}(\epsilon)]}{\log \epsilon}, \tag{9}$$

where $H_\mu[\mathcal{F}(\epsilon)] = - \sum_{C_i \in \mathcal{F}(\epsilon)} \mu(C_i) \log \mu(C_i)$ is the Shannon entropy over the set $\mathcal{F}(\epsilon)$ of cells that cover attractor \mathcal{R} with respect to μ .

B. Dimensions and scaling of HMCs

These dimensions give two complementary resource-scaling laws for HMC-generated processes. Rearranging Eq. (8), we see that the number of mixed states in our finite-state approximation to $\mathcal{U}(M)$ scales algebraically with \mathcal{R} ’s box-counting dimension,

$$|\mathcal{F}(\epsilon)| \sim \epsilon^{-d_0(\mathcal{R})}. \tag{10}$$

In other words, for an uncountably infinite MSP, the exponential growth rate of mixed states is $d_0(\mathcal{R})$.

Similarly, the entropy of the Blackwell measure scales with the information dimension. Rearranging Eq. (9) shows that the state entropy of the finite-state approximation to $\mathcal{U}(M)$ scales logarithmically with \mathcal{R} ’s information dimension with respect to the Blackwell measure,

$$H_\mu[\mathcal{F}] \sim d_1(\mu_B) \cdot \log \epsilon. \tag{11}$$

As $\epsilon \rightarrow 0$, $|\mathcal{F}|$ and $H_\mu(\mathcal{F})$ diverge and d_0 and d_1 are the divergence rates, respectively. The remainder focuses on d_1 as applied to the ϵ -machine for which it describes the rate of divergence of statistical complexity C_μ .

V. STATISTICAL COMPLEXITY DIMENSION

We refer to the information dimension $d_1(\mu)$ of the ϵ -machine the *statistical complexity dimension* d_μ . Applying d_1 to the Blackwell measure, $\mu_B(\mathcal{R})$ gives the rate of divergence of C_μ as one constructs increasingly better finite-state approximations to the infinite-state ϵ -machine. In this way, d_μ describes the divergence of memory resources when attempting to optimally predict a process that requires an uncountably infinite number of predictive features. This is a unique, minimal description of the process’s structural complexity. This solves the challenge posed in the introduction: quantifying structure for a broad class of truly complex systems.

When a process may be optimally predicted with a finite number of predictive features, the statistical complexity dimension vanishes. In this case, the more relevant complexity measure is the original ϵ -machine statistical complexity C_μ , which is finite. When a process requires uncountably infinite causal states but may be generated with a minimal N -state (nonunifilar) HMC, the statistical complexity is less than or equal to $N - 1$. This is the associated IFS's *embedding dimension* since the mixed states lie in a space of dimension $N - 1$.

Unfortunately, directly calculating the information dimension using Eq. (9)—and, therefore, calculating the statistical complexity dimension d_μ —is nontrivial, as it requires estimating a fractal measure. Fortunately, to calculate the d_μ of the mixed-state attractor \mathcal{R} , we can leverage the associated generating dynamical system (see Sec. III C).

A. Dimension from dynamical (in)stabilities

We can link the information dimension of an MSP's mixed-state set \mathcal{R} to the stability properties of the associated IFS. This starts with determining the local time-average stability and instability of orbits within an attractor via the *spectrum of Lyapunov characteristic exponents* $\Gamma = \{\lambda_1, \dots, \lambda_N : \lambda_i \geq \lambda_{i+1}\}$.^{42,43} Individual LCEs λ_i measure the average local growth or decay rate of orbit perturbations. The net result is a list of quantities that indicate long-term orbit instability ($\lambda_i > 0$) and orbit stability ($\lambda_i < 0$) in complementary directions. Usefully, their sum gives the net state-space divergence—volume loss for dissipative systems. The sum of the positive LCEs is the dynamical system's entropy rate^{41,42,44–46}—the net information generation.

To motivate our present use of the Lyapunov spectrum, it will help to develop a simple intuition for how the LCEs relate to dimension. Both of our previously defined dimensional quantities—Eqs. (9) and (8)—depended on the growth rate of cells needed to cover the attractor as we take the side length of the cells to zero. First, imagine the attractor of a two-dimensional map with LCEs $\lambda_1 > 0 > \lambda_2$ and whose state space is covered with equally spaced squares of side length ϵ . After iterating the map q times, for ϵ small enough, the local action of the map is approximately linear. From the LCE definition, this means that it takes the initial square cells to rectangles of average length $(e^{\lambda_1 q})\epsilon$ and average width $(e^{\lambda_2 q})\epsilon$. Now, consider covering the attractor with a new set of squares of side length $(e^{\lambda_2 q})\epsilon$. We can see by inspection that this requires roughly $e^{(\lambda_1/\lambda_2)q}$ squares per rectangle. In this way, the Lyapunov exponents relate to the scalings measured by dimension quantities.⁴⁷

Indeed, this relationship has resulted in the definition of a *Lyapunov dimension* d_Γ in terms of the spectrum Γ ,⁴⁸

$$d_\Gamma = \begin{cases} k + \frac{\Lambda(k)}{|\lambda_{k+1}|}, & \Lambda(N) < 0, \\ N, & \Lambda(N) \geq 0, \end{cases} \quad (12)$$

where we introduce the Lyapunov spectrum partial sum $\Lambda(m) = \sum_{i=1}^m \lambda_i$ and $k = 0, 1, 2, \dots, N$ is the largest index for which $\Lambda(k) > 0$. If $\lambda_1 < 0$, we take $\Lambda(0) = 0$ and $d_\Gamma = 0$. Its name helps distinguish the conditions under which the relationships between the various dimensions actually hold. (There are many

conditions and system classes that this summary necessarily leaves out.)

It has been shown that for any ergodic, invariant probability measure μ ,

$$d_1(\mu) \leq d_\Gamma,$$

with equality when μ is a Sinai–Bowen–Ruelle (SBR) measure.⁴⁹ Furthermore, it was conjectured that for “typical dynamical systems,” the Lyapunov dimension d_Γ equals the information dimension d_1 .⁴⁸ This remarkable relationship directly relates a system's dynamics to the geometry and natural measure of its attractor. Additionally, and usefully, Eq. (12) gives us a tractable method to find the information dimension of an attractor generated by a dynamical system when the equality holds and an upper bound when it does not.

B. Calculating statistical complexity dimension

We now have in hand two important pieces. The first is the definition of statistical complexity dimension d_μ as the information dimension of an ϵ -machine. Second, we have a bound on the information dimension of an attractor, given knowledge of the generating system's dynamics. To complete our picture, we now address the final puzzle piece.

As Sec. V A discussed, there is a direct relationship between the information dimension of a chaotic attractor and the dynamics of the system to which the attractor belongs. Furthermore, as discussed in Sec. III C, every HMC has an associated random dynamical system—the iterated function system (IFS)—which has the HMC's set of mixed states \mathcal{R} as its unique attractor. Combining these two facts allows us to exactly calculate d_μ in many cases of interest.

An advantage of working with IFSs defined by HMCs is a clean division exhibited by their Lyapunov spectrum. It has been shown that the entropy rate of the generated process is equivalent to the largest Lyapunov exponent;^{50,51} conceptually, this appends h_μ to the Lyapunov spectrum as λ_0 . Calculating h_μ was the main topic of the prequel to the present work.³² Furthermore, due to the contractivity of the IFS mapping functions, all other Lyapunov exponents (λ_i , $i = 1, 2, \dots, N$) will necessarily be negative. For a review of calculating the Lyapunov exponents for IFSs, see Appendix E.

Therefore, the Lyapunov dimension of an IFS is

$$\tilde{d}_\Gamma = \begin{cases} k + \frac{\Lambda(k) + h_\mu}{|\lambda_{k+1}|}, & -\Lambda(N) > h_\mu, \\ N, & -\Lambda(N) \leq h_\mu, \end{cases} \quad (13)$$

where $k = 0, 1, 2, \dots, N$ is now the largest index for which $-\Lambda(k) < h_\mu$. In contrast to Eq. (12), in this case, $\Lambda(m) < 0$ for all $m = 1, 2, \dots, N$. We still take $\Lambda(0) = 0$, but the dimension does not necessarily go to zero: $\tilde{d}_\Gamma > 0$ when $h_\mu > 0$. Readers familiar with the Lyapunov dimension should take care as this expression effectively re-indexes the traditional presentation of d_Γ as given in Eq. (12).

Under specific technical conditions to be discussed shortly, the IFS d_Γ is exactly equal to the information dimension of the IFS's attractor and, therefore, is d_μ .⁵² In general, relaxing those conditions,

\tilde{d}_Γ upper bounds the statistical complexity dimension,

$$\tilde{d}_\Gamma \geq d_\mu. \tag{14}$$

Assembling these pieces together determines the basic algorithm to calculate (or bound) the statistical complexity dimension:

1. For an N -state HMC M with $|\mathcal{A}| = k$, write down the associated IFS with k symbol-labeled mapping functions $f^{(x)}$ and probability functions $p^{(x)}$.
2. Calculate the entropy rate h_μ using the Blackwell limit (see Ref. 32).
3. Calculate the negative Lyapunov exponents $\{\lambda_1, \dots, \lambda_{N-1}\}$ (see Appendix E).
4. Compute the Lyapunov dimension d_Γ using Eq. (13).

As mentioned, in specific cases, the Lyapunov dimension is exactly equal to the statistical complexity dimension, and our task is complete. However, there are major technical concerns with when we have only the bound in Eq. (14) and with its tightness then.

C. The overlap problem

A subtle disadvantage of working with IFSs is a direct result of the stochastic nature of them as random dynamical systems. We must consider the *overlap problem*, which concerns the ranges of the symbol-labeled mapping functions $f^{(x)}$, illustrated in Fig. 4. Specifically, the problem means that we must distinguish between IFSs that meet the open set condition^{53,54} and those that do not.

Definition 6. An iterated function system with mapping functions $f^{(x)} : \Delta \rightarrow \Delta$ satisfies the open set condition (OSC) if there exists a nonempty open set $U \in \Delta$ such that for all $x, x' \in \mathcal{A}$,

$$f^{(x)}(U) \cap f^{(x')}(U) = \emptyset, \quad x \neq x'.$$

IFSs that meet the OSC are nonoverlapping IFSs.

When the images of the symbol-labeled mappings overlap, the inequality in Eq. (14) is strict. To briefly outline the consequences, for an overlapping IFS, the entropy rate h_μ does not accurately capture state-space expansion. This causes the IFS d_Γ [Eq. (13)] to overestimate the information dimension. As a rule of thumb, the

degree to which the mappings overlap determines the magnitude of the bound's error. The impact of overlaps is significant. It is explored both in Sec. VI, where we calculate the statistical complexity dimension for HMCs with and without overlap, as well as in the sequel, which diagnoses the problem's origins and provides an algorithmic solution.

For now, to give a workable approach, we simply introduce two extra steps to the d_μ algorithm from Sec. V B:

1. Determine if the open set condition is met using the mapping functions $f^{(x)}$.
2. If the OSC is not met, estimate the degree of overlap to determine the closeness of the bound on d_μ .

D. Statistical complexity dimension for processes generated by two-state HMCs

Finally, we analyze two-state HMCs for which Eq. (13) simplifies significantly. When there is no overlap in the maps, we have exact results for d_μ .

For two-state nonunifilar HMCs, the mixed-state set lives on the 1-simplex Δ^1 —the unit interval from $\eta = (0, 1)$ to $\eta = (1, 0)$. Mixed states $\eta \in \mathcal{R}$ and the dynamic on them exist in a one-dimensional space, and, thus, there is a single negative Lyapunov exponent $\lambda_1 < 0$.

In this case, the calculation of the negative Lyapunov exponent is particularly direct since the maps are all one-dimensional. The negative Lyapunov exponent for an one-dimensional map $\eta_{n+1} = f(\eta_n)$ is

$$\lambda(\eta_0) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=0}^{N-1} \log \left| \frac{df(\eta_i)}{d\eta} \right|$$

for an orbit starting at η_0 . For an IFS with a set of mapping functions $\{f^{(x)}\}$, we find λ as the weighted average of the Lyapunov exponents of each map,

$$\lambda_\mu = \int \sum_x p^{(x)}(\eta) \log \left| \frac{df^{(x)}(\eta)}{d\eta} \right| d\mu,$$

where μ is the IFS's Blackwell measure. We can apply ergodicity to transform this into a summation over time for ease of calculation.

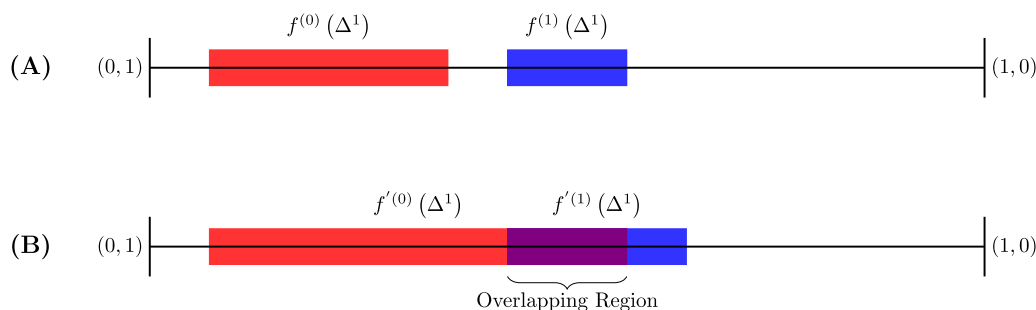


FIG. 4. Overlap problem on the 1-simplex Δ^1 : Two distinct IFSs are considered, each with two mapping functions. The images of the mapping functions over the entire simplex are depicted in red and blue. (a) Images of the mapping functions $f^{(0)}$ and $f^{(1)}$ do not overlap—every mixed state $\eta_i \in \mathcal{R}$ has a unique pre-image. (b) Images of the mapping functions overlap (purple overlapping region)—there exist $\eta_1, \eta_2 \in \mathcal{R}$ such that $f^{(0)}(\eta_1) = f^{(1)}(\eta_2) = \eta_3$. This case is an *overlapping IFS*.

If a two-state HMC has an MSP with an uncountable infinity of mixed states and its corresponding IFS satisfies the OSC, there is a simple relationship between the entropy rate, the Lyapunov exponent, and the statistical complexity dimension. This is given by

$$d_\mu(\mu) = -\frac{h_\mu}{\lambda_\mu}, \tag{15}$$

recalling that $h_\mu > 0$ so that the dimension is always positive. Failing OSC, this ratio is an upper bound on the dimension of the measure μ .⁵⁵ For a discussion of the intuition behind this formula, see Appendix C.

VI. MULTI-STATE HMC EXAMPLES

Notably, the Lyapunov dimension [Eq. (15)] for more-than-two-state HMCs is easily shown to be correct when the maps are similitudes and the probability functions are constant. The latter is seen, for example, with the Sierpinski triangle, as discussed in Appendix D.

However, we are generally interested in multi-state HMCs that do not produce perfectly self-similar fractals. Furthermore, we are often interested in considering physical systems described by parameterized HMCs, such as those that arose in the two pre-quals on quantum measurement processes and information engine functionality.^{18,33} In such cases, an HMC determined by an application may meet the OSC in some regions of the parameter space and fail to do so in others. Let us consider an HMC that spans the breadth of these possible behaviors, from zero overlap to complete overlap. This will demonstrate the range of applicability of our statistical complexity dimension algorithm. Ultimately in a sequel,⁵⁶ it leads to a practical algorithm and a conjecture for general HMCs.

Consider the following HMC with three symbols and three states:

$$T^\square = \begin{pmatrix} \alpha y & \beta x & \beta x \\ \alpha x & \beta y & \beta x \\ \alpha x & \beta x & \beta y \end{pmatrix}, \quad T^\Delta = \begin{pmatrix} \beta y & \alpha x & \beta x \\ \beta x & \alpha y & \beta x \\ \beta x & \alpha x & \beta y \end{pmatrix},$$

and

$$T^\circ = \begin{pmatrix} \beta y & \beta x & \alpha x \\ \beta x & \beta y & \alpha x \\ \beta x & \beta x & \alpha y \end{pmatrix}, \tag{16}$$

with $\beta = (1 - \alpha)/2$ and $y = 1 - 2x$. By inspection, we see that α takes on any value from 0 to 1 and x may range from 0 to 1/2.

Figure 5 shows how the MSP attractors change across the (α, x) parameter space. Each black dot is a generated mixed state, while the colored regions show the range of each symbol-labeled map.

For example, on one hand, in the top left corner with $\alpha = 0.01$ and $x = 0.01$, we find an attractor that fills the simplex, with moderate amounts of overlap. On the other, $\alpha = 0.79$ and $x = 0.11$ produce an attractor with no overlap and clearly defined regions.

Moreover, for any α , choosing $x = 1/3$ leads the MSP attractor to collapse to a finite three-state HMC since the symbol-labeled mapping functions become constant functions. In this case, there is no overlap, as each symbol-labeled map takes on a different constant value.

However, when $\alpha = \beta = 1/3$, all symbol-labeled mapping functions are identical. Therefore, the attractor is the single fixed point shared by all three maps—a single-state HMC. This is a case of maximal possible overlap. Along both lines in the parameter space, the MSP collapses to a finite-state HMC, so $d_\mu = 0$, by definition. However, these different mechanisms of state-collapse are relevant in calculation of d_Γ via Eq. (13).

First, consider Fig. 6(a), which illustrates the estimated area on the simplex taken up by the attractor across the parameter space. For a discussion of how the attractor area was estimated, see Appendix F. This figure matches the grid in Fig. 5: lower values of x produce larger attractors excepting the region near $\alpha = 1/3$, where the area drops to zero.

Now, the area taken up by an attractor is not a good proxy for dimension—we may have very small-in-size attractors with a large dimension. Indeed, this is the case for most values of α near 1/3, where the attractors are very small but whose $d_\mu \rightarrow 2$. However, we know from analyzing the attractor grid that for HMCs lying exactly on this line, the attractor is an instantaneously finite state. Therefore, the statistical complexity dimension d_μ must discontinuously drop to zero. However, this is not accurately reflected by d_Γ , as seen in Fig. 6(b).

That said, d_Γ clearly smoothly approaches zero as $x \rightarrow 1/3$. Along the vertical line, d_Γ is correctly exactly zero. This is the other line in the parameter space where the MSP is finite state and d_μ is analytically known to be zero. The disparity, in both correctness and continuity, is due to the different mechanisms driving the collapse noted above.

As $x \rightarrow 1/3$, the slopes of the symbol-labeled mapping functions approach zero. When $x = 1/3$, the symbol-labeled mapping functions become constant values, as reflected in the Lyapunov exponents and consequently in d_Γ . The constant functions have negative Lyapunov exponents of negative infinity, sending Eq. (13) to zero.

In contrast, along the $\alpha = 1/3$ line, the contraction in the state space is a result of the maps instantaneously sharing a fixed point. For $\alpha = 1/3 \pm \epsilon$, the attractor is not finite. In this region of the parameter space, symbol-labeled maps are not infinitely contracting; therefore, d_Γ badly overestimates d_μ along the $\alpha = 1/3$ line. This illustrates the importance of the OSC on the bound.

This poses the question: which regions in the HMC parameter space exhibit overlap IFS maps? Figure 6(d) depicts parameter space regions in terms of overlap or no overlap. Figure 6(c) shows this as a percentage of the total attractor area. For a discussion of how overlap was determined, please see Appendix F. Comparing the two, we see that there is a significant region over which overlap does exist for $x < 0.15$ and a smaller region where $x > 0.48$. However, for much of that region, the attractor’s overlap area is relatively small. As a rule of thumb, the gap between d_Γ and d_μ for the mixed-state set is determined by the percentage of the attractor that is affected by overlap. If the overlap region is relatively small, in comparison with the size of the attractor, d_Γ may be very close to d_μ .

However, if the overlap is very large, d_Γ may be a dramatic overestimation of d_μ . This occurs when $\alpha = 1/3$ and $x < 0.15$. The statistical complexity dimension d_μ vanishes; yet, the Lyapunov dimension saturates at $d_\Gamma = 2.0$.

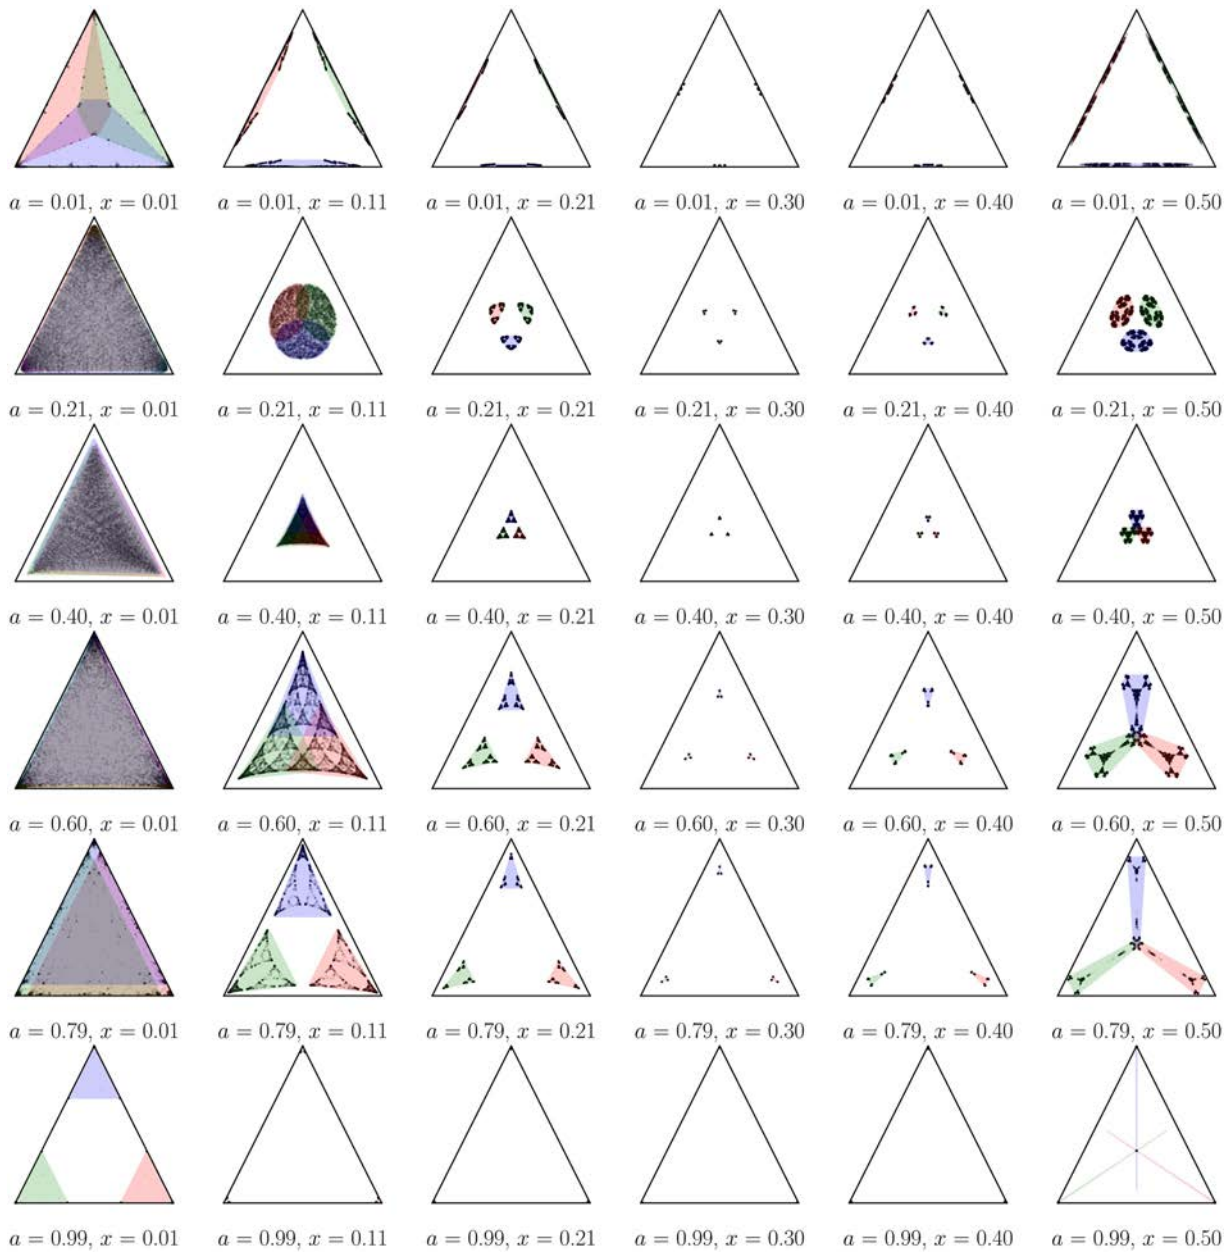


FIG. 5. Mixed-state attractors generated by a three-state HMC parameterized over $\alpha \in [0, 1]$ and $x \in [0, 0.5]$. The HMC itself is given in Eq. (16). 10^5 mixed states are plotted for each attractor, with the initial 5×10^4 states thrown away as transients. The ranges of the symbol-labeled maps are color shaded, revealing regions of their image overlap on the attractor. Comparing to Eq. (16), the red, blue, and green regions represent the images of the mapping functions defined by T^\square , T^Δ , and T° , respectively.

We also note that the mechanism driving the collapse of the MSP attractor at $\alpha = 1/3$ is a discontinuity in the parameter space, as compared to $x = 1/3$. This is because state-space collapse due to overlap requires the maps to be identical, and even minute differences in the symbol-labeled transition matrices will produce an

uncountably infinite MSP, potentially with $d_\mu = 2.0$. This encourages us to consider not just the statistical complexity dimension, but also the area of the attractor and the nearby regions in the parameter space for a clearer understanding of the underlying HMC. In this, the tools developed here, by allowing (computationally efficient)

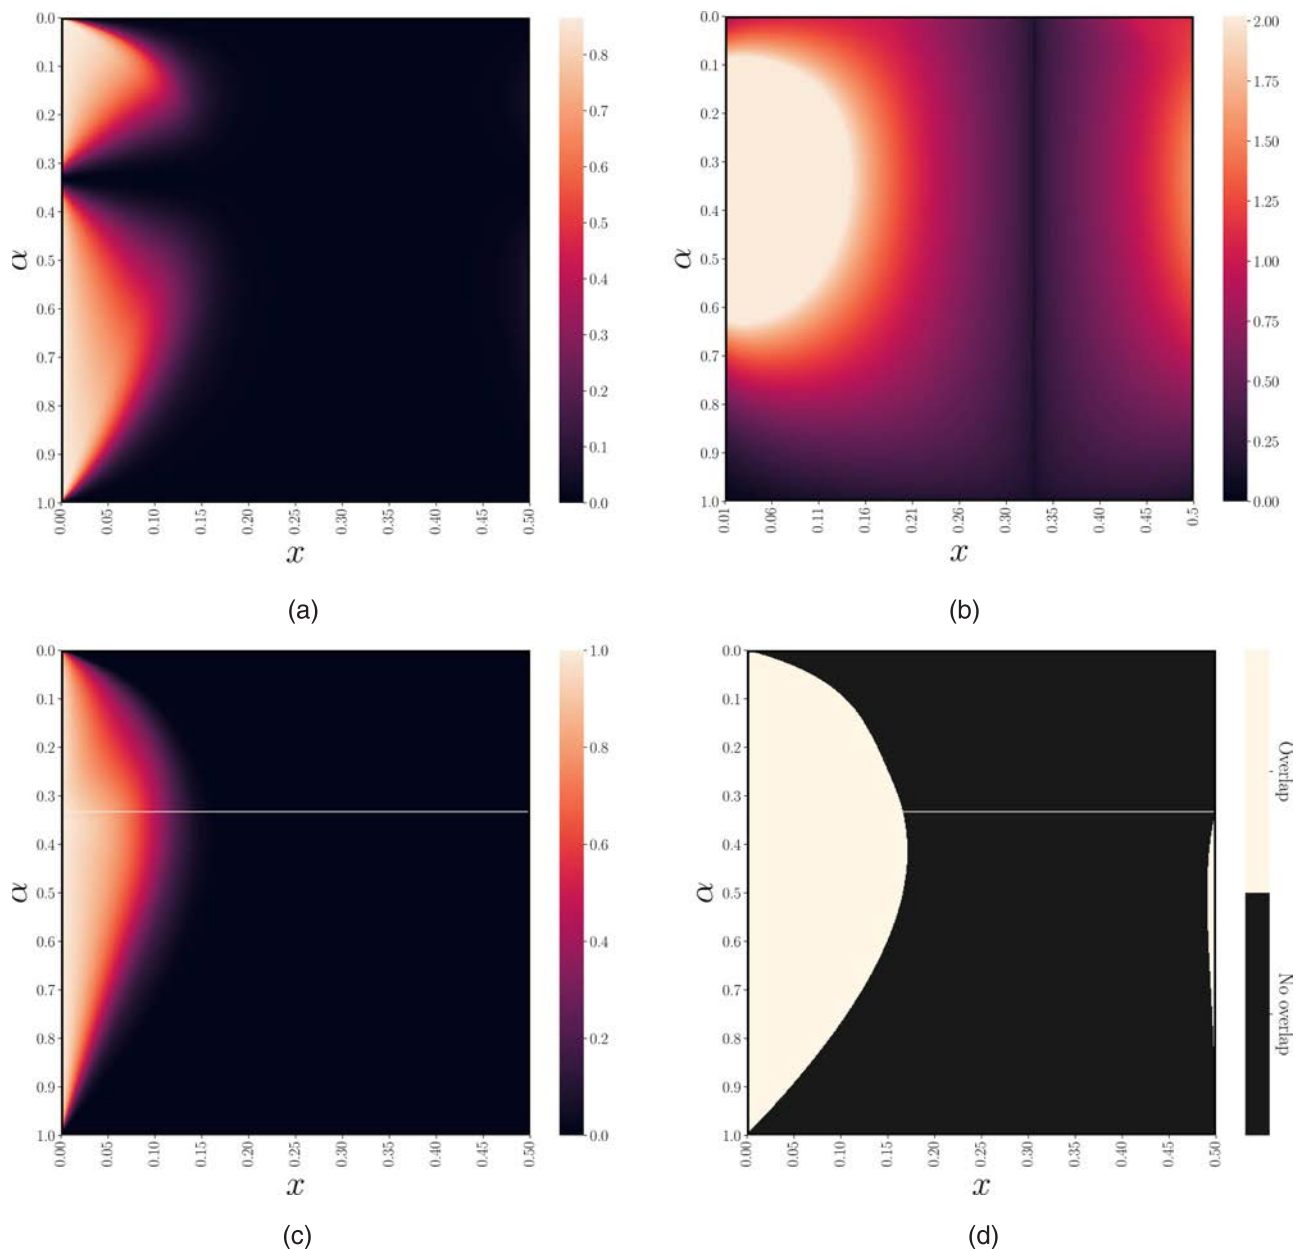


FIG. 6. Attractor area, overlap regions, and Lyapunov dimension of the mixed-state attractors shown in Fig. 5, parameterized by $\alpha = [0, 1]$ and $x = [0, 0.5]$. The HMC itself is given in Sec. VI. See Appendixes E and F for a detailed discussion of the production of these plots. (a) Mixed-state attractor area estimated across α and x . The area is minimized along lines of $\alpha = 1/3$ and $x = 1/3$, where the attractor becomes point-like. (b) Mixed-state attractor Lyapunov dimension does not detect the collapse to zero dimension along $\alpha = 1/3$, which is due to overlaps. (c) Percentage of an attractor area in which there is overlap. (d) Overlap in the attractor area. Comparing with (c), for much of this area, the overlap is very small.

surveys of large regions of the parameter space, are particularly useful.

Thus, while there are wide parameter regions in which the analysis developed here is correct and efficient, this is not the entire story. Overlap must be addressed for full generality.

Our sequel shows how to correct for overestimating a statistical complexity dimension, allowing accurate calculation across the entire parameter space.⁵⁶ However, the exploratory observations outlined here are a crucial diagnostic guide to those further developments.

VII. CONCLUSION

Our development opened by considering the challenge of quantifying the structure of complex systems. For well over a half a century, since Kolmogorov and Sinai introduced Shannon's mathematical theory of communication to dynamical systems, the entropy rate stood as the standard to quantify randomness in time series and in chaotic dynamical systems. Quantifying observable patterns remained a more elusive goal. However, with developments from computational mechanics, it has become possible to answer questions of structure and pattern, at least for stochastic processes generated by finite-state predictive machines, including the symbolic dynamics generated by chaotic dynamical systems.

To handle processes generated by finite-state nonpredictive (nonunifilar) hidden Markov chains, we developed the mixed-state presentation. This *unifilarizes* general HMCs, giving a predictive presentation that itself generates the process. However, adopting a unifilar presentation came at a heavy cost: Generically, they are infinite state and so previous structural measures diverge. Nonetheless, we showed how to work constructively with these infinite mixed-state presentations. In particular, we showed that they fall into a common class of dynamical system: The mixed-state presentation is a place-dependent iterated function system. Due to this, a number of results from dynamical systems theory can be applied to more fully describe the original stochastic process.

Previously, others considered the IFS-HMC connection.^{57,58} Complementing those efforts, we expanded the role of the mixed-state presentation to calculate entropy rate and demonstrated its usefulness in determining the underlying structural properties of the generated process. Indeed, Figs. 3 and 5 show how visually striking—and distinct—mixed-state sets generated by HMCs are.

Here, moving in a new direction beyond previous efforts, we established that the information dimension of the mixed-state attractor is exactly the divergence rate of the *statistical complexity*⁵⁹—a measure of a process's structural complexity that tracks memory. Thus, processes in this class effectively increase their use of memory, “creating” mixed or causal states, on the fly. Furthermore, we introduced a method to calculate the information dimension of the mixed-state attractor from the Lyapunov spectrum of the mixed-state IFS. In this way, we demonstrated that coarse-graining the mixed-state simplex—the previous method for studying the structure of infinite-state processes⁶⁰—can be avoided altogether. This greatly improves accuracy and computational speed and deepens our understanding of the origins of complexity in stochastic processes.

During the development, we noted several obstacles. Most importantly, the presence of *overlap* and failure to meet the open set condition causes the Lyapunov dimension to be a strict upper bound, and sometimes quite a poor one, on the statistical complexity dimension. The final work of our trilogy⁵⁶ introduces a measure of how badly the entropy rate overestimates the expansion of the mixed-state set. Combining this measure with the Lyapunov-information dimension conjecture finally yields a correct equation (13) to apply to HMCs with overlap, that is, for processes generated by all general HMCs.

To close, we note that the structural tools introduced here and the entropy-rate method introduced previously have been put to practical use in two previous works. One diagnosed the origin of

randomness and structural complexity in quantum measurement.³³ The other exactly determined the thermodynamic functioning of Maxwellian information engines¹⁸ when there had been no previous method for this kind of detailed and accurate analysis. At this point, however, we leave the full explication of these techniques and further analysis on how mixed states reveal the underlying structure of processes generated by hidden Markov chains to the sequel.⁵⁶

AUTHORS' CONTRIBUTIONS

All authors contributed equally to this work.

ACKNOWLEDGMENTS

The authors thank Sam Loomis, Sarah Marzen, Ariadna Venegas-Li, Nicolas Brodu, Alec Boyd, and Ryan James for helpful discussions and the Telluride Science Research Center for hospitality during visits and the participants of the Information Engine Workshops there. J.P.C. acknowledges the kind hospitality of the Santa Fe Institute, the Institute for Advanced Study at the University of Amsterdam, and the California Institute of Technology for their hospitality during visits. This material is based upon work supported by, or in part by, FQXi grants (Nos. FQXi-RFP-IPW-1902 and FQXi-RFP-CPW-2007) and the U.S. Army Research Laboratory and the U.S. Army Research Office (USARO) grants (Nos. W911NF-18-1-0028 and W911NF-21-1-0048).

APPENDIX A: NONUNIFILAR HMC EXAMPLES

We reproduce here the HMCs used to create Fig. 3. First, the “alpha” HMC, from Fig. 3(a), is given by

$$T^{\square} = \begin{pmatrix} 2.734 \times 10^{-2} & 0.392 & 1.924 \times 10^{-2} \\ 0.475 & 2.176 \times 10^{-2} & 2.766 \times 10^{-4} \\ 0.224 & 2.711 \times 10^{-3} & 0.236 \end{pmatrix},$$

$$T^{\Delta} = \begin{pmatrix} 1.845 \times 10^{-3} & 0.133 & 0.259 \\ 3.913 \times 10^{-2} & 0.315 & 2.789 \times 10^{-2} \\ 0.467 & 1.015 \times 10^{-2} & 4.699 \times 10^{-3} \end{pmatrix},$$

$$T^{\circ} = \begin{pmatrix} 9.782 \times 10^{-2} & 3.374 \times 10^{-2} & 3.644 \times 10^{-2} \\ 5.422 \times 10^{-2} & 6.503 \times 10^{-2} & 2.090 \times 10^{-3} \\ 5.328 \times 10^{-2} & 1.278 \times 10^{-3} & 8.778 \times 10^{-4} \end{pmatrix}.$$

Figure 3(b) is given by Eq. (16), at $\alpha = 0.6$ and $x = 0.1$. The “beta” HMC, in Fig. 3(c), is given by

$$T^{\square} = \begin{pmatrix} 5.001 \times 10^{-2} & 0.388 & 4.251 \times 10^{-2} \\ 0.464 & 4.484 \times 10^{-2} & 2.495 \times 10^{-2} \\ 0.232 & 2.720 \times 10^{-2} & 0.243 \end{pmatrix},$$

$$T^{\Delta} = \begin{pmatrix} 1.708 \times 10^{-3} & 0.123 & 0.240 \\ 3.623 \times 10^{-2} & 0.292 & 2.583 \times 10^{-2} \\ 0.432 & 9.397 \times 10^{-3} & 4.351 \times 10^{-3} \end{pmatrix},$$

$$T^{\circ} = \begin{pmatrix} 9.0576 \times 10^{-2} & 3.124 \times 10^{-2} & 3.374 \times 10^{-2} \\ 5.020 \times 10^{-2} & 6.021 \times 10^{-2} & 1.935 \times 10^{-3} \\ 4.933 \times 10^{-2} & 1.183 \times 10^{-3} & 8.127 \times 10^{-4} \end{pmatrix}.$$

Due to finite numerical accuracy, reproduction of the attractors using these specifications may differ slightly from Fig. 3.

APPENDIX B: MIXED-STATE PRESENTATION MINIMALITY

Given an HMC M , minimality of infinite-state mixed-state presentations $\mathcal{U}(M)$ is an open question. MSPs are not guaranteed to be minimal. In fact, it is possible to construct MSPs with an uncountably infinite number of states for a process that requires only one state to optimally predict, as seen with the Cantor machine process in Figs. 1 and 2. Note that while this HMC generates an uncountable number of mixed states, each one has the same emitted-symbol probability distribution, indicating that all states can be merged into a single state with no loss of predictability. Indeed, the ϵ -machine for the HMC depicted in Fig. 1 is simply the single-state Fair Coin HMC.

A proposed solution to this needless presentation verbosity is a short and simple check on *mergeability* of mixed states. This refers to any two distinct mixed states that have the same conditional probability distribution over future strings, i.e., any two mixed states η_0 and ζ_0 for which

$$\Pr(X_{0:L}|\eta_0) = \Pr(X_{0:L}|\zeta_0) \tag{B1}$$

for all $L \in \mathbb{N}^+$.

A benefit of the IFS formalization of the MSP is the ability to directly check for duplicated states and, therefore, determine if the MSP is nonminimal. We check this by considering, for an $N + 1$ state HMC M with alphabet $\mathcal{A} = \{0, 1, \dots, k\}$, the dynamic not only over mixed states, but probability distributions over symbols. Let

$$P(\eta) = (p^{(0)}(\eta), \dots, p^{(k-1)}(\eta)) \tag{B2}$$

and consider Fig. 7. For each mixed state $\eta \in \Delta^N$, Eq. (B2) gives the corresponding probability distribution $\rho(\eta) \in \Delta^k$ over the symbols $x \in \mathcal{A}$. Let M emit symbol x , then the dynamic from one such probability distribution $\rho \in \Delta^k$ to the next is given by

$$\begin{aligned} g^{(x)}(\rho_t) &= P \circ f^{(x)} \circ P^{-1}(\rho) \\ &= \rho_{t+1,x}. \end{aligned} \tag{B3}$$

From this, we see that if Eq. (B3) is invertible, $g^{(x)} : \Delta^k \rightarrow \Delta^k$ is well defined and has the same functional properties as $f^{(x)}$. In other words, in this case, it is not possible to have two distinct mixed states $\eta, \zeta \in \Delta^N$ with the same probability distribution over symbols. The probability distributions can only converge under the action of $g^{(x)}$ if the mixed states also converge under the action of $f^{(x)}$. When every mixed state has a unique outgoing probability distribution, these states are also the causal states, and the MSP is the process's ϵ -machine. Our companion work⁶¹ elaborates on this and the implications for identifying the embedding dimension of minimal generators.

APPENDIX C: CORRESPONDENCE WITH THE BAKER'S MAP

The simple dimension formula in Eq. (15) may not seem easily motivated. Particularly, considering that, in general, both positive and negative Lyapunov exponents are required to have a nontrivial attractor. However, for iterated function systems defined by ergodic HMCs, all Lyapunov exponents are negative, and the expansive role




FIG. 7. Commuting diagram for probability functions $P = \{p^{(x)}\}$, mixed-state mapping functions $f^{(x)}$, and proposed symbol-distribution mapping functions $g^{(x)}$.

played by positive Lyapunov exponents is instead played by an IFS's stochastic map selection, as measured by the entropy rate h_μ .

This is more intuitively appreciated by comparing the two-state IFS with the baker's map. Consider the baker's map,

$$\begin{aligned} x_{n+1} &= \begin{cases} x_n, & y < p \\ \frac{x_n + s_1 - 1}{s_1}, & y \geq p \end{cases} \text{ and} \\ y_{n+1} &= \begin{cases} \frac{y_n}{p}, & y < p \\ \frac{y_n - p}{p-1}, & y \geq p. \end{cases} \end{aligned}$$

It has LCE spectrum $\Lambda = \{\lambda_1, \lambda_2\}$, where

$$\begin{aligned} \lambda_1 &= p \log(p) + (1 - p) \log(1 - p), \\ \lambda_2 &= p \log(1/s_0) + (1 - p) \log(1/s_1). \end{aligned}$$

Note that $\lambda_1 > 0$ and $\lambda_2 < 0$. Then, the Lyapunov dimension is

$$d_\Gamma = 1 - \frac{\lambda_1}{\lambda_2}.$$

To compare this to an IFS, take

$$\{f(x)\} = \left\{ \frac{x_n}{s_0}, \frac{x_n + s_1 - 1}{s_1} \right\}$$

and

$$\{p(x)\} = \{p, 1 - p\}.$$

Thus, we identify the y coordinate as controlling the stochastic map choice. The dynamic over position in the y direction exactly determines the IFS entropy rate. Since the baker's map is volume preserving in y , the extra dimension always contributes a plus one in the dimension formula. In other words, the dimension along a slice of constant y equals the IFS dimension.

APPENDIX D: SIERPINSKI'S TRIANGLE

The Sierpinski triangle is a canonical Cantor set in two dimensions. An HMC that generates an MSP attractor, which is the

Sierpinski triangle, is

$$T^{(0)} = \begin{pmatrix} a & 0 & a(s-1) \\ 0 & a & a(s-1) \\ 0 & 0 & as \end{pmatrix},$$

$$T^{(1)} = \begin{pmatrix} \frac{1-as}{2} & 0 & 0 \\ \frac{(1-as)(s-1)}{2s} & \frac{1-as}{2s} & 0 \\ \frac{(1-as)(s-1)}{2s} & 0 & \frac{1-as}{2s} \end{pmatrix},$$

and

$$T^{(2)} = \begin{pmatrix} \frac{1-as}{2s} & \frac{(1-as)(s-1)}{2s} & 0 \\ 0 & \frac{1-as}{2} & 0 \\ 0 & \frac{(1-as)(s-1)}{2s} & \frac{(1-as)(s-1)}{2s} \end{pmatrix}, \tag{D1}$$

where s controls the contraction coefficient and a controls the probability of selecting the maps. This HMC produces constant probability functions,

$$p^{(0)} = as, \quad p^{(1)} = \frac{1-as}{2}, \quad \text{and} \quad p^{(2)} = \frac{1-as}{2},$$

and, therefore, linear mappings since $f^{(0)} = \frac{\langle \eta | T^{(x)} \rangle}{p^{(x)}(\eta)}$. The constant probability functions make the entropy rate trivial to calculate. The linearity of the mappings does the same for the Lyapunov exponents.

Setting $s = 2$ and $a = 1/6$ results in equal probability for all maps and gives the standard Sierpinski triangle shown in Fig. 8. In this case, the entropy rate is $h_\mu = \log_2 3$, and the Lyapunov exponents are both $-\log_2 2$. Plugging this into Eq. (13) returns the well-known fractal dimension of the Sierpinski triangle: $\log_2 3 / \log_2 2 \approx 1.585$.

APPENDIX E: LYAPUNOV EXPONENTS

A *Lyapunov characteristic exponent* for a dynamical system measures the exponential rate of separation of trajectories that begin infinitesimally close. Since, typically, the separation rate depends on the direction of the initial separation, we use a spectrum of Lyapunov exponents, with one exponent for each state-space dimension. In a chaotic dynamical system, at least one Lyapunov exponent is positive. In general, the Lyapunov exponent spectrum for an N -dimensional dynamical system with mapping $x_{n+1} = F(x_n)$ depends on the initial condition x_0 . However, here, we consider ergodic systems for which the spectrum does not.

Consider the map's *Jacobian* matrix,

$$J = \frac{\partial F}{\partial x},$$

and the evolution of vectors in the tangent space, controlled by

$$\dot{Y} = YJ,$$

where $Y(0) = \mathbb{I}_N$ and $Y(t)$ describes how an infinitesimal change in $x(0)$ has propagated to $x(t)$. Let $\{y_1, \dots, y_N\}$ be the eigenvalues of the

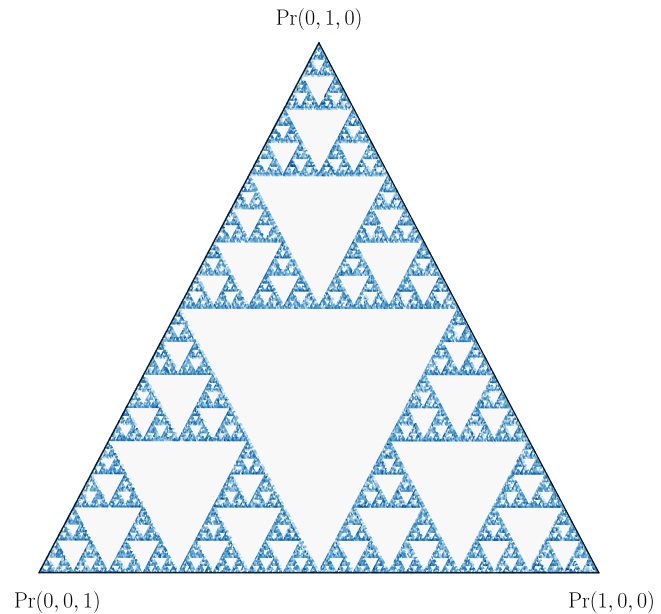


FIG. 8. Attractor of the three-state, three-symbol machine specified in Eq. (D1), with $s = 2$ and $a = \frac{1}{6}$.

matrix $Y(t)Y(t)^T$. Then, the Lyapunov exponents are

$$\lambda_i = \lim_{t \rightarrow \infty} \frac{1}{2t} \log y_i.$$

The *Lyapunov numbers* were introduced and proven to exist by Oseledets.⁶² The Lyapunov exponents are merely the logarithms of Lyapunov numbers.

The most common way of calculating an IFS's Lyapunov spectrum, employed to produce the results in Fig. 6, is the *pull-back* method. The basic idea is that for an IFS in the $N - 1$ simplex, defined by an N -state HMC, there will be $N - 1$ independent directions of contraction. These directions are represented by a coordinate frame of $N - 1$ vectors that are kept orthogonal and normalized. This coordinate frame is carried along a long orbit of the IFS. At each time step, the Jacobian is used to evolve the frame. We track the contraction rate of each vector, and this becomes the estimation of the Lyapunov exponents.

In contrast to applying this approach to deterministic dynamical systems, as is more familiar, an IFS's stochastic nature introduces additional error. Since the Jacobian varies not just across the simplex, but also for the selected maps, the orbit must be long enough to sample from the IFS attractor's distribution accurately for each possible mapping function. This being said, it has been established that the pull-back method works for IFS spectra given a sufficiently long orbit.⁶³

The prequel, on estimating the entropy rate of HMC processes, made use of error-bounding techniques from Markov chain Monte Carlo (MCMC).³² Since here we are estimating Lyapunov exponents by sampling from the Blackwell distribution, similar error-bounding techniques apply. In this analysis, there are two fundamental sources

of estimation error. First is due to *initialization bias* or undesired statistical trends introduced by the initial transient data produced by the Markov chain before it reaches the desired stationary distribution. Second, there are errors induced by *autocorrelation in equilibrium*. That is, the samples produced by the Markov chain are correlated. The consequence is that the statistical error cannot be estimated by $1/\sqrt{N}$, as done for N independent samples.

Bounding these error sources requires estimating the autocorrelation function, which can be done from long sequences of samples. If we have the nonunifilar model in hand, it is a simple matter of sweeping through increasingly long sequences of generated samples until we observe convergence of the autocorrelation function. An alternative method of approximating the infinite-state HMC with a finite-state approximation is discussed in detail in our previous work.³² The upshot is that the method here generally efficiently leads to accurate estimates of the LCE spectrum.

For completeness, we note that there are alternative methods to calculate Lyapunov exponents; see, e.g., Refs. 64 and 65. These methods may be more appropriate in specific applications. That said, the accuracy and applicability of Lyapunov exponent estimation is not the focus here.

APPENDIX F: OVERLAP ESTIMATION

To estimate the size of a mixed-state attractor and overlap of mapping functions in Fig. 6, a combination of techniques were used. We will briefly summarize the method here.

First, 250 000 different HMCs were generated using a 500×500 parameter grid over $\alpha = [0, 1]$ and $x = [0, 0.5]$. Each HMC was defined by plugging the appropriate parameter values into the symbol-labeled transition matrices in Eq. (16). From this HMC, the mapping and probability functions were defined [see Eq. (7) and (6)], producing a place-dependent IFS.

For each IFS and associated HMC, 10 000 mixed states were generated from an initial randomized state, throwing away the first 5000 as transients. Using a spatial algorithm from the SciPy Python package, a convex hull was drawn around this set of points, with a small buffer. This convex hull (the attractor “outline”) was converted into a polygon. This polygon was then evolved independently by each symbol-labeled mapping function, producing three polygons, each associated with a symbol. This may be visualized by referencing Fig. 5, where the evolved polygons are depicted on the top of the mixed-state attractor, each with a different color. We can see that the combination of these polygons must necessarily cover the attractor.

These three symbol-labeled polygons were then combined into a single polygon or multipolygon (a polygon with “holes” that are themselves polygons) using the geometry-processing module Shapely.⁶⁶ This produces a more accurate outline of the attractor than the convex hull. This process may then be repeated with the new outline for as many iterations as desired until a polygon or multipolygon that covers the mixed-state attractor with the desired level of accuracy is produced. The same result could be achieved by beginning with the entire simplex as the initial outline, without any production of mixed-states. However, the step of estimating the convex hull sharply reduces the number of required iterations and, more importantly, makes the required number more equal across the parameter space. To see this, consider that attractors taking up

less of the simplex require several iterations to converge to the small size of the attractor. By initializing with the convex hull, the process of converging to the attractor’s basic shape is skipped, and the iterations are merely refinements.

In producing Fig. 6, we found that we could determine good outlines across the parameter space by evolving the convex hull three times. To produce Fig. 6(a), the area of the resultant polygon or multipolygon was found using Shapely. To produce Figs. 6(d) and 6(e), the outline was evolved one more time by each map, and the resultant polygons and/or multipolygons were checked for intersection. For the binary overlap/no-overlap plot in Fig. 6(d), only the existence of overlap somewhere on the attractor was considered. For the percentage overlap in Fig. 6(c), the area of the total outline that was comprised of overlapping polygons—whether only two or all three—was compared to the total area. The subtlety of whether an overlap region included two or three maps was largely ignored here but will be analyzed in future explorations of the overlap problem.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- 1 A. Turing, “On computable numbers, with an application to the Entscheidungsproblem,” *Proc. Lond. Math. Soc.* **s2-42**, 230–265 (1937).
- 2 C. E. Shannon, “A universal Turing machine with two internal states,” in *Automata Studies*, Annals of Mathematical Studies Vol. 34, edited by C. E. Shannon and J. McCarthy (Princeton University Press, Princeton, NJ, 1956), pp. 157–165.
- 3 M. Minsky, *Computation: Finite and Infinite Machines* (Prentice-Hall, Englewood Cliffs, NJ, 1967).
- 4 C. E. Shannon, “A mathematical theory of communication,” *Bell Syst. Tech. J.* **27**, 379–423 (1948).
- 5 A. N. Kolmogorov, *Foundations of the Theory of Probability*, 2nd ed. (Chelsea Publishing Company, New York, 1956).
- 6 A. N. Kolmogorov, “Three approaches to the concept of the amount of information,” *Probl. Inf. Transm.* **1**, 1 (1965).
- 7 A. N. Kolmogorov, “Combinatorial foundations of information theory and the calculus of probabilities,” *Russ. Math. Surv.* **38**, 29–40 (1983).
- 8 A. N. Kolmogorov, “Entropy per unit time as a metric invariant of automorphisms,” *Dokl. Akad. Nauk. SSSR* **124**, 754 (1959); (*Russian*) *Math. Rev.* **21**(2035b).
- 9 Ja. G. Sinai, “On the notion of entropy of a dynamical system,” *Dokl. Akad. Nauk. SSSR* **124**, 768 (1959).
- 10 J. P. Crutchfield, “Between order and chaos,” *Nat. Phys.* **8**(January), 17–24 (2012).
- 11 J. P. Crutchfield and D. P. Feldman, “Regularities unseen, randomness observed: Levels of entropy convergence,” *Chaos* **13**(1), 25–54 (2003).
- 12 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 C. R. Shalizi and J. P. Crutchfield, “Computational mechanics: Pattern and prediction, structure and simplicity,” *J. Stat. Phys.* **104**, 817–879 (2001).
- 14 D. Blackwell, “The entropy of functions of finite-state Markov chains,” in *Transactions of the First Prague Conference on Information Theory, Statistical Decision Functions, Random Processes* (Publishing House of the Czechoslovak Academy of Sciences, Prague, 1957), Vol. 28, pp. 13–20.
- 15 J. P. Crutchfield and K. Young, “Computation at the onset of chaos,” in *Entropy, Complexity, and the Physics of Information*, SFI Studies in the Sciences of Complexity Vol. VIII, edited by W. Zurek (Addison-Wesley, Reading, MA, 1990), pp. 223–269.

- ¹⁶N. Travers and J. P. Crutchfield, "Infinite excess entropy processes with countable-state generators," *Entropy* **16**, 1396–1413 (2014).
- ¹⁷L. Debowski, "On hidden Markov processes with infinite excess entropy," *J. Theor. Probab.* **27**(2), 539–551 (2012).
- ¹⁸A. Jurgens and J. P. Crutchfield, "Functional thermodynamics of Maxwellian ratchets: Constructing and deconstructing patterns, randomizing and derandomizing behaviors," *Phys. Rev. Res.* **2**(3), 033334 (2020).
- ¹⁹D. L. Turcotte, *Fractals and Chaos in Geology and Geophysics*, 2nd ed. (Cambridge University Press, Cambridge, UK, 1997).
- ²⁰J. M. Beggs and D. Plenz, "Neuronal avalanches in neocortical circuits," *J. Neurosci.* **23**(35), 11167–11177 (2003).
- ²¹L. Debowski, "Excess entropy in natural language: Present state and perspectives," *Chaos* **21**(3), 037105 (2011).
- ²²I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, "Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization," *Chaos* **17**(2), 026103 (2007).
- ²³J. P. Crutchfield, "The calculi of emergence: Computation, dynamics, and induction," *Physica D* **75**, 11–54 (1994).
- ²⁴B. Marcus, K. Petersen, and T. Weissman, in *Entropy of Hidden Markov Process and Connections to Dynamical Systems*, Lecture Notes Series Vol. 385 (London Mathematical Society, 2011).
- ²⁵Y. Ephraim and N. Merhav, "Hidden Markov processes," *IEEE Trans. Inf. Theory* **48**(6), 1518–1569 (2002).
- ²⁶J. Bechhoefer, "Hidden Markov models for stochastic thermodynamics," *New J. Phys.* **17**, 075003 (2015).
- ²⁷L. R. Rabiner and B. H. Juang, "An introduction to hidden Markov models," *IEEE ASSP Mag.* **3**, 4–16 (1986).
- ²⁸E. Birney, "Hidden Markov models in biological sequence analysis," *IBM J. Res. Dev.* **45**(3.4), 449–454 (2001).
- ²⁹S. Eddy, "What is a hidden Markov model?," *Nat. Biotechnol.* **22**, 1315–1316 (2004).
- ³⁰C. Bretó, D. He, E. L. Ionides, and A. A. King, "Time series analysis via mechanistic models," *Ann. Appl. Stat.* **3**(1), 319–348 (2009).
- ³¹T. Rydén, T. Teräsvirta, and S. Åsbrink, "Stylized facts of daily return series and the hidden Markov model," *J. Appl. Econ.* **13**, 217–244 (1998).
- ³²A. Jurgens and J. P. Crutchfield, "Shannon entropy rate of hidden Markov processes," *J. Stat. Phys.* **183**, 32 (2020).
- ³³A. Venegas-Li, A. Jurgens, and J. P. Crutchfield, "Measurement-induced randomness and structure in controlled qubit processes," *Phys. Rev. E* **102**(4), 040102 (2020).
- ³⁴T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley-Interscience, New York, 2006).
- ³⁵R. G. James, J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield, "Many roads to synchrony: Natural time scales and their algorithms," *Phys. Rev. E* **89**, 042135 (2014).
- ³⁶J. P. Crutchfield, P. Riechers, and C. J. Ellison, "Exact complexity: Spectral decomposition of intrinsic computation," *Phys. Lett. A* **380**(9–10), 998–1002 (2016).
- ³⁷A. Renyi, "On the dimension and entropy of probability distributions," *Acta Math. Hung.* **10**, 193 (1959).
- ³⁸B. B. Mandelbrot, *The Fractal Geometry of Nature* (W. H. Freeman and Company, San Francisco, CA, 1982).
- ³⁹G. A. Edgar, *Measure, Topology, and Fractal Geometry* (Springer-Verlag, New York, 1990).
- ⁴⁰K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications* (John Wiley, Chichester, 1990).
- ⁴¹Ya. B. Pesin, in *Dimension Theory in Dynamical Systems: Contemporary Views and Applications*, Chicago Lectures in Mathematics (University of Chicago Press, Chicago, IL, 1997).
- ⁴²I. Shimada and T. Nagashima, "A numerical approach to ergodic problem of dissipative dynamical systems," *Prog. Theor. Phys.* **61**, 1605 (1979).
- ⁴³G. Benettin, L. Galgani, A. Giorgilli, and J.-M. Strelcyn, "Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them," *Meccanica* **15**, 9 (1980).
- ⁴⁴Ja. B. Pesin, "Lyapunov characteristic exponents and ergodic properties of smooth dynamical systems with an invariant measure," *Dokl. Akad. Nauk.* **226**, 774–777 (1976).
- ⁴⁵D. Ruelle, "Sensitive dependence on initial conditions and turbulent behavior of dynamical systems," *Ann. N. Y. Acad. Sci.* **316**(1), 408–416 (1977).
- ⁴⁶G. Benettin, L. Galgani, and J.-M. Strelcyn, "Kolmogorov entropy and numerical experiments," *Phys. Rev. A* **14**(6), 2338–2345 (1976).
- ⁴⁷J. D. Farmer, E. Ott, and J. A. Yorke, "The dimension of chaotic attractors," *Physica D* **7**, 153 (1983).
- ⁴⁸J. Kaplan and J. Yorke, "Chaotic behavior of multidimensional difference equations," in *Functional Differential Equations and Approximation of Fixed Points*, Lecture Notes in Mathematics Vol. 730 (Springer, 1979), pp. 204–227.
- ⁴⁹F. Ledrappier and L. S. Young, "The metric entropy of diffeomorphisms," Part II: Relations between entropy, exponents and dimension," *Ann. Math.* **122**, 540–574 (1985).
- ⁵⁰P. Jacquet, G. Seroussi, and W. Szpankowski, "On the entropy of a hidden Markov process," *Theor. Comput. Sci.* **395**, 203–219 (2008).
- ⁵¹T. Holliday, A. Goldsmith, and P. Glynn, "Capacity of finite state channels based on Lyapunov exponents of random matrices," *IEEE Trans. Inf. Theory* **52**, 3509–3532 (2006).
- ⁵²B. Barany, "On the Ledrappier-Young formula for self-affine measures," *Math. Proc. Cambridge Philos. Soc.* **159**(3), 405–432 (2015).
- ⁵³P. A. P. Moran, "Additive functions of intervals and Hausdorff measure," *Math. Proc. Cambridge Philos. Soc.* **42**(1), 15–23 (1946).
- ⁵⁴C. Bandt, N. Viet Hung, and H. Rao, "On the open set condition for self-similar fractals," *Proc. Am. Math. Soc.* **134**(5), 1369–1374 (2006).
- ⁵⁵J. Jaroszewska and M. Rams, "On the Hausdorff dimension of invariant measures of weakly contracting on average measurable IFS," *J. Stat. Phys.* **132**, 907 (2008).
- ⁵⁶A. Jurgens and J. P. Crutchfield, "Backwards entropy rate of hidden Markov processes" (unpublished) (2021).
- ⁵⁷M. Rezaeian, "Hidden Markov process: A new representation, entropy rate and estimation entropy," [arXiv:0606114](https://arxiv.org/abs/0606114) (2006).
- ⁵⁸W. Słomczyński, J. Kwapien, and K. Życzkowski, "Entropy computing via integration over fractal measures," *Chaos* **10**(1), 180–188 (2000).
- ⁵⁹J. P. Crutchfield and K. Young, "Inferring statistical complexity," *Phys. Rev. Lett.* **63**, 105–108 (1989).
- ⁶⁰S. E. Marzen and J. P. Crutchfield, "Nearly maximally predictive features and their dimensions," *Phys. Rev. E* **95**(5), 051301 (2017).
- ⁶¹A. Jurgens and J. P. Crutchfield, "Minimal embedding dimension of minimally infinite hidden Markov processes" (unpublished) (2021).
- ⁶²V. I. Oseledets, "A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems," *Trans. Moscow Math. Soc.* **19**, 197–231 (1968).
- ⁶³J. H. Elton, "An ergodic theorem for iterated maps," *Ergod. Theory Dyn. Syst.* **7**, 481–488 (1987).
- ⁶⁴G. Froyland and K. Aihara, "Rigorous numerical estimation of Lyapunov exponents and invariant measures of iterated function systems and random matrix products," *Intl. J. Bifurcation Chaos* **10**(1), 103–122 (2000).
- ⁶⁵A. Boyarsky and Y. S. Lou, "A matrix method for approximating fractal measures," *Intl. J. Bifurcation Chaos* **02**, 167–175 (1992).
- ⁶⁶Shapely, see <https://pypi.org/project/Shapely/> for "Geometric Objects, Predicates, and Operations," version 1.7.1. <https://doi.org/10.21105/joss.00738>.