

Maximally predictive ensemble dynamics from data

Antonio C. Costa^{a*} and Tosif Ahamed^{b,c}, David Jordan^d, Greg J. Stephens^{a,c}

^a*Department of Physics and Astronomy, Vrije Universiteit Amsterdam, 1081HV Amsterdam, The Netherlands*

^b*Lunenfeld-Tanenbaum Research Institute, Mount Sinai Hospital, Toronto, Canada*

^c*Biological Physics Theory Unit, OIST Graduate University, Okinawa 904-0495, Japan*

^d*Gurdon Institute, University of Cambridge, United Kingdom*

We leverage the interplay between microscopic variability and macroscopic order to connect physical descriptions across scales directly from data, without underlying equations. We reconstruct a state space by concatenating measurements in time, building a maximum entropy partition of the resulting sequences, and choosing the sequence length to maximize predictive information. Trading non-linear trajectories for linear, ensemble evolution, we analyze reconstructed dynamics through transfer operators. The evolution is parameterized by a transition time τ : capturing the source entropy rate at small τ and revealing timescale separation with collective, coherent states through the operator spectrum at larger τ . Applicable to both deterministic and stochastic systems, we illustrate our approach through the Langevin dynamics of a particle in a double-well potential and the Lorenz system. Applied to the behavior of the nematode worm *C. elegans*, we derive a “run-and-pirouette” navigation strategy directly from posture dynamics. We demonstrate how sequences simulated from the ensemble evolution recover effective diffusion in the worm’s centroid trajectories and introduce a top-down, operator-based clustering which reveals subtle subdivisions of the “run” behavior.

INTRODUCTION

The principled integration of fine-scale degrees of freedom yielding coarse-grained theories that successfully capture large-scale structure is a remarkable achievement in physics and an important example for the quantitative understanding of complex systems, from the ensembles of statistical mechanics to the kinetic theory of fluids (see e.g. [1]). Powerful as they are, such examples require detailed knowledge about the underlying dynamics, symmetries and conservation laws, or a parameterization of the models for tools such as the renormalization group [2]. How can we emulate this success in situations for which we lack comparable understanding?

One intriguing approach is to apply statistical inference to characterize how model selection, assessed through the Fisher information matrix, can systematically vary across scales [3, 4]. An alternative is to work directly from data, for example by seeking a renormalization-type scaling [5, 6]. In complex systems, those with many components heterogeneously interacting across multiple scales, working directly with data is appealing but also requires caution; even as new machine learning approaches allow for unprecedented quantitative access (see e.g. [7, 8]), we often lack basic understanding such as the correct class of models for the phenomena of interest.

Timescale separation offers an important and fundamental guide; coarse-grained degrees of freedom often vary slowly compared to their microscopic constituents, such as the seconds and longer lifetime of currents in a

river against the $\sim 10^{-15}$ sec of molecular collisions. A timescale separation is integral to the Langevin and Einstein approach to Brownian motion (see e.g. [9]), where collisions experienced by a particle in a heat bath are treated as Gaussian white noise, allowing for the prediction of larger scale diffusive dynamics. Indeed, long-lived dynamics have been the focus of a wide range of previous work from hydrodynamic Lyapunov modes in high dimensional dynamical systems [10] to neuroscience [11] and molecular dynamics [12].

Long timescales are of direct interest in organism-scale behaviors [13–16] where navigation provides a clear example; how do persistent states underlying search strategies emerge from much faster posture dynamics? An organism must successfully navigate its environment to improve its chances of feeding, mating, and survival, and substantial theoretical work has investigated the principles of environmental search, including chemotaxis, infotaxis and levy flights [17–19]. While environmental search is usually considered at the level of movement of the body centroid, differing modes of movement are generated by modulating neuromuscular control systems and their resulting postural dynamics. Human running and walking, for example, have discernibly different patterns of limb movement resulting in different locomotory speeds. From sequences of body posture it should be possible, in principle, to identify persistent states, as well as the behavioral strategies which modulate their frequency, thus linking distinct dynamics across scales.

In the context of systems whose fundamental equations are unknown, we introduce a framework for the principled extraction of coarse-grained, slow dynamics directly from time series data. We combine delay embedding, fine-scale partitioning, and a Markov approximation of the resulting sequence entropy to reconstruct a maximally predictive state space. We analyze dynam-

* Current address: Laboratoire de Physique de l’Ecole normale supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, F-75005 Paris, France

ics through ensemble evolution directed by transfer operators approximated from transition probabilities between partitions, and identify long-lived states through the eigenvalues and eigenfunctions of these operators. We illustrate our approach in stochastic and deterministic model systems, and in the undulatory posture movement of the nematode *C. elegans*, for which the fundamental dynamics are unknown. We derive “run-and-pirouette” navigation dynamics directly from the worm’s posture, connecting behavior across spatiotemporal scales differing by multiple orders of magnitude, and revealing subtle behavioral motifs.

THE TRANSFER OPERATOR PERSPECTIVE

We consider the general setting of the time evolution of a system with both deterministic and noisy influences through a set of Ito stochastic differential equations for the evolution of the system’s state $x \in \mathbb{R}^D$,

$$dx_t = F(x_t)dt + \gamma(x_t)dW_t, \quad (1)$$

where $F : \mathbb{R}^D \rightarrow \mathbb{R}^D$, $\gamma : \mathbb{R}^D \rightarrow \mathbb{R}^D$ and dW_t is the increment of a Wiener process. Natural systems are typically composed of a large number of interacting degrees of freedom ($D \gg 1$) and both F and γ can be complicated non-linear functions. While such equations provide important information about the fine-scale properties of the dynamics, extracting longer time scale, coarse-grained properties from them is often challenging, indeed this is the primary challenge of kinetic theory (see e.g. [1]). Even when D is small, there are large scale dynamical patterns, such as unstable periodic orbits (UPOs) in chaotic dynamical systems [20], whose properties we cannot immediately infer from the knowledge of the equations of motion.

The motion of a particle trapped in a double-well potential at finite temperature provides an explanatory model dynamics, Fig. 1(A-left). If the barrier height is large compared to the temperature, the particle will hop between wells with almost vanishing probability, see e.g. [21]. Studying a single trajectory for only a short amount of time is misleading, as the particle mostly fluctuates around only one of the wells. If, however, we consider an *ensemble* of trajectories, the two-well nature of the system becomes apparent, Fig. 1(A-middle). For both this example and the general dynamics of Eq. (1), state space densities evolve according to a Fokker-Planck equation,

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (F\rho) + \frac{1}{2} \nabla \cdot [\nabla \cdot (\Gamma\rho)], \quad (2)$$

where $\Gamma = \gamma\gamma^\top$. By defining the linear operator \mathcal{L} such that $\mathcal{L}\rho = -\nabla \cdot (F\rho) + \frac{1}{2} \nabla \cdot [\nabla \cdot (\Gamma\rho)]$, the ensemble

evolution equation can be written as,

$$\begin{aligned} \frac{\partial \rho}{\partial t} &= \mathcal{L}\rho \\ \rho(\tau) &= e^{\mathcal{L}\tau} \rho_0 \\ \rho(\tau) &= \mathcal{P}_\tau \rho_0, \end{aligned} \quad (3)$$

where ρ_0 is an initial density, and \mathcal{P}_τ is the Perron-Frobenius (PF) operator [22–25] that “transfers” the ensemble density from time $t = 0$ to a finite time τ in the future. We note that \mathcal{L} and therefore \mathcal{P}_τ are not restricted to evolving dynamics of the Fokker-Planck class, extending equally to other stationary white noise processes and even deterministic dynamics [24–26], such as the Liouville operator of classical Hamiltonian dynamics. We leverage this universality to study stochastic and deterministic systems through a single framework.

Finding tractable analytic expressions of the PF operator is difficult in all but toy systems [21]. However, such operators can be effectively approximated by a Markov transition matrix between state space partitions [25, 27], and we exploit such an approximation to bridge the gap between fine scale and coarse-grained phenomenological descriptions of complex dynamical systems.

In particular, we focus on the operator eigenvalue spectrum, which, in the context of our Markov approximation, naturally orders the timescales of the system. In the two-well example, Fig. 1(A-middle), the equilibrium distribution emerges as the eigenfunction with unit eigenvalue of \mathcal{P}_τ , $\mathcal{P}_\tau \pi = \pi$, while the exponentially-long timescale for hopping between wells is captured by the next-largest eigenvalue [21]. In general, the operator approach provides a hierarchy of timescales, which can naturally reveal a timescale separation and thus an effective phenomenological description of the system through the associated long-lived eigenfunctions. These eigenfunctions can capture important emergent coarse-grained properties of the dynamics such as transitions between multiple metastable states, Fig. 1(A-right). In the context of chemical kinetics, for example, such eigenfunctions provide ideal reaction coordinates or order parameters [12, 28, 29], which are analogous to the committor function of transition path theory [29–31].

PREDICTIVE INFORMATION FOR STATE SPACE RECONSTRUCTION

Measurements of real-world dynamical systems rarely capture the full set of variables that constitute the state space, which evolves according to a first order differential equation or discrete map [32, 33]. Such incomplete information results in non-Markovian effects that are reflected in the increased unpredictability of the time series: the dynamics within a state space is Markovian and maximally predictive, since we need only know the current state to predict the next. Markovian dynamics is rigorously applicable in fundamental physical systems

and we capture this ideal with an approach to state space reconstruction of unknown dynamics that maximizes the predictive information of the state. Building upon the results of Takens and others [34–38], who have shown that past information can be leveraged to reconstruct a Markovian state [39], we include previous time points and examine how the predictability improves.

Our starting point is a set of incomplete measurements of an unknown dynamical system, $\vec{y} = M(\vec{x})$, where M is a measurement function mapping the underlying state space $\vec{x} \in \mathbb{R}^D$ into our measurements $\vec{y} \in \mathbb{R}^d$, for which typically $d < D$, Fig. 1(B-left). We expand the putative state by adding $K - 1$ time delays to the measurement time series, sampled on a time scale δt , yielding a candidate state space $X_K \in \mathbb{R}^{d \times K}$. We quantitatively characterize the unpredictability of the candidate state space through the entropy rate of the symbolic dynamics resulting from partitioning each $K \times d$ -dimensional putative state space into N Voronoi cells through clustering (Methods). With a partitioning, the reconstructed dynamics are encoded as a row-stochastic transition probability matrix P which evolves a state-space density p by a time δt ,

$$p_i(t + \delta t) = P_{ji} p_j(t), \quad (4)$$

where we sum over repeated indices (Einstein’s summation convention). The entropy rate of the source is approached by estimating the entropy rate of the associated Markov chain for increasing values of K ,

$$h_{P_N}(K) = - \sum_{ij} \pi_i P_{ij} \log P_{ij} \quad (5)$$

where π is the estimated stationary distribution of the Markov chain P . The Markov approximation of the entropies provides an estimate of the conditional entropies, $\langle -\log [p(s_j | s_i)] \rangle$, where $i, j \in \{1, \dots, N\}$ and s_i are the obtained discrete states. Each discrete state contains a population of delay vectors $\vec{y}^K = \{\vec{y}_i, \dots, \vec{y}_{i+K-1}\}$, and therefore the entropy of the Markov chain provides an estimate of the conditional entropy,

$$\begin{aligned} h_{P_N}(K) &\approx \langle -\log [p_N(\vec{y}_{i+K} | \vec{y}_i \dots \vec{y}_{i+K-1})] \rangle \\ &= H_{K+1}(N) - H_K(N) \\ &= h_K(N) \end{aligned}$$

where $H_K(N)$ is the entropy of the symbolic sequence built by discretizing the \vec{y} space into N partitions. Note that the entropy rate is a non-decreasing function on the number of partitions, N , which is equivalent to a typical state-space scale $\epsilon \sim 1/N$. For deterministic dynamical systems, the Kolmogorov-Sinai (KS) entropy rate is defined in the limit of $K \rightarrow \infty$, $\delta t \rightarrow 0$ and as a supremum over all possible partitions of the state space, which is usually approached by taking $N \rightarrow \infty$ ($\epsilon \rightarrow 0$) [40, 41]. Deterministic periodic systems are completely

predictable $h_{KS} = 0$ and chaotic systems possess a characteristic finite entropy rate that is equal to the sum of positive Lyapunov exponents when the conditions of Pesin’s theorem are satisfied [42]. For stochastic processes $h_{KS} = \infty$ as the system possesses information on all length scales [41]. Importantly, we seek to preserve as much information as possible in the discretized state space, setting the number of partitions as the largest N after which the entropy rate stops increasing due to finite-size effects Fig. 1(B-middle), yielding the maximum entropy rate with respect to the number of partitions. Besides the KS entropy, other quantities such as information or correlation dimensions can in principle be obtained by studying the scaling of the measure with the partition size [40, 43, 44].

Our estimate of the entropy rate is analogous to the Cohen-Procaccia (CP) correlation entropy rate estimate [40] for finite K . In the CP method, sequence entropies are approximated by considering time series segments of duration K , and estimating the probability of finding pairs of sequences within a distance ϵ of each other. The KS entropy limit is approached with $\epsilon \sim 1/N \rightarrow 0$, $K \rightarrow \infty$ and $\delta t \rightarrow 0$. Correlation entropies were introduced to avoid counting long sequences of discrete states needed to reach the KS limit ($\epsilon \rightarrow 0$) which renders impractical the direct estimation of the sequence entropies for large K (the number of possible sequences scales with the length as N^K). As in the CP method, instead of counting sequences of discrete states, we first extend the state by adding K time delays, but then return to the partitioned space to estimate the entropy rate through the associated Markov chain, also avoiding the explicit counting of long sequences of discrete states. In the Methods we provide further discussion on the relationship between approaches to the estimation of entropy rates.

With increasing K , we approach the partition entropy rate of the source from above so that the difference,

$$\delta h_{P,N}(K) = h_{P_N}(K - 1) - h_{P_N}(K),$$

is a non-increasing function of K : $\delta h_{P_N}(K) \geq 0$. $\delta h_{P,N}(K)$ has been used to define measures of forecasting complexity in dynamical systems [45] and is the amount of information that has to be kept in the $K - 1$ time delays for an accurate forecast of the next time step. When K is too short, $\delta h_{P_N}(K)$ will be large, meaning that a large amount of information is required in order to make an accurate prediction. In systems with finite range correlations, there is a K^* for which $\delta h_{P,N}(K^*) = 0$, in which K^* corresponds to the amount of memory sitting in the measurement time series. Our state space reconstruction seeks K^* such that $h_{P_N}(K)$ is minimized, i.e.,

$$\partial_K h_{P_N}(K^*) = 0, \quad (6)$$

which corresponds to maximizing the predictive information. We recall the general definition of predictive

information [46], $I_{\text{pred}} = \left\langle \log \frac{p(x_{\text{future}}|x_{\text{past}})}{p(x_{\text{future}})} \right\rangle$. Defining x_{past} as the first $K - 1$ time steps in the time series and x_{future} as the K th time step,

$$I_{\text{pred}} = \langle -\log p(\vec{y}_{i+K}) \rangle - \langle -\log p(\vec{y}_{i+K}|\vec{y}_i \dots \vec{y}_{i+K-1}) \rangle \\ = H(\vec{y}) - h_K,$$

and thus with respect to a partition into N states, we have

$$I_{\text{pred}}(N) = H_1(N) - h_{P_N}(K). \quad (7)$$

$I_{\text{pred}}(N)$ thus quantifies the reduction in the topological entropy of the measurement time series that results from the knowledge of the transition probabilities. The predictive information is maximized when $h_{P_N}(K)$ is minimized, which for a system with finite correlations is attained when $\partial_K h_{P_N}(K^*) = 0$.

Given a reconstructed state X_{K^*} , the transition matrix P_{ij} (defined with timestep $\tau = \delta t$), is an example of an approximated *transfer operator* \mathcal{P}_τ , Eq. (3) [27], Fig. 1(B-right). The transfer operator dynamics solely depends on the topology of the state space trajectories, which are guaranteed to be preserved by a state space embedding [35, 38]. In that sense, \mathcal{P}_τ is in principle exactly preserved, and maintains all the properties of the underlying dynamics. In contrast, the set of non-linear equations of motion driving the dynamics of the reconstructed state $\dot{X} = F(X)$, which are for instance required to obtain estimates of the local Jacobian, strongly depend on the geometric properties of the space (such as dimensions, metric, etc.), making it non-trivial to accurately approximate the underlying dynamics $\dot{x} = f(x)$, Fig. 1(A,B-right). The transfer operator formalism is therefore complementary to trajectory based approaches, providing a means to study large scale properties of the dynamics while being robust to the precise geometric properties of the reconstructed state.

EXTRACTING COARSE-GRAINED DYNAMICS

Instantaneous ensemble dynamics is characterized by the generator \mathcal{L} of Eq. 3, Fig. 1(C-left). However, when working directly from measurements sampled at discrete times, the finite-time transfer operator \mathcal{P}_τ is more immediately available. The operators \mathcal{L} and \mathcal{P}_τ share the same set of eigenfunctions, while the eigenvalues $\lambda_i(\tau)$ of \mathcal{P}_τ are exponential functions of the generator eigenvalues Λ_i , $\lambda_i(\tau) = e^{\Lambda_i \tau}$. When the eigenvalues λ_i are real, the relaxation of the density dynamics is characterized by the implied relaxation times corresponding to each eigenvalue,

$$t_i^{\text{imp}}(\tau) = \left| \hat{\Lambda}_i^{-1} \right|(\tau) = \frac{-\tau}{\log \lambda_i(\tau)}. \quad (8)$$

\mathcal{P}_τ lives in an infinite dimensional functional space, which we discretize using N basis functions: in our partitioned space, the basis functions are characteristic functions and the measure is piecewise constant (Methods). The truncation at finite N erases fine-scale information within the partition, and so the ability of the transfer operator approximation P_{ij} to capture long-time dynamics depends on the transition time τ , a parameter that we vary. For the state space reconstruction described in the previous section, we chose τ as the sampling time $\tau = \delta t$ in order to maximize short-time predictability and to capture the Kolmogorov-Sinai limit of the entropy rate of the dynamics, $\epsilon \sim 1/N \rightarrow 0$, $\delta t \rightarrow 0$ [41]. However, to accurately capture longer-time dynamics and metastable states, we vary τ and study changes in the inferred spectrum, Fig. 1(C-middle).

For $\tau \rightarrow 0$, P_{ij} is close to an identity matrix (the system has little time to leave the current partition) and all eigenvalues are nearly degenerate and close to unity, Fig. 1(C-middle). For τ much longer than the mixing time of the system $\tau \gg t_{\text{mix}}$, the eigenvalues start collapsing and in the limit of $\tau \rightarrow \infty \Rightarrow P_{ij}(\tau) \sim \pi_j$: the probability of two subsequent states separated by such τ becomes independent, the eigenvalues of $P_{ij}(\tau)$ stop changing and t_{imp} grows linearly with τ for all eigenfunctions. In this limit, the transition probability matrix contains noisy copies of the invariant density akin to a shuffle of the symbolic sequence. This yields an effective noise floor, which is observed earlier (shorter τ) for faster decaying eigenfunctions. For intermediate τ there is a spectral gap, indicating that the fast dynamics have relaxed and we can isolate the slow dynamics. For such τ , the inferred relaxation times reach a constant value that matches the underlying long time scale Λ_2^{-1} , indicating that such dynamics are Markovian: the Chapman-Kolmogorov equation (see e.g. [47]) is verified $\mathcal{P}_{n\tau}\rho = \mathcal{P}_\tau^n \rho$, and thus $t_i^{\text{imp}}(n\tau) = \frac{-n\tau}{\log(\lambda_i(\tau)^n)} = t_i^{\text{imp}}(\tau)$ [48]. Notably however, the dynamics on the discretized state space is not generally Markovian on all time scales since the discretization erases fine scale information. Nonetheless, a Markov model can approximate the slow dynamics arbitrarily well, provided that the discretization can be made arbitrarily fine and the transition time scale τ sufficiently long [49, 50]. We choose τ^* after the initial transient, which corresponds to the regime in which the spectral gap is the largest and timescale separation is evident (green shaded region).

Within our reconstructed ensemble evolution, coarse-grained dynamics can be identified between state space regions where the system becomes temporarily trapped, these are metastable states or almost invariant sets [51, 52]. We search for collections of states that move coherently with the flow: their relative future states belong to the same macroscopic set. Since the slowest decay to the stationary distribution is captured by the first non-trivial eigenfunction of the transfer operator, we search for the state space subdivision along this eigenfunction

that maximizes the coherence of the resulting macroscopic sets, a heuristic previously introduced in [51].

A set S is coherent when a state is more likely to remain within the set than it is to leave it within a time τ . We quantify this intuition through $\chi_{\mu,\tau}(S) = \frac{\mu(S \cap \Phi_{-\tau} S)}{\mu(S)} = \frac{\mu(\Phi_{\tau} S \cap S)}{\mu(S)}$, where μ is the invariant measure preserved by the invertible flow Φ_{τ} . Given an inferred transfer operator $P_{ij}(\tau)$ and its associated stationary eigenvector π , we can immediately compute χ (Methods),

$$\chi_{\pi,\tau}(S) = \frac{\sum_{i,j \in S} \pi_i P_{ij}(\tau)}{\sum_{i \in S} \pi_i}, \quad (9)$$

and we identify a coherent set when $\chi_{\mu,\tau}(S)$ is large. The approximation of transfer operators for the extraction of long-lived states has been successfully applied in a number of systems: from identifying metastable protein conformations in molecular dynamics simulations [12], to determining coherent structures in oceanic flows [26, 53, 54] (for reviews, see [27, 31, 55–58]).

To identify optimally-coherent metastable states through spectral analysis, we define a time-symmetric (reversibilized) transfer operator,

$$\mathcal{P}_r \equiv \frac{\mathcal{P} + \mathcal{P}^\dagger}{2}, \quad (10)$$

where \mathcal{P}^\dagger is the dual operator to \mathcal{P} , pulling the dynamics backward in time; see Methods for the discrete numerical approximation, P_r . While transfer operators describing ensemble dynamics are not generally symmetric [51, 59], the definition of coherence is invariant under time reversal: since the measure is time invariant, it does not matter in which direction we look for mass loss from a set [31, 60]. Besides the invariance property, an important benefit of working with P_r is that its second eigenvector ϕ_2 provides an *optimal* subdivision of the state space into almost invariant sets, as shown in [60]. Given τ^* and the corresponding $P_r(\tau^*)$, we identify metastable sets by choosing a subdivision along ϕ_2 that maximizes an overall measure of coherence,

$$\chi = \min\{\chi_{\mu,\tau^*}(S^+), \chi_{\mu,\tau^*}(S^-)\}, \quad (11)$$

where $\{S^+, S^-\}$ result from a partition at ϕ_2^c : the optimal almost invariant sets are identified with respect to the sign of $\phi_2 - \phi_2^c$, Fig. 1(C-right).

We use the real eigenvalues of P_r to probe the Markovianity of the inferred transfer operator across transition times τ and to choose τ^* . However, for out-of-equilibrium systems $t_{\text{imp}}(\tau^*)$ is not generally interpretable as a kinetic time scale of the underlying dynamics, even as the coherent sets are accurate, a consideration that will primarily apply to our Lorenz and posture dynamics applications in later sections. We will further address the use of P_r in the Discussion. Note that the validity of the Markov approximation can also be assessed *a posteriori* by analyzing the sensitivity of the identified metastable states and their kinetics to changes in τ .

RESULTS

Brownian particle in a double well potential

We consider the Langevin dynamics of a particle in a double well potential $V(x) = (x^2 - 1)^2$ in thermal equilibrium, Fig. 2(A-left). The equation of motion for the position x of the particle is

$$m\ddot{x} = -\gamma\dot{x} - 4x(x^2 - 1) + \sqrt{2\gamma\beta^{-1}}\eta(t), \quad (12)$$

where $\beta^{-1} = k_B T$, m , γ , k_B and T are the mass, damping coefficient, Boltzmann constant and temperature, respectively, and $\eta(t)$ is a Gaussian white random variable capturing the random collisions with the heat bath. The dynamics is second order in the position x , and so we introduce the velocity ($v = \dot{x}$) degree of freedom to rewrite Eq. (12) as a system of first order Ito stochastic differential equations evolving the state (x, v) in time (we use $m = \gamma = 1$ throughout),

$$\begin{cases} dx_t = v_t dt \\ dv_t = -v_t dt - 4x_t(x_t^2 - 1)dt + \sqrt{2\beta^{-1}}dW_t, \end{cases} \quad (13)$$

where dW_t is a Wiener process. The dynamics of (x, v) is first order and therefore fundamentally Markovian and memoryless. To emulate a real data example for which only a subset of the degrees of freedom is available, we measure only the position x of the particle, and seek to reconstruct the state space by adding time delays.

We generate long simulations of 10⁶ s sampled at $\delta t = 0.05$ s, with temperatures ranging between $\beta^{-1} = 0.5$ J and $\beta^{-1} = 2.5$ J (Methods): an example trajectory for $\beta^{-1} = 0.5$ J is shown in Fig. 2(A-right). We concatenate the x measurement with K delays, partition the resulting state with N partitions and show the entropy rate as a function of increasing K and N , Fig. 2(B-left). Notably, the behavior of $h(K, N)$ reveals fundamental information about the properties of the measurement time series. Since the generating equations are a continuous-time stochastic process, $\lim_{N \rightarrow \infty} h(K, N) = \infty$. However, the rate of approach to infinity depends on the complexity of the random process [41]. For a Markov-chain approximation of the Langevin dynamics we expect $h \propto \log(N)$, as observed for $K \gtrsim 2$ frames, a result which we derive from the Cohen-Procaccia $h(\epsilon, \tau)$ estimate of the entropy translated to our partition-based approach assuming $\epsilon \sim 1/N$ (Eq. 3.71 in [41] for fixed τ). For small N , there are not enough partitions to see the difference between candidate state spaces. For large N , finite-sampling effects produce an underestimation of the entropy [40, 41], a drop we see most visibly for $K = 1$ frame. At intermediate N , where estimation effects are negligible, there is an abrupt change in the entropy rate from $K = 1$ to $K > 2$ frames = 0.1 s, Fig. S5(A), indicative of the finite memory sitting in the

measurement time series from the missing degree of freedom. Notably, the dynamics is not fully memoryless with $K = 2$ delays, as one would naively expect, a result due to the Euler-scheme update which induces memory effects on the infinitesimal propagator $P(x_{k+1}|x_k, x_{k-1})$ [61–64]. We choose $K^* = 7$ frames = 0.35 s for subsequent analysis, which recovers the equipartition theorem, showing that momentum information is accurately captured by the reconstructed state, Fig. 2(C). The behavior of the entropy with K and N is qualitatively conserved across the range of temperatures studied here, Fig. S1(A).

Using the reconstructed state space with $K^* = 7$ frames, we approximate the transfer operator by making a Voronoi tessellation through k-means clustering (Methods), and counting the transition probabilities between discrete states as a function of transition time scale τ . We construct the reversibilized transition matrix P_τ and show the 15 slowest inferred relaxation times for $N = 1000$ partitions and transition times τ , Fig. 3(A). The longest relaxation time decays as the system thermalizes, and reaches its asymptotic limit after $\tau \sim 5$ s. As we increase τ to a few multiples of the hopping time scale, the eigenvalues stop changing and t_{imp} simply grows linearly with τ ; faster decaying eigenfunctions exhibit this behavior earlier, Fig. 3(A).

To gain further intuition into the meaning of the longest implied timescale, we simulate the ensemble dynamics using the inferred transfer operator P_{ij} for $\beta^{-1} = 0.5$ J. In Fig. 3(B) we show the decay of second eigenfunction as a function of time and the coordinates (u_1, u_2) , which are projections along the first two singular vectors of the state space embedding that match the position and velocity degrees of freedom, Fig. S1(B). We initiate the system with the density ρ concentrated in one of the $N = 1000$ partitions and use $\tau^* = 10$ s as the transition time. After $t = 10$ s the ensemble has distributed itself in one of the wells and at $t = 40$ s, it has almost completely relaxed to the invariant density, which here corresponds to the Boltzmann distribution.

The state space structure of the second eigenvector of the reversibilized transition matrix ϕ_2 reveals a collective coordinate capturing transitions between the wells. In Fig. 3(C) we show a contour plot of ϕ_2 in the (u_1, u_2) space for $\beta^{-1} = 0.5$ J where it is apparent that the sign of the eigenvector effectively splits the reconstructed state space into the two wells of the system. In fact, the maximum of χ , Eq. (11), is obtained at $\phi_2^C \sim 0$ across temperatures, Fig. S2(B). Notably, due to the underdamped nature of the dynamics the separation between potential wells includes the high velocity transition regions $|u_2|$, as expected from theory [65]. Indeed, $|\hat{\Lambda}_2|$ provides an excellent fit to the hopping rate which we estimate empirically by fitting an exponential function to the dwell time distributions in both wells, Fig. 3(D). We note that in order to examine different temperatures the choice of τ^* should change accordingly to reflect the different resulting dynamics. Increasing the temperature reduces the

hopping rates and therefore the amount of time it takes before the system relaxes to the stationary distribution. Nonetheless, the qualitative behavior of the longest inferred implied relaxation times is conserved across temperatures, Fig. S2(A): a short τ transient quickly converges to the time scale of hopping between wells, before the system completely mixes and the implied time scales grow linearly. To choose τ^* consistently across temperatures we find the transition time that minimizes the longest reversibilized $t_{\text{imp}}(\tau)$. As expected the resulting τ^* is reduced as we increase the temperature, reflecting the faster nature of the dynamics.

Lorenz system

Our approach to state space reconstruction and extraction of coarse-grained dynamics is equally applicable to deterministic systems, which we demonstrate using the Lorenz equations,

$$\begin{cases} \dot{x} = \sigma(y - x) \\ \dot{y} = x(\rho - z) - y, \\ \dot{z} = xy - \beta z \end{cases} \quad (14)$$

in the standard chaotic regime $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$, Fig. 4(A). We take measurements of only the x variable (sampled at $\delta t = 0.01$ s for 10^5 s), reconstruct the state space, and extract coarse-grained dynamics through the transfer operator. To reconstruct a maximally-predictive state space, we add delays to the measurement series to minimize the entropy rate, shown in Fig. 4(B) for different numbers of delays K and partitions N . We use $K^* = 12$ frames, an embedding window of 0.12 s for which the entropy rate is constant, Fig. S5(B), to define the state space. Unlike the stochastic dynamics within the double-well, the entropy rate of this deterministic dynamical system reaches an asymptotic value: this is the Kolmogorov-Sinai entropy h_{KS} [66] (see also [41] for subtleties), and our result is in excellent agreement with the maximum Lyapunov exponent, as expected for this system where Pesin’s theorem [42] applies, $h_{KS} = \sum_i \lambda_i$, where λ_i are the positive Lyapunov exponents. In addition to the entropy rate, our approach can easily convey other important invariants, such as the information dimension, Fig. S4.

We further explore the connection to ergodic properties by showing a contour plot of the inferred first non-trivial eigenfunction ϕ_2 of the reversibilized transfer operator P_τ , Eq. 10, constructed with $N = 10^{3.5}$ partitions and transition time $\tau^* = 0.1$ s, and shown in the space of first two largest singular modes (u_1, u_2) , Fig. 4(C). The sign of ϕ_2 divides the state space into its almost invariant sets, Fig. S3(B), which are partially split along the shortest unstable periodic orbit (UPO) [67, 68], identified here through recurrences [69], Fig. S3(C) (Methods). Chaotic dynamics generally exhibit an intricate

interplay of multiple timescales, which is apparent in the eigenvalue structure of P_τ for different transition times, Fig. S3(A), which is more complicated than that of the double-well. Indeed, in the chaotic regime, the Lorenz dynamics can be described through a skeleton of UPOs [70], which are reflected in the periodic behavior of the eigenvalues as a function of τ .

Movement behavior across scales: “run-and-pirouette” navigation from posture dynamics in *C. elegans*

We use the operator approach to demonstrate a direct bridge between multiple behavioral scales in the foraging dynamics of the nematode worm *C. elegans*, an important model organism in genetics and neuroscience [71, 72]. We show that time scale separation and coarse graining at the level of postural dynamics can reliably identify long-lived states, previously described as “runs” and “pirouettes”. Further, we show how the dynamics of switching between these states leads to a ballistic-to-diffusive transition in centroid trajectories. Guided by heavy tails in the dwell time distribution of the emergent “run” and “pirouette” states, we detail an operator-based clustering to find that the “run” state is better described as compound, composed of four previously unidentified subdivisions.

We employ a previously analysed dataset [73] composed of 35 minute recordings of $N = 12$ N2 lab-strain worms freely moving on an agar plate, sampled at $\delta t = 1/16$ s. We measure the worm’s body posture using a low dimensional representation of the worm’s centerline, expressed as five “eigenworm” coefficients $\vec{a} \in \mathbb{R}^5$ [73, 74], Fig. 5(A). To reconstruct the state space and transfer operators, we stack K delays of the measurement time series, and estimate the entropy rate as a function of K and the number of partitions N , Fig. 5(B). After $N \sim 1000$ partitions we start observing finite-size effects as the entropy rate stops increasing, specially for lower K . Nonetheless, after $K \sim 8$ frames = 0.5 s the entropy rate curves start collapsing, and we choose $K^* = 11$ frames = 0.6875 s to define the reconstructed space X_{K^*} as this corresponds to a regime in which the entropy rate is roughly constant, Fig. S5(C). We choose $N = 1000$ partitions so as to maximize the entropy rate and extract the $K \rightarrow \infty$ limit of the entropy rate by extrapolation [75]: $h_\infty = 1.19 \pm (1.11, 1.25)$ nat/s, obtained by extracting the offset of a linear fit in the interval $1/K = [1, 8]$ s⁻¹. The early K behavior h is proportional to $1/K$, indicating that the system has finite range correlations [46, 75], Fig. 5(B, inset). The large K decrease of h results from sampling error in the high-dimensional space [75]. Our estimated entropy rate is in good agreement with a previous calculation using estimated Lyapunov exponents from the same data [76]. Given K^* , we approximate the ensemble dynamics with $N = 1000$ partitions, and estimate the relaxation times

of the first ten eigenfunctions of the reversibilized transition matrix, Fig. 5(C). We choose $\tau^* = 0.75$ s (dashed line) for subsequent analysis, the earliest τ where an approximately-Markov regime appears in the long-lived dynamics.

To identify the longer-time behavior apparent in the transfer operator, we project the dynamics along ϕ_2 (the 2nd eigenvector of the reversibilized operator) and use this mode to color the body wave phase velocities ω and mean body curvature $\langle |\theta| \rangle$, averaged in $2\tau^* = 1.5$ s windows, Fig. 5(D). Positive values (blue) generally align with negative phase velocities and large turning amplitudes, indicative of “pirouettes”, while negative values (red) correspond to positive phase velocities and low turning amplitudes, indicative of “forward runs”. In the inset we show an example 10 minute long centroid trajectory color coded by the projection along ϕ_2 . Negative projections occur during “runs”, while positive values are found when worms abruptly change the direction of locomotion. We cluster ϕ_2 into two coarse-grained states by maximizing the coherence measure χ , Eq. (11), and find that the cumulative distribution of resulting run lengths $1 - P(t_{\text{state}} < t)$ is roughly characterized by two time scales, Fig. 5(E), fit by a sum of exponential functions and in excellent agreement with previous phenomenological observations [77].

The “run” and “pirouettes” state transition dynamics are out of equilibrium. The transition rates imply a relaxation time of $|\Lambda_2^{-1}| = (1/\tau_P + 1/\tau_R)^{-1} \sim 3.31$ (2.99, 3.72) s, which is much shorter than the implied time scale of $|\hat{\Lambda}_2^{-1}|(\tau^*) = 8.42$ (7.40, 9.52) s, computed from the reversibilized dynamics. Indeed, positive entropy production has been reported in the short-time posture dynamics [76]. For out-of-equilibrium systems, we generally expect a discrepancy in eigenvalues (timescales) between P and P_τ , an aspect previously raised in the section “Extracting Coarse-Grained Dynamics” and to which we return in Discussion.

Transfer operator simulations predict both fine scale posture dynamics and large scale diffusive properties of foraging trajectories

We use the inferred transfer operator, P_{ij} , to simulate symbolic sequences of the worm’s posture movement: probing both fine scale posture dynamics and large scale properties of the foraging trajectories. We start from the initial discrete state of an individual worm and simulate symbolic sequences by sampling from the conditional probability distribution $P(s_j|\hat{s}_i)$, Fig. 6(A), where \hat{s}_i is the current state, s_j are all possible future states after a time scale τ^* and $P(s_j|\hat{s}_i)$ is the i -th row of the inferred transfer operator matrix $P_{ij}(\tau^*)$. The result is a sequence of microstates with the same duration as the worm trajectories, but with a sampling time $\delta t = \tau^*$. These dynamics are effectively diffusive in the state space: hopping between microstates according

to the Markov dynamics and random selection from the population of state space points X_i within each visited partition.

From the diffusion in state space, where each state corresponds to a sequence of postures, we obtain a time series of “eigenworm” coefficients $\{\vec{a}(t)\}$ by sampling the first posture in each state: we thus obtain a simulated time series of postures of the same duration as the experimental data, sampled at $\delta t = \tau^*$. We estimate the autocorrelation function of the simulated time series, showing that the operator derived dynamics accurately predicts the fine scale posture movements across worms, Fig. 6(B). On larger scales, we assign each microstate to a “run” or “pirouette” state as before, and estimate the average kinetic transition rates from runs-to-pirouettes $\kappa_{R \rightarrow P}$ and from pirouettes-to-runs $\kappa_{P \rightarrow R}$, finding close agreement between data and simulations Fig. 6(C).

The accuracy of the Markov dynamics suggests that it is also possible to recover the diffusive properties of foraging trajectories. At the coarse-grained level, the symbolic sequences become binary sequences of run and pirouette states, with differing durations Δt . We generate centroid trajectories by first sampling from the set of run and pirouette trajectories for each duration Δt , S7(A). We then transform symbolic sequences into centroid trajectories by aligning and appending run and pirouette trajectories, sampled from the space of centroid trajectories with the appropriate duration (Methods). This yields simulated centroid trajectories that maintain the trajectory characteristics within each macrostate, but otherwise destroy correlations between them. In Fig. 6(D-left), we show an example centroid trajectory from the data (black), an example trajectory generated from the Markov model (blue) and an example trajectory generated from a random shuffle that obeys the stationary distribution π (green): the model generated trajectory is qualitatively more realistic than that of the shuffle, even though the shuffle respects the underlying fraction of runs and pirouettes.

To quantify the similarity between the simulated centroid trajectories and the data, we estimate the mean squared displacement $\langle (\vec{X}(t+\tau) - \vec{X}(t))^2 \rangle$ (MSD), which exhibits a transition between super-diffusive (nearly ballistic) and diffusive behavior between 10s and 100s [78–81], Fig. 6(D-middle), S7(B). The operator-based simulations accurately capture the MSD across a wide range of scales, including the ballistic-to-diffusive transition. To further assess the quality of the simulations, we estimate an effective diffusion coefficient by fitting $\text{MSD} = D\tau$ in the linear regime [82] and find that, across worms, the resulting diffusion coefficients obtained from simulations closely follow the data, Fig. 6(D-right). In contrast, centroid trajectories obtained from symbolic sequences that respect the underlying invariant density π but are otherwise random fail to mimic the extent to which worms explore the environment Fig. 6(D-left), which results in an underestimation of the effective diffusion coefficients, Fig. 6(D-right). Sampling from the underlying centroid

trajectories is equivalent to having an accurate mechanical model of the interaction between the body wave and the medium and our results demonstrate that it is possible to go from microscopic posture dynamics to diffusive properties in a living organism. Despite fine-scale differences in the dynamics of each worm, a single τ^* allows for an accurate prediction of the diffusive properties across the population.

Subdivisions in the long-lived states: “run(s)-and-pirouette”

While the operator derived dynamics accurately capture the average kinetic rates in the 2-state coarse-graining, the dwell time distribution of run states exhibits, in most worm recordings, heavy tails not captured by a 2-state subdivision, as exemplified in Fig. 7(A) and summarized in Fig. S8(A). To explore the possibility that these heavy tails are indicative of a finer decomposition relevant for worm behavior, we systematically subdivide the metastable states. At each step, the metastable state with the largest measure is subdivided along the first non-trivial eigenvector of the reversibilized transition matrix conditioned solely on the microstates within the metastable state, Fig. 7(B) and (Methods). As before, we partition according to the corresponding first non-trivial eigenfunction, maximizing the coherence measure χ , Eq. (11), defined only on the partitions of the highest measure macrostate. This yields a subdivision of the state space that obeys the structure of the free energy landscape; at each iteration, we subdivide the system along the largest energy barrier within the highest measure basin. Similar ideas have been used previously [83], where the notion of relatively coherent sets has been introduced. We note that our metastate subdivision process proceeds *from the longest-lived states down* rather than from the shortest-lived states up, where the latter is more common in temporal clustering approaches [84–86].

Applied to *C. elegans*, the next iterations, beyond the initial subdivision into runs and pirouettes, yield 4 distinct run states: first the run state is split into fast and slow runs, and then each of these is subdivided into two distinct states. This 5-state subdivision allows us to better capture the dwell time distributions Fig. 7(C), with a Kolmogorov-Sinai test across worms showing that the dwell time distributions obtained from simulations are indistinguishable from the data with a 5-state partition, whereas with a 2-state partition some dwell time distributions in the run state are significantly different from the data, Fig. S8(A).

The four “run” states reveal important fine scale differences in the body wave dynamics, Fig. 7(D, E), further detailed in Fig. S8(B, C). The subdivision of the fast states yields distinct gaits with distinct body wavelengths but comparable speeds; we show the body curvature as function of time for two example runs of the fast state in Fig. 7(D-left) and the respective speeds Fig. 7(E-

top). The longer wavelength state typically results in higher curvature runs compared to the other fast state, Fig. 7(E-bottom). Accordingly, the largest wavelength fast state exhibits larger coiling amplitudes when compared to the second fast state. The slow states exhibit larger amplitude head and tail excitations (as measured by the fourth eigenworm coefficient, a_4), when compared to the fast states. In addition, the first slow state exhibits a dorsal bias (negative a_3), which results in the gradual reorientation of the worm through head-casts, Fig. 7(D-right). In contrast, the second slow state is more akin to a pure pause state, for which there are only minor posture changes that do not result in a significant displacement of the worm, Fig. 7(D-right). Such fine-grained states are essential in capturing the posture kinematics, but have not yet been characterized in the *C. elegans* literature, and add an important component for further quantitative understanding in this important and accessible model organism.

DISCUSSION

When seeking order in complex systems, many simplifications are possible, and this choice has consequences for both the nature and complexity of the resulting model. Often such simplifications are chosen *a priori*, and the complexity of the model is taken to be a fact of the underlying dynamics. But this is not generally true; the definition of the states and the resulting dynamics between them are inextricably linked. Here we choose the simplification informed by their resulting model dynamics; we define states in an attempt to achieve a Markov model of the state dynamics. While this class may seem restrictive, we find models of the dynamics that are accurate across a range of time scales and systems. Even for dynamics that do not ultimately permit a Markovian description, we offer a principled coarse graining to independently investigate long time scales.

Applicable to both deterministic and stochastic systems, we combine state space reconstruction with ensemble dynamics to infer probabilistic state space transitions directly from data, which we apply to the Langevin dynamics of a particle in a double-well, the Lorenz system, and the wiggling of the nematode worm *C. elegans*. Rather than seek low-dimensional descriptions of the data directly, as is often the choice in cluster-based analysis (e.g. [80, 87–89]), we instead first *expand* in representation complexity: enlarging the putative state to short sequences and constructing a *maximum* entropy partition to include as much predictive information as possible in the embedding. For dynamical systems with a continuous state space, maximizing the partition entropy can also approximate a generating partition, which preserves ergodic characteristics [90].

By characterizing nonlinear dynamics through transitions between state space partitions, we trade trajectory-based techniques for the analysis of linear operators.

This approach is complementary to previous methods of state space reconstruction which focus on geometric and topological quantities such as Lyapunov exponents, the Kolmogorov-Sinai entropy and dimensions [34, 69, 76, 91, 92]. In particular, trajectory-based techniques of ergodic analysis rely on precise estimates of dimension and of the local Jacobian, which can be challenging in systems with unknown equations.

In our approach, we recognize and address the mutual dependence of state space reconstruction and Markovian evolution: reconstructing the state space is required for an effective Markovian description and the framework of ensemble dynamics provides the principled, information-theoretic measure of memory used to optimize the reconstructed state. Previous work in molecular and neural dynamics [28, 93] have also implicitly expanded states in time in order to improve the accuracy of Markov models. While a Markov approximation is valid for a wide class of physical systems, hidden or slow variables can render the data non-stationary, in which case the transfer operator has an explicit time dependence $\mathcal{P}_\tau(t)$. Nonetheless, a τ -parametrized family of coherent sets can capture *moving* regions of the state space that remain coherent within a time scale τ [56, 94–98]. Even for non-Markovian systems our approach can help identify and isolate the long-lived dynamics [99, 100].

To optimally identify almost invariant sets, we use the reversibilized operator P_τ , Eq. (10), which is generally different than P . For an overdamped system in thermodynamic equilibrium, this symmetrization simply enforces detailed balance and $|\hat{\Lambda}_2^{-1}|$ is directly related to the hopping rates between metastable states. Irreversible dynamics, however, result in complex eigenvalues, which cannot be directly associated with relaxation time scales. Generally, the dynamics of P_τ will relax more slowly than those of P , providing an upper bound to the true relaxation times [101, 102]. We thus assume that the Markovianity of the metastable dynamics of P is reflected in P_τ , and we find that is a good approximation in our examples, Fig. S9. It will also be interesting to explore approaches based directly on P [56].

When the state-space is known or its reconstruction decoupled from the ensemble dynamics, mesh-free discretizations have been used to characterize ensemble evolution, including diffusion maps [93, 103] and other kernel based approaches, such as Reproducing Kernel Hilbert Spaces [104] or Extended Dynamic Mode Decomposition [105]. Though powerful, such methods also require subtle choices in kernels, neighborhood length scales and transition times, for which we lack guiding principles.

A particularly interesting focus of maximally-predictive ensemble dynamics is in the quantitative understanding of animal behavior [13, 14, 106]. This may seem paradoxical; organism-scale behavior arises as the cumulative output from a wide range of influences, including genetic, neural, and bio-mechanical systems, all of which may in principle introduce important hidden-

states (see e.g. [107, 108]) and thus non-Markovian dynamics [88, 109]. Yet, the process of constructing behavioral states has often been decoupled from the analysis of the dynamics and “under-embedding” could easily introduce apparent memory.

In the behavior of the nematode *C. elegans*, maximally predictive ensemble dynamics connects sub-second posture dynamics to long-lasting, compound behavioral states. The worm embedding combines postures across roughly a quarter of the duration of a worm’s typical body wave and results in delay and entropy estimates in agreement with previous work [76]. On longer timescales, the posture-based “run-and-pirouette” navigation strategy [77, 110] derived from the operator dynamics provide an accurate and principled discretization of foraging behavior, disentangling motions that are confounded by centroid-derived measurements (see e.g. [111]). For example, the different “gaits” that result from the subdivision of the run state exhibit comparable centroid speeds and are only clearly distinguished due to the different body wavelengths. These emergent Markov dynamics offer a promising and powerful demonstration of quantitative connections across the hierarchy of movement behavior generally exhibited by all organisms [88, 112]. Particularly interesting future directions include the analysis of even longer dynamics in *C. elegans* [113–116], for which no canonical behavioral states have been described.

The principled integration of fluctuating, microscopic dynamics resulting in coarse-grained but effective theories is a remarkable success of statistical physics. Conceptually similar, our work here is designed towards systems sampled from data whose fundamental dynamics are unknown. We leverage, rather than ignore, small-scale variability to subsume nonlinear dynamics into linear, ensemble evolution, enabling the principled identification of coarse-grained, long-timescale processes, which we expect to be informative in a wide variety of systems.

ACKNOWLEDGEMENTS

We thank Massimo Vergassola and Federica Ferretti for comments. This work was supported by OIST Graduate University (TA, GJS), a program grant from the Netherlands Organization for Scientific Research (AC, GJS), by the Herchel Smith Fund (DJ), and by Vrije Universiteit Amsterdam (AC, GJS).

METHODS

Software and data availability: Code for reproducing our results is publicly available: https://github.com/AntonioCCosta/predictive_ensemble_dynamics/. Data can be found in [117].

State space reconstruction: Given a measurement

time series, $\vec{y}(t)$, with $t \in \{\delta t, \dots, T\delta t\}$ and $\vec{y} \in \mathbb{R}^d$, we build a state trajectory matrix by stacking K time-shifted copies of \vec{y} to give a $(T - K) \times Kd$ state matrix X_K . For each K , we partition the candidate state space and estimate the entropy rate of the associated Markov chain (see below). We choose K^* such that $\partial_K h_{PN}(K^*) \sim 0$.

State space partitioning: We partition the state space into N Voronoi cells, $A_i, i \in \{1, \dots, N\}$, through k-means clustering with a k-means++ initialization using scikit-learn [118].

Approximation of the Perron-Frobenius operator: We build a finite dimensional approximation of the Perron-Frobenius operator using an Ulam-Galerkin discretization. A Galerkin projection takes the infinite dimensional operator onto an $N \times N$ operator of finite rank by truncating an infinite dimensional set of basis functions at N . Ulam’s method uses characteristic functions as the basis for this projection,

$$\zeta_i(x) = \begin{cases} 1, & \text{for } x \in A_i \\ 0, & \text{otherwise} \end{cases}. \quad (15)$$

Our characteristic functions are implicitly defined through the k-means discretization of the space. We thus partition the space into N connected sets with nonempty and disjoint interior that covers M : $M = \cup_{i=1}^N A_i$, and approximate the operator as a Markov chain by counting transitions from A_i to A_j in a finite time τ . Given T observations, a set of N partitions, and a transition time τ , we compute

$$C_{ij}(\tau) = \sum_{t=0}^{T-\tau} \zeta_i(x(t))\zeta_j(x(t+\tau)),$$

where $\zeta_i(x)$ are the Ulam basis functions, Eq. (15). The maximum likelihood estimator of the transition matrix is obtained by simply row normalizing the count matrix,

$$P_{ij}(\tau) = \frac{C_{ij}(\tau)}{\sum_j C_{ij}(\tau)},$$

which yields an approximation of the Perron-Frobenius operator.

Invariant density estimation: Given a transition matrix P , the invariant density is obtained through the left eigenvector of the non-degenerate eigenvalue 1 of P , $\pi P = \pi$: π_i is the probability of finding the system in a partition A_i .

Entropy estimation: Given a partition of the state space into regions of size ϵ , the probability of the system visiting M consecutive cells is given by the joint probability $p(s_{i_1}, s_{i_2}, \dots, s_{i_M})$, where $i_k \in \{0, \dots, N\}$ and N is

the number of partitions of the state space. The M -block Shannon entropies are defined as the sum over possible sequences, $H_M = -\sum p(s_{i_1}, \dots, s_{i_M}) \log p(s_{i_1}, \dots, s_{i_M})$, measured in *nats*. As in a thermodynamic system, for which entropy grows with the system size, the entropy of a symbolic sequence grows with the sequence length. The rate of entropy growth is defined through the entropy of the condition probability of the system visiting the cell s_{M+1} given that it has visited the previous M cells: $p(s_{M+1}|s_{i_1}, \dots, s_{i_M})$,

$$\begin{aligned} h_M &= \langle -\log [p(s_{M+1}|s_{i_1}, \dots, s_{i_M})] \rangle \\ &= H_{M+1} - H_M. \end{aligned} \quad (16)$$

The limit $h = \lim_{M \rightarrow \infty} h_M$ is the entropy rate of the source and measures the overall unpredictability of the dynamics. It is a non-decreasing function of number of partitions N , being bounded from above by $h \leq \log N$. To characterize the underlying dynamical process, one takes the supremum of the entropy rate over all possible partitions, thus obtaining the Kolmogorov-Sinai entropy rate,

$$h_{KS} = \sup_{\mathbb{P}} \lim_{M \rightarrow \infty} h_M(\mathbb{P}), \quad (17)$$

where \mathbb{P} is a candidate partition of the state space.

One approach to the estimation of the entropy rate, more common among the information theory community, is to estimate the M -block entropy rates for increasing M , given a fixed small number of partitions N . This entails counting all sequences of length M and estimating the joint probabilities $p(s_{i_1}, \dots, s_{i_M})$. However, high dimensional non-linear dynamical systems generally require a large number of partitions to approach the Kolmogorov-Sinai limit. Given this constraint, computing sequence probabilities for reasonably large sequence lengths M quickly becomes infeasible since the number of possible sequences grows with N^M .

Another approach, more common in the dynamical systems literature, is to build blocks of K delays and estimate correlation entropies, introduced by Cohen and Procaccia (CP) [40]. Given a measurement time series $\{y(t)\}$ where $t \in \{\tau, 2\tau, \dots, T\tau\}$, ϵ -tubes are built around sequences of length $K\tau$, $y^K(t)$, and the entropy is estimated as,

$$H_{K\tau}^C(\epsilon) = -\left\langle \log C_{K\tau} \left(\frac{\epsilon}{2} \right) \right\rangle, \quad (18)$$

where $C_{K\tau}(\frac{\epsilon}{2})$ is the correlation function of K -dimensional ϵ -tubes built from time series measurements with a sampling time of τ , which essentially measures the probability of finding two $y^K(t)$ vectors within a distance ϵ of each other. This approach avoids the counting problem of direct estimates for large N (small ϵ) and K .

The correlation entropy rate can then be obtained by estimating,

$$h_\tau^C(\epsilon) = \lim_{K \rightarrow \infty} \left[H_{(K+1)\tau}^C(\epsilon) - H_{K\tau}^C(\epsilon) \right], \quad (19)$$

which converges to the Kolmogorov-Sinai entropy in the limit $\epsilon \rightarrow 0$ and $\tau \rightarrow 0$. It is important to note that correlations are typically measured using a Chebyshev distance, $d(X, Y) = \max_i^D |X_i - Y_i|$ for $X, Y \in \mathbb{R}^D$ and $|\cdot|$ the absolute value, and so the precise values of $H_{K\tau}^C(\epsilon)$ may depend on the chosen metric.

As described in the main text, we present here a hybrid approach to the estimation of entropy. As with the CP method, we add time into the definition of state (instead of increasing the length of the sequence of discrete states), but then partition that state space and use the inferred Markov chain to estimate the entropy rate of the dynamics for increasing K , Eq. (5). We illustrate the similarities between our partition based approach and the geometric approach using an autoregressive process of order p , AR(p), as in [119]. The behavior of the entropy as a function of the number of delays is an indication of the memory of the autoregressive process. We use a simple AR(2) model, with

$$x_{t+1} = 1.88x_t - 0.95x_{t-1} + \eta_t, \quad (20)$$

where η_t is a delta correlated Gaussian random variable with unit variance. In Fig. S10 we show the resulting estimates of the entropy rate and predictive information using the CP method of Eq. (19) [40] and our partition based approach Eq. (5). The behavior of the entropy as a function of the number of delays is qualitatively similar between the two methods, despite the dramatic differences in metrics and estimation procedures. The entropy rate grows faster with the decreasing length scale for $K = 1$ and the behavior of h with the length scale for $K \gtrsim 2$ is $h(K, \epsilon) = h(\infty, \epsilon) - \log \epsilon$ with identical $h(\infty, \epsilon)$. The difference between $h(K = 1, \epsilon)$ and $h(K = 2, \epsilon)$ stems from the memory of the process. Similar conclusions can be taken by looking at the predictive information estimates, $I_{\text{pred}} = H - h$: the system exhibits a higher degree of predictability when $K > 1$, indicating that it is a second order process. We can therefore recover the dimensionality of the underlying state space, which in this case simply corresponds to the order of the autoregressive process.

Identifying metastable states: Metastable states correspond to regions of the state space that the system visits often, separated by regions where transitions occur. We therefore search for collections of states that move coherently with the flow: their relative future states belong to the same macroscopic set. As discussed in the Introduction, we leverage a heuristic introduced in [51], which makes use of the first non-trivial eigenfunction of the transfer operator to identify almost invariant

sets. As in [67], we will work with a time reversibilized transfer operator \mathcal{P}_r : the coherence properties are invariant to this transformation [60] and the analysis is simplified due to the optimality properties of reversible Markov chains. In discrete time and space, P_r is defined as

$$P_r(\tau) = \frac{P(\tau) + P(-\tau)}{2}, \quad (21)$$

where,

$$P_{ij}(-\tau) = \frac{\pi_j P_{ji}(\tau)}{\pi_i}.$$

is the stochastic matrix governing the time-reversal of the Markov chain. The first nontrivial ($\lambda < 1$) right eigenvector of P_r , ϕ_2 , allows us to define the macrostates as

$$S^+ := \bigcup_{i:\phi_2 \geq \phi_2^c} A_i, S^- := \bigcup_{i:\phi_2 < \phi_2^c} A_i, \quad (22)$$

where ϕ_2^c is chosen so as to maximize χ , Eq. 11, and $\chi_{\mu,\tau}$ are computed in a discretized state space,

$$\chi_{\pi,\tau}(S) = \frac{\sum_{i,j \in S} \pi_i P_{ij}(\tau)}{\sum_{i \in S} \pi_i}. \quad (23)$$

In practice, we compute Eq. 11 for all ϕ_2^c (S^+ and S^- are implicit functions of ϕ_2^c) and find the global maximum: this is an inexpensive calculations since χ can be obtained by matrix multiplications and ϕ_2^c can only take N different values, where N is the number of partitions. We note that for an overdamped thermodynamic system in equilibrium, the eigenvectors ϕ_k are discrete approximations of the eigenfunctions of the Koopman operator [105], which are nearly constant within a metastable set thus simplifying the clustering into almost invariant sets.

Operator-based state space subdivision: We leverage the notion of relatively coherent sets [83] to subdivide the state space. However, instead of subdividing both metastable state at each iteration, we identify the state with the most measure S' and build a new transition matrix only with partitions belonging to that state,

$$P_{S'}(\tau) = p(s_j(t + \tau) | s_i(t)), i, j \in S'.$$

From $P_{S'}$ we proceed as before: we compute the stationary distribution of S' through the first left eigenvector of $P_{S'}$, π'_i , build the corresponding reversibilized transition matrix $P_{r,S'}$ and identify relatively metastable states through its first non-trivial eigenvector by maximizing Eq. (9) where π_i and $P_{ij}(\tau)$ are replaced by their

relative counterparts π'_i and $P_{S'}$. The iterative process is illustrated in Fig. 7(A).

Choice of transition time τ^* : We choose τ^* as the shortest transition time scale after which the inferred implied relaxation times reach a plateau. For τ too short, the approximation of the operator will yield a transition matrix that is nearly identity (due to the finite size of the partitions and too short transition time), which results in degenerate eigenvalues close to $\lambda \sim 1$: an artifact of the discretization and not reflective of the underlying dynamics. For τ too large, the transition probabilities become indistinguishable from noisy estimates of invariant density, which results in a single surviving eigenvalue $\lambda_1 = 1$ while the remaining eigenvalues converge to a noise floor resulting from a finite sampling of the invariant density. Between such regimes, we find a region with the largest time scale separation (as illustrated in Fig. 1(C-middle)) which also corresponds to the regime for which the longest relaxation times t_{imp} , Eq. 8, are robust to the choice of τ . We compute t_{imp} using the eigenvalues of the reversibilized transition matrix, P_r , which only gives an upper bound to the relaxation dynamics. In addition, we assume that P_r can only be Markovian if also P is Markovian, a statement that holds for the systems studied here, Fig. S9.

Matrix diagonalization: The high dimensionality and the sparsity of the transition matrices for large N results in numerical errors when using a naive estimator for the full spectrum of eigenvalues. In addition, since we are interested in the longest lived dynamics, we focus on finding only the n_{modes} largest magnitude real eigenvalues using the ARPACK [120] algorithm.

Periodic Orbit identification: We identify the shortest period unstable periodic orbit of the Lorenz system, Eq. (14), by studying the distribution of recurrence times. We set a short distance ϵ and look for the times p at which $\|X_{t+p} - X_t\| < \epsilon$, where $\|\cdot\|$ represents the Euclidean distance. In practice, we compute $1/\|X_{t+p} - X_t\|^2$ and find peaks with height larger than $1/\epsilon^2$ using the `find_peaks` function from the `scipy.signal` package [121]. For short enough ϵ ($\epsilon = 5 \times 10^{-6}$ in our case), the distribution of recurrence times has its first peak at the period of the shortest unstable periodic orbit [122]. A trajectory corresponding to a shadow of this unstable periodic orbit is shown in Figs. 4C, S3C [123, 124].

C. *elegans* centroid trajectory simulations: Given a symbolic sequence, we generate centroid trajectories by appending and aligning “run” and “pirouette” trajectories independently drawn from the set of all possible trajectories of a given duration. In the rare occasion that the generated symbolic sequences yield state durations that are longer than all the observed trajectories, we append multiple randomly sampled trajectories up to the corresponding duration. We align the beginning of a new sampled centroid trajectory with the end of the previous one in the following way: we average the velocity

vectors in the $2\tau^* = 1.5$ s segment preceding the sampled trajectory, estimate its orientation angle and rotate the trajectory such that the initial orientation matches the angle defined from the ending $2\tau^*$ segment of the previous trajectory. Trajectories with a duration less than $2\tau^*$ were discarded as they do not allow for an appropriate alignment. This process generates continuous centroid trajectories with the same overall duration as the experiments, but with independent segments randomly sampled from the population of “run” or “pirouette” trajectories of a given duration.

Relaxation times for the non-reversibilized dynamics: We choose τ^* through the *real* spectrum of the reversibilized transition matrix P_r , looking for a constant implied relaxation time scale $|\Lambda_2^{-1}|$, which sets an upper bound to the underlying relaxation dynamics. We thus make an indirect test for the Markovianity of P , assuming that P_r is only Markovian when P is also Markovian. Estimating time scales directly from P is non-trivial due to the possibility of complex eigenvalues. However, we can probe the relaxation dynamics by estimating the first passage time distribution between coarse-grained states directly from the full Markov chain. Coarse-graining along the first non-trivial eigenvector of the reversibilized transition matrix ϕ_2 , we obtain two macrostates a and b with measures π_a and π_b respectively. The dwell time distribution in these states $f_\alpha(t)$ can be obtained directly from the full non-reversibilized Markov chain P : the system would have to hop from β to α , stay in α for $t = n\tau$ steps and then hop back to state β , which yields $f_\alpha(n\tau) = (P_{\beta\alpha}P_{\alpha\alpha}^n P_{\alpha\beta}\pi_\beta)/(P_{\alpha\beta}\pi_\beta)$, where $P_{\beta\alpha}(\tau)$ is the probability of transitioning from a microstate in macrostate α to a microstate in macrostate β in a transition time τ , $\alpha, \beta \in \{a, b\}$. The mean first passage time can then be obtained by integrating $t_{\text{MFP}} = \int_0^\infty t f(t) dt$ and the decay rate to the invariant density is given by $|\Lambda_2| = 1/t_{a \rightarrow b} + 1/t_{b \rightarrow a}$. However, this calculation is challenging to do numerically, due to the large number of matrix multiplications needed to reach the $t = n\tau \rightarrow \infty$ limit and the low accuracy of the Riemann integral for large τ . An equivalent yet simpler calculation can be done by leveraging the eigenvalue decomposition of the infinitesimal dynamics: the first passage time distribution simply becomes $f(t) \propto \sum_i^N e^{\Lambda_i t}$, and the mean first passage time yields $t_{\text{MFP}} = \int_0^\infty t f(t) dt \propto \sum_i^N \Lambda_i^{-2}$, where $\Lambda_i \in \mathbb{R}^+$ are the eigenvalues of the infinitesimal generator of the dynamics. For a discrete spectrum with a time scale separation, the mean first passage time will be governed by the first non-trivial eigenvalue, $\Lambda_2 \sim 0$, $t_{\text{MFP}} \sim \Lambda_2^{-2}$, since the other eigenvalues $\Lambda_i \ll 0$ for $i > 2$ only add minor contributions. Therefore, a simpler yet equivalent estimate of the slow kinetic rates can be obtained by first coarse-graining the system into its macrostates a and b , and then computing the first non-trivial eigenvalue of the coarse-grained two-state Markov chain P_c . The corresponding relaxation times for our example applications are shown in Fig. S9.

Double-well simulation: We use an Euler-Maruyama integration scheme to simulate Eq. (13) and generate a $T = 10^6$ s long trajectory of a particle in a double-well potential with m and γ set to unity and $\beta^{-1} = \{0.5, 0.75, 1.00, 1.25, 1.5, 1.75, 2.00, 2.25, 2.5\}$ J. We first sampled at 100 Hz, and then downsampled to a sampling time of $\delta t = 0.05$ s. In addition, the first 1000 s were discarded to avoid transients. In Fig. 2(B) and S2(C), we project the Boltzmann distribution into the SVD space spanned by $[u_1, u_2]$ by learning a linear mapping θ , between $[u_1, u_2]$ and $[x, v]$: $[u_1, u_2] = \theta \cdot [x, v]$. This allows us to project the centroids of each Voronoi cell in the $[u_1, u_2]$ space to the $[x, v]$ space, and estimate the corresponding Boltzmann weight $p([u_1, u_2] \cdot \theta^{-1}) = p([x, v]) = \exp\{-\beta [(x^2 - 1)^2 + v^2/2]\}/Z$, where Z represent the normalization.

Lorenz system simulation: We use scipy’s `odeint` package [121] to generate a $T = 5 \times 10^5$ s long trajectory of the Lorenz system in the standard chaotic regime $(\sigma, \rho, \beta) = (10, 28, 8/3)$, sampled at $\delta t = 10^{-2} \text{ s}^{-1}$. We discard the first 10^3 s to avoid transients.

C. elegans foraging dataset: We used a previously-analyzed dataset [74], in which N2-strain *C. elegans* were imaged at $f = 32$ Hz with a video tracking microscope on a food-free plate and downsampled to $f = 16$ Hz to incorporate coiled postures [73]. Worms were grown at 20°C under standard conditions [125]. Before imaging, worms were removed from bacteria-strewn agar plates using a platinum worm pick, and rinsed from *E. coli* by letting them swim for 1 min in NGM buffer. They were then transferred to an assay plate (9 cm Petri dish) that contained a copper ring (5.1 cm inner diameter) pressed into the agar surface, preventing the worm from reaching the side of the plate. Recording started approximately 5 min after the transfer, and lasted for 2100 s, for a total of $T = 33600$ frames.

- [1] L. E. Reichl, *A modern course in statistical physics; 4th ed.* (Wiley-VCH, Weinheim, 2016).
- [2] N. D. Goldenfeld, *Lectures on phase transitions and the renormalization group*, Frontiers in Physics (Addison-Wesley, Reading, MA, 1992) this book has also been published by CRC Press in 2018.
- [3] B. B. Machta, R. Chachra, M. K. Transtrum, and J. P. Sethna, *Parameter Space Compression Underlies Emergent Theories and Predictive Models*, *Science* **342**, 604 (2013).
- [4] M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, *Perspective: Slowness and emergent theories in physics, biology, and beyond*, *The Journal of Chemical Physics* **143**, 010901 (2015).
- [5] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, *Coarse graining, fixed points, and scaling in a large population of neurons*, *Physical review letters* **123**, 178103 (2019), pMC7335427[pmcid].
- [6] S. Bradde and W. Bialek, *PCA Meets RG*, *Journal of Statistical Physics* **167**, 462 (2017).
- [7] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, *A high-bias, low-variance introduction to machine learning for physicists*, *Physics Reports* **810**, 1 (2019).
- [8] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Machine learning and the physical sciences*, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [9] N. G. Van Kampen, *Stochastic processes in physics and chemistry* (North-Holland, Amsterdam, 1981).
- [10] H. A. Posch and R. Hirschl, Simulation of billiards and of hard body fluids, in *Hard Ball Systems and the Lorentz Gas*, edited by D. Szász (Springer Berlin Heidelberg, Berlin, Heidelberg, 2000) pp. 279–314.
- [11] L. Wiskott and T. J. Sejnowski, *Slow feature analysis: Unsupervised learning of invariances*, *Neural Computation* **14**, 715 (2002).
- [12] J. D. Chodera and F. Noé, *Markov state models of biomolecular conformational dynamics*, *Current Opinion in Structural Biology* **25**, 135 (2014).
- [13] G. J. Berman, *Measuring behavior across scales*, *BMC Biol.* **16** (2018), 10.1186/s12915-018-0494-7.
- [14] A. X. Brown and B. de Bivort, *Ethology as a physical science*, *Nature Physics* **14**, 653 (2018).
- [15] S. R. Datta, D. J. Anderson, K. Branson, P. Perona, and A. Leifer, *Computational Neuroethology: A Call to Action*, *Neuron* **104**, 11 (2019).
- [16] A. J. Calhoun and M. Murthy, *Quantifying behavior to solve sensorimotor transformations: advances from worms and flies*, *Curr. Opin. Neurobiol.* **46**, 90 (2017).
- [17] F. Bartumeus, M. G. E. da Luz, G. M. Viswanathan, and J. Catalan, *Animal search strategies: A quantitative random-walk analysis*, *Ecology* **86**, 3078 (2005).
- [18] M. Vergassola, E. Villermaux, and B. I. Shraiman, *Infotaxis as a strategy for searching without gradients*, *Nature* **445**, 406 (2007).
- [19] V. Sourjik and N. S. Wingreen, *Responding to chemical gradients: bacterial chemotaxis*, *Current Opinion in Cell Biology* **24**, 262 (2012).
- [20] D. Auerbach, P. Cvitanović, J.-P. Eckmann, G. Gunaratne, and I. Procaccia, *Exploring chaotic motion through periodic orbits*, *Physical Review Letters* **58**, 2387 (1987).
- [21] H. Risken and H. Haken, *The Fokker-Planck Equation: Methods of Solution and Applications Second Edition* (Springer, 1989).
- [22] P. Gaspard, *Chaos, Scattering and Statistical Mechanics*, Cambridge Nonlinear Science Series (Cambridge University Press, 1998).
- [23] D. Givon, R. Kupferman, and A. Stuart, *Extracting macroscopic dynamics: Model problems and algorithms*, *Nonlinearity* **17**, 1 (2004).
- [24] G. Pavliotis, *Stochastic Processes and Applications Diffusion Processes, the Fokker-Planck and Langevin Equations* (Springer, 2014).
- [25] A. Lasota and M. Mackey, *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*, 2nd ed., Vol. 97 (Springer-Verlag New York, 1994).
- [26] E. M. Boltt, A. Luttmann, S. Kramer, and R. Basnayake, *Measurable Dynamics Analysis of Transport in the Gulf of Mexico During the Oil Spill*, *International Journal of Bifurcation and Chaos* **346** (2011), 10.1142/s0218127412300121.
- [27] E. M. Boltt and N. Santitissadeekorn, *Applied and computational measurable dynamics* (Society for Industrial and Applied Mathematics, Philadelphia, United States, 2013).
- [28] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *Identification of slow molecular order parameters for markov model construction*, *The Journal of Chemical Physics* **139**, 015102 (2013).
- [29] A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz, and C. Schütte, *Transition Manifolds of Complex Metastable Systems: Theory and Data-Driven Computation of Effective Dynamics*, *Journal of Nonlinear Science* **28**, 471 (2018).
- [30] W. E and E. Vanden-Eijnden, *Transition-path theory and path-finding algorithms for the study of rare events*, *Annual Review of Physical Chemistry* **61**, 391 (2010), pMID: 18999998.
- [31] G. Froyland, G. A. Gottwald, and A. Hammerlindl, *A Computational Method to Extract Macroscopic Variables and Their Dynamics in Multiscale Systems*, *SIAM Journal on Applied Dynamical Systems* **13**, 1816 (2014).
- [32] S. H. Strogatz, *Nonlinear Dynamics and Chaos with Student Solutions Manual: With Applications to Physics, Biology, Chemistry, and Engineering* (Westview Press, a member of the Perseus Books Group, Boulder, Colorado, USA, 2015).
- [33] D. D. Nolte, *The tangled tale of phase space*, *Physics Today* **63**, 33 (2010).
- [34] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, *Geomery from a Time Series*, *Physical Review Letters* **45**, 712 (1980).
- [35] F. Takens, in *Dynamical Systems and Turbulence, Warwick 1980*, edited by D. Rand and L.-S. Young (Springer Berlin Heidelberg, Berlin, Heidelberg, 1981) pp. 366–381.

- [36] M. Casdagli, S. Eubank, J. D. Farmer, and J. Gibson, *The theory of state space reconstruction in the presence of noise*, *Physica D* **51**, 52 (1991).
- [37] J. Stark, *Delay embeddings for forced systems. I. Deterministic forcing*, *Journal of Nonlinear Science* **9**, 255 (1999).
- [38] J. Stark, D. S. Broomhead, M. Davies, and J. Huke, *Delay embeddings for forced systems. II. Stochastic forcing*, *Journal of Nonlinear Science* **13**, 519 (2003).
- [39] Although first introduced in the context of autonomous deterministic systems, these results have been extended to other types of dynamical systems, including stochastic dynamics [37, 38].
- [40] A. Cohen and I. Procaccia, *Computing the Kolmogorov entropy from time signals of dissipative and conservative dynamical systems*, *Physical Review A* **31**, 1872 (1985).
- [41] P. Gaspard and X.-J. Wang, *Noise, chaos, and (epsilon, tau)-entropy per unit time*, *Physics Reports* **235**, 291 (1993).
- [42] Y. B. Pesin, *Characteristic Liapunov exponents and smooth ergodic theory*, *Russian Mathematical Surveys* **32**, 55 (1977).
- [43] J. Farmer, E. Ott, and J. A. Yorke, *The dimension of chaotic attractors*, *Physica D: Nonlinear Phenomena* **7**, 153 (1983).
- [44] P. Grassberger and I. Procaccia, *Measuring the strangeness of strange attractors*, *Physica D: Nonlinear Phenomena* **9**, 189 (1983).
- [45] P. Grassberger, *Toward a quantitative theory of self-generated complexity*, *International Journal of Theoretical Physics* **25**, 907 (1986).
- [46] W. Bialek, I. Nemenman, and N. Tishby, *Predictability, complexity, and learning*, *Neural Computation* **13**, 2409 (2001).
- [47] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 2nd ed. (McGraw-Hill, New York, 1984) pp. 392–393.
- [48] We note that this is not a sufficient condition for Markovianity, as also the eigenvectors need to be constant with τ .
- [49] N. Djurdjevac, M. Sarich, and C. Schütte, *Estimating the eigenvalue error of Markov State Models*, *SIAM Multiscale Model. Simul.* **10**, 61 (2012).
- [50] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, F. Noé, J.-h. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, *Markov models of molecular kinetics : Generation and validation Markov models of molecular kinetics : Generation and validation*, *The Journal of Chemical Physics* **134** (2011), 10.1063/1.3565032.
- [51] M. Dellnitz and O. Junge, *On the Approximation of Complicated Dynamical Behavior*, *SIAM Journal on Numerical Analysis* **36**, 491 (1999).
- [52] P. Deuffhard, W. Huisinga, A. Fischer, and C. Schütte, *Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains*, *Linear Algebra and Its Applications* **315**, 39 (2000).
- [53] G. Froyland, K. Padberg, M. H. England, and A. M. Treguier, *Detection of coherent oceanic structures via transfer operators*, *Phys. Rev. Lett.* **98**, 224503 (2007).
- [54] M. Dellnitz, G. Froyland, C. Horenkamp, K. Padberg-Gehle, and A. Sen Gupta, *Seasonal variability of the subpolar gyres in the southern ocean: a numerical investigation based on transfer operators*, *Nonlinear Processes in Geophysics* **16**, 655 (2009).
- [55] S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé, *Data-Driven Model Reduction and Transfer Operator Approximation*, *Journal of Nonlinear Science* **28**, 985 (2018).
- [56] K. Fackeldey, P. Koltai, P. Névir, H. Rust, A. Schild, and M. Weber, *From metastable to coherent sets - Time-discretization schemes*, *Chaos* **29** (2019), 10.1063/1.5058128.
- [57] C. Schütte, W. Huisinga, and P. Deuffhard, in *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems*, edited by B. Fiedler (Springer Berlin Heidelberg, Berlin, Heidelberg, 2001) pp. 191–223.
- [58] G. R. Bowman, V. S. Pande, and F. Noé, eds., *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*, *Advances in Experimental Medicine and Biology*, Vol. 797 (Springer Netherlands, Dordrecht, 2014).
- [59] G. Froyland and M. Dellnitz, *Detecting and locating near-optimal almost-invariant sets and cycles*, *SIAM Journal on Scientific Computing* **24**, 1839 (2003).
- [60] G. Froyland, *Statistically optimal almost-invariant sets*, *Physica D: Nonlinear Phenomena* **200**, 205 (2005).
- [61] A. Gloter, *Parameter estimation for a discretely observed integrated diffusion process*, *Scandinavian Journal of Statistics* **33**, 83 (2006).
- [62] G. A. Pavliotis and A. M. Stuart, *Parameter estimation for multiscale diffusions*, *Journal of Statistical Physics* **127**, 741 (2007).
- [63] F. Ferretti, V. Chardès, T. Mora, A. M. Walczak, and I. Giardina, *Building general langevin models from discrete datasets*, *Phys. Rev. X* **10**, 031018 (2020).
- [64] D. B. Brückner, P. Ronceray, and C. P. Broedersz, *Inferring the dynamics of underdamped stochastic systems*, *Phys. Rev. Lett.* **125**, 058103 (2020).
- [65] H. Risken and K. Voigtlaender, *Eigenvalues and eigenfunctions of the Fokker-Planck equation for the extremely underdamped Brownian motion in a double-well potential*, *Journal of Statistical Physics* **41**, 825 (1985).
- [66] J. P. Eckmann and D. Ruelle, *Ergodic theory of chaos and strange attractors*, *Reviews of Modern Physics* **57**, 617 (1985).
- [67] G. Froyland and K. Padberg, *Almost-invariant sets and invariant manifolds - Connecting probabilistic and geometric descriptions of coherent structures in flows*, *Physica D: Nonlinear Phenomena* **238**, 1507 (2009).
- [68] S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. Nathan Kutz, *Chaos as an intermittently forced linear system*, *Nat. Commun.* **8**, 1 (2017).
- [69] D. P. Lathrop and E. J. Kostelich, *Characterization of an experimental strange attractor by periodic orbits*, *Phys. Rev. A* **40**, 4028 (1989).
- [70] J. D. Farmer and J. J. Sidorowich, *Predicting chaotic time series*, *Phys. Rev. Lett.* **59**, 845 (1987).
- [71] S. Brenner, *The Genetics of Caenorhabditis elegans*, *Genetics* **77**, 71 (1974).
- [72] C. I. Bargmann and E. Marder, *From the connectome to brain function*, *Nat Methods* **10** (2013), 10.1016/j.cub.2012.01.061.
- [73] O. D. Broekmans, J. B. Rodgers, W. S. Ryu, and G. J.

- Stephens, *Resolving coiled shapes reveals new reorientation behaviors in C. elegans*, eLife **5**(e17227) (2016), 10.7554/eLife.17227.
- [74] G. J. Stephens, B. Johnson-Kerner, W. Bialek, and W. S. Ryu, *Dimensionality and dynamics in the behavior of C. elegans*. PLoS Comput. Biol. **4**, e1000028 (2008).
- [75] S. P. Strong, R. Koberle, R. R. de Ruyter van Steveninck, and W. Bialek, *Entropy and Information in Neural Spike Trains*, Phys. Rev. Lett. **80**, 197 (1998).
- [76] T. Ahamed, A. C. Costa, and G. J. Stephens, *Capturing the continuous complexity of behaviour in Caenorhabditis elegans*, Nature Physics **17**, 275 (2021).
- [77] J. T. Pierce-Shimomura, T. M. Morse, and S. R. Lockery, *The fundamental role of pirouettes in Caenorhabditis elegans chemotaxis*. J. Neurosci. **19**, 9557 (1999).
- [78] S. J. Helms, W. M. Rozemuller, A. C. Costa, L. Avery, G. J. Stephens, and T. S. Shimizu, *Modelling the ballistic-to-diffusive transition in nematode motility reveals variation in exploratory behaviour across species*, Journal of The Royal Society Interface **16**, 20190174 (2019).
- [79] G. J. Stephens, B. Johnson-Kerner, W. Bialek, and W. S. Ryu, *From modes to movement in the behavior of Caenorhabditis elegans*. PLoS One **5**, e13914 (2010).
- [80] L. C. Salvador, F. Bartumeus, S. A. Levin, and W. S. Ryu, *Mechanistic analysis of the search behaviour of Caenorhabditis elegans*, Journal of the Royal Society Interface **11** (2014), 10.1098/rsif.2013.1092.
- [81] L. G. Alves, P. B. Winter, L. N. Ferreira, R. M. Briemann, R. I. Morimoto, and L. A. Amaral, *Long-range correlations and fractal dynamics in C. elegans: Changes with aging and stress*, Physical Review E **96**, 1 (2017).
- [82] We note that on longer time scales the MSD exhibits the behavior of a confined random walk due to the rigid boundaries of the agar plate, which makes it non-trivial to accurately estimate the diffusion coefficient [126]. Nonetheless, the estimated effective diffusion coefficients provide a measure of the extent to which worms explore the environment.
- [83] T. Ma and E. M. Bollt, *Relatively coherent sets as a hierarchical partition method*, International Journal of Bifurcation and Chaos **23**, 1330026 (2013).
- [84] A. E. X. Brown, E. I. Yemini, L. J. Grundy, T. Jucikas, and W. R. Schafer, *A dictionary of behavioral motifs reveals clusters of genes affecting Caenorhabditis elegans locomotion*, Proceedings of the National Academy of Sciences of the United States of America **110**, 791 (2013).
- [85] G. J. Berman, D. M. Choi, W. Bialek, and J. W. Shaevitz, *Mapping the stereotyped behaviour of freely moving fruit flies*. J. Royal Soc. Interface **11**, 1 (2014).
- [86] G. Reddy, L. Desban, H. Tanaka, J. Roussel, O. Mirat, and C. Wyart, *A lexical approach for identifying behavioral action sequences*, bioRxiv (2020), 10.1101/2020.08.27.270694.
- [87] S. W. Flavell, N. Pokala, E. Z. Macosko, D. R. Albrecht, J. Larsch, and C. I. Bargmann, *Serotonin and the neuropeptide PDF initiate and extend opposing behavioral states in C. elegans*, Cell **154**, 1023 (2013).
- [88] G. J. Berman, W. Bialek, and J. W. Shaevitz, *Hierarchy and predictability in Drosophila behavior*, Proceedings of the National Academy of Sciences **104**, 20167 (2016).
- [89] S. Gupta and A. Gomez-Marin, *A context-free grammar for Caenorhabditis elegans behavior*, bioRxiv (2019), 10.1101/708891.
- [90] N. Rubido, C. Grebogi, and M. S. Baptista, *Entropy-based generating markov partitions for complex systems*, Chaos: An Interdisciplinary Journal of Nonlinear Science **28**, 033611 (2018).
- [91] J. D. Farmer, *Information dimension and the probabilistic structure of chaos*, Zeitschrift für Naturforschung A **37**, 1304 (1982).
- [92] X. Pei and F. Moss, *Characterization of low-dimensional dynamics in the crayfish caudal photoreceptor*, Nature **379**, 618 (1996).
- [93] C. Brennan and A. Proekt, *A quantitative model of conserved macroscopic dynamics predicts future motor commands*, eLife **8** (2019).
- [94] G. Froyland, S. Lloyd, and N. Santitissadeekorn, *Coherent sets for nonautonomous dynamical systems*, Physica D: Nonlinear Phenomena **239**, 1527 (2010).
- [95] G. Froyland, N. Santitissadeekorn, and A. Monahan, *Transport in time-dependent dynamical systems: Finite-time coherent sets*, Chaos: An Interdisciplinary Journal of Nonlinear Science **20**, 043116 (2010).
- [96] H. Wang and C. Schütte, *Building markov state models for periodically driven non-equilibrium systems*, Journal of Chemical Theory and Computation **11**, 1819 (2015), PMID: 26889513.
- [97] P. Koltai, G. Ciccotti, and C. Schütte, *On metastability and markov state models for non-stationary molecular dynamics*, The Journal of Chemical Physics **145**, 174103 (2016).
- [98] P. Koltai, H. Wu, F. Noé, and C. Schütte, *Optimal data-driven estimation of generalized markov state models for non-equilibrium dynamics*, Computation **6**, 1 (2018).
- [99] R. Friedrich, J. Peinke, M. Sahimi, and M. Reza Rahimi Tabar, *Approaching complexity by stochastic methods: From biological systems to turbulence*, Physics Reports **506**, 87 (2011).
- [100] J. L. Callahan, J.-C. Loiseau, G. Rigas, and S. L. Brunton, *Nonlinear stochastic modeling with langevin regression*, (2020), arXiv:2009.01006 [cond-mat.stat-mech].
- [101] J. A. Fill, *Eigenvalue Bounds on Convergence to Stationarity for Nonreversible Markov Chains, with an Application to the Exclusion Process*, The Annals of Applied Probability **1**, 62 (1991).
- [102] A. Ichiki and M. Ohzeki, *Violation of detailed balance accelerates relaxation*, Phys. Rev. E **88**, 020101 (2013).
- [103] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis, *Diffusion maps, spectral clustering and reaction coordinates of dynamical systems*, Applied and Computational Harmonic Analysis **21**, 113 (2006), special Issue: Diffusion Maps and Wavelets.
- [104] S. Klus, I. Schuster, and K. Muandet, *Eigendecompositions of Transfer Operators in Reproducing Kernel Hilbert Spaces*, Journal of Nonlinear Science **30**, 283 (2020).
- [105] S. Klus, P. Koltai, and C. Schütte, *On the numerical approximation of the Perron-Frobenius and Koopman operator*, Journal of Computational Dynamics **3**,

- 51 (2016).
- [106] W. Bialek, *What do we mean by the dimensionality of behavior?* (2020), arXiv:2008.09574 [q-bio.NC].
- [107] A. J. Calhoun, J. W. Pillow, and M. Murthy, *Unsupervised identification of the internal states that shape natural behavior*, *Nature Neuroscience* **22**, 2040 (2019).
- [108] L. Tao, S. Ozarkar, J. M. Beck, and V. Bhandawat, *Statistical structure of locomotion and its modulation by odors*, *eLife* **8**, e41235 (2019).
- [109] V. Alba, G. J. Berman, W. Bialek, and J. W. Shaevitz, *Exploring a strongly non-markovian animal behavior*, (2020), arXiv:2012.15681 [q-bio.NC].
- [110] H. C. Berg and D. A. Brown, *Chemotaxis in Escherichia coli analysed by Three-dimensional Tracking*, *Nature* **239**, 500 (1972).
- [111] W. M. Roberts, S. B. Augustine, K. J. Lawton, T. H. Lindsay, T. R. Thiele, E. J. Izquierdo, S. Faumont, R. A. Lindsay, M. C. Britton, N. Pokala, *et al.*, *A stochastic neuronal model predicts random search behaviors at multiple spatial scales in C. elegans*, *Elife* **5**, e12572 (2016).
- [112] N. Tinbergen, *On aims and methods of Ethology*, *Zeitschrift für Tierpsychologie* **20**, 410 (1963).
- [113] L. Hebert, T. Ahamed, A. C. Costa, L. O’Shaughnessy, and G. J. Stephens, *WormPose: Image synthesis and convolutional networks for pose estimation in C. elegans*, *PLoS Computational Biology* **17**, e1008914 (2021).
- [114] C. N. Martineau, A. E. X. Brown, and P. Laurent, *Multidimensional phenotyping predicts lifespan and quantifies health in Caenorhabditis elegans*, *PLOS Computational Biology* **16**, 1 (2020).
- [115] Y. Ikeda, P. Jurica, H. Kimura, H. Takagi, Z. R. Struzik, K. Kiyono, Y. Arata, and Y. Sako, *C. elegans episodic swimming is driven by multifractal kinetics*, *Scientific Reports* **10**, 14775 (2020).
- [116] K. N. Le, M. Zhan, Y. Cho, J. Wan, D. S. Patel, and H. Lu, *An automated platform to monitor long-term behavior and healthspan in caenorhabditis elegans under precise environmental control*, *Communications Biology* **3**, 297 (2020).
- [117] A. C. Costa, T. Ahamed, D. Jordan, and G. Stephens, *Maximally predictive ensemble dynamics from data*, <https://doi.org/10.5281/zenodo.4778712> (2021).
- [118] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Scikit-learn: Machine learning in Python*, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [119] H. Kantz and E. Olbrich, *Coarse grained dynamical entropies: Investigation of high-entropic dynamical systems*, *Physica A: Statistical Mechanics and its Applications* **280**, 34 (2000).
- [120] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK Users’ Guide* (Society for Industrial and Applied Mathematics, 1998) <https://epubs.siam.org/doi/pdf/10.1137/1.9780898719628>.
- [121] E. Jones, T. Oliphant, P. Peterson, and et al., *SciPy: Open source scientific tools for Python*, (2001–).
- [122] R. Barrio, A. Dena, and W. Tucker, *A database of rigorous and high-precision periodic orbits of the Lorenz model*, *Computer Physics Communications* **194**, 76 (2015).
- [123] S. M. Hammel, J. A. Yorke, and C. Grebogi, *Do numerical orbits of chaotic dynamical processes represent true orbits?* *Journal of Complexity* **3**, 136 (1987).
- [124] H. E. Nusse and J. A. Yorke, *Is every approximate trajectory of some process near an exact trajectory of a nearby process?* *Communications in Mathematical Physics* **114**, 363 (1988).
- [125] J. E. Sulston and S. Brenner, *The DNA of Caenorhabditis elegans*. *Genetics* **77**, 95 (1974).
- [126] V. Artale, G. Boffetta, A. Celani, M. Cencini, and A. Vulpiani, *Dispersion of passive tracers in closed basins: Beyond the diffusion coefficient*, *Physics of Fluids* **9**, 3162 (1997).
- [127] J. C. Sprott, *Chaos and Time-Series Analysis* (Oxford University Press, Inc., USA, 2003).
- [128] H. S. Kaplan, O. Salazar Thula, N. Khoss, and M. Zimmer, *Nested neuronal dynamics orchestrate a behavioral hierarchy across timescales*, *Neuron* **105**, 562 (2020).
- [129] Y. Iino and K. Yoshida, *Parallel Use of Two Behavioral Mechanisms for Chemotaxis in Caenorhabditis elegans*, *J. Neurosci.* **29**, 5370 (2009).
- [130] M. Peliti, J. S. Chuang, and S. Shaham, *Directional locomotion of C. elegans in the absence of external stimuli*, *PLoS ONE* **8** (2013), 10.1371/journal.pone.0078535.

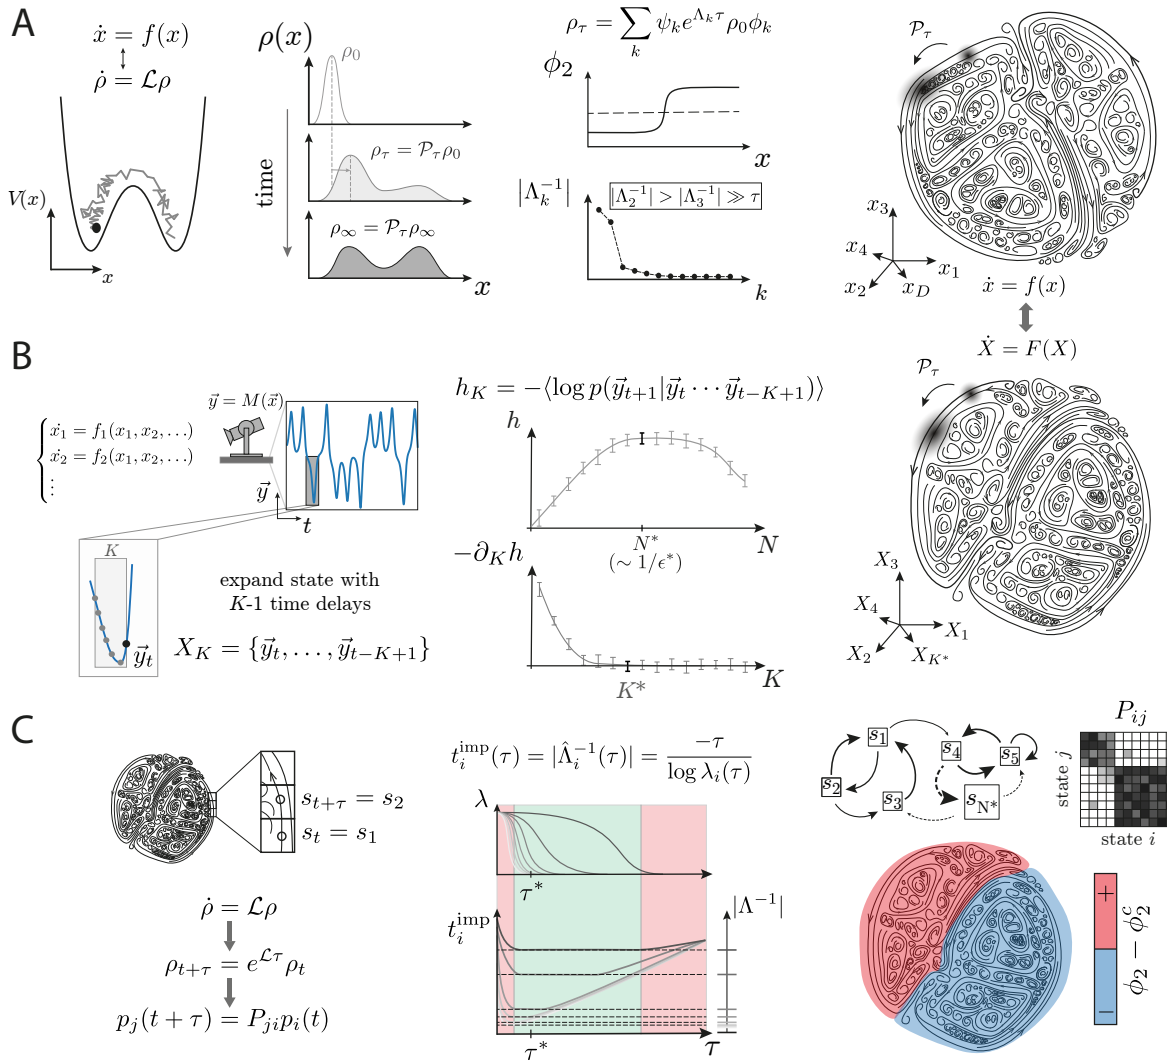


FIG. 1. Maximizing predictive information for state space reconstruction and ensemble dynamics. (A) The ensemble approach naturally identifies large scale structures in systems with a large degree of unpredictability, for which individual short trajectories are less informative (left, particle in a double-well at finite-temperature). We trade non-linear dynamics in the state space $\dot{x} = f(x)$ for equivalent linear dynamics in density space $\dot{\rho} = \mathcal{L}\rho$. In discrete time, the ensemble evolution is dictated by the action of transfer operators \mathcal{P}_τ , evolving densities in time by a time scale τ . We achieve coarse-graining through the eigenfunctions and eigenvalues $\lambda_k = e^{\Lambda_k \tau}$ of such transfer operators $\mathcal{P}_\tau = e^{\mathcal{L}\tau}$; the eigenfunction with $\lambda = 1$ (equivalently $\Lambda = 0$) corresponds to the invariant distribution while the others decay on timescales $|\Lambda_k^{-1}|$. The eigenvalue gap is a principled guide to timescale separation: ϕ_2 splits the invariant density into the two wells and $|\Lambda_2^{-1}|$ captures the hopping timescale. We leverage transfer operators to extract principled, coarse-grained descriptions of complex systems exhibiting intricate state-space structure (right). (B) We reconstruct the state space by concatenating measurements in time to maximize predictability. For each space of K delays we find the number of bins N^* that yield a maximum entropy partition and add delays until $\partial_K h(K^*) \approx 0$, where h is the entropy rate. Our reconstructed states $X = \{y_t, \dots, y_{t-K+1}\}$ preserve the ergodic properties of the system sampled from the measurement time series, including the density dynamics \mathcal{P}_τ . (C) (left) We extract ensemble dynamics through the Perron-Frobenius operator \mathcal{P}_τ approximated as the Markov transition matrix P_{ij} between state space partitions $\{s_i\}$ with transition time τ . P_{ij} provides a predictive model of ensemble dynamics applicable to both deterministic and stochastic systems. (middle) The inferred relaxation times of the underlying dynamics, t_{imp}^i , implicitly depend on the choice of τ . For $\tau \rightarrow 0$, t_{imp}^i exhibit a short transient; P_{ij} is close to the identity matrix and the eigenvalues are nearly degenerate about 1 (left red shaded region). For $\tau \gg \tau^*$ the eigenvalues are approximately constant as P_{ij} is determined primarily by the invariant distribution and relaxation times grow linearly (right red region). We choose τ^* after the initial transient, when the regime in which the long-time dynamics are approximately Markovian settles and the spectral gap is the largest (green shaded region). The eigenvalues of $P_{ij}(\tau^*)$ and the corresponding eigenvectors ϕ reveal coherent structures and the transitions between them (right, red and blue); ϕ_2^c defines the subdivision along ϕ_2 that maximizes the coherence of a metastable set (Methods).

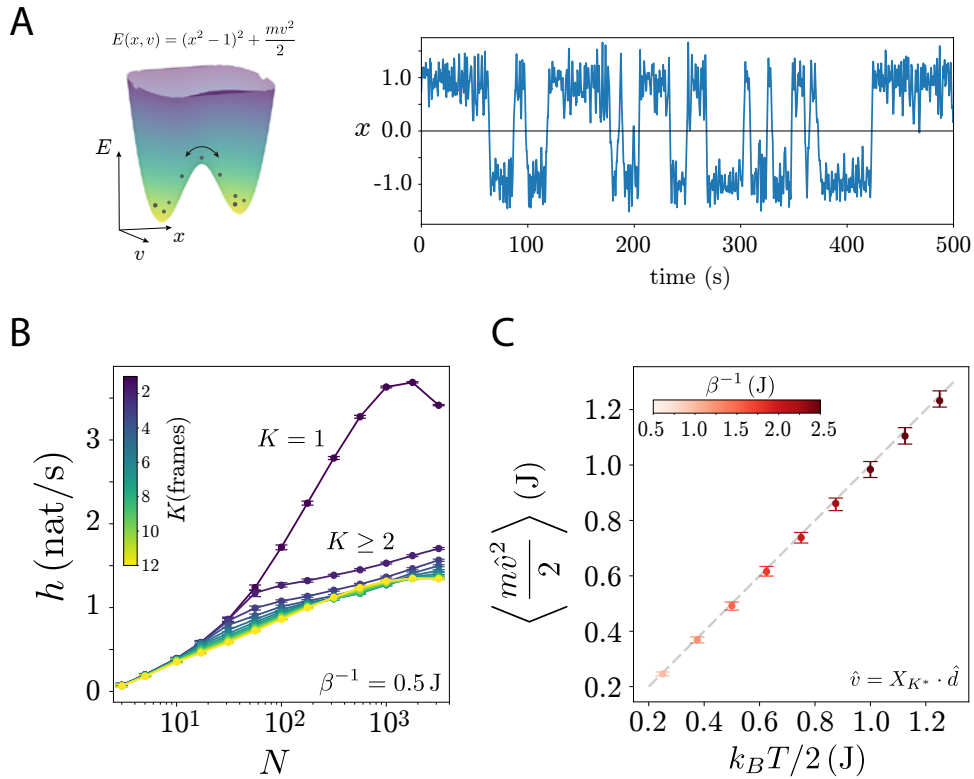


FIG. 2. Maximally predictive state space reconstruction of a Brownian particle in a double-well potential with viscous damping. (A) We use Langevin dynamics Eq. (12) to simulate a long (10^6 s) trajectory of an underdamped particle in a double-well potential (left). To emulate the challenge of data from systems where the fundamental dynamics are unknown, we measure only the position $x(t)$ and show a sample of the trajectory with effective temperature $\beta^{-1} = 0.5$ J (right). (B) We concatenate the measurement with K delays, partition the resulting state with N partitions and show the entropy rate as a function of K and N . The behavior of $h(K, N)$ reveals important information about the measurement time series. The entropy rate changes dramatically as missing and predictive velocity information is incorporated into the state ($K > 2$ frames). Notably, the reconstructed state is not fully Markovian at $K = 2$ frames, a memory effect induced by the Euler-scheme used for simulation [61, 63, 64]. As expected for this continuous stochastic process $\lim_{N \rightarrow \infty} h(K, N) = \infty$. Furthermore, once the additional velocity information is incorporated, we see the approximate logarithmic divergence with N as is expected for a Markov process. Also note the underestimation of the entropy that is expected for large N . Error bars are 95% confidence intervals across 200 non-overlapping trajectory segments of duration 5000 s. We choose $K^* = 7$ frames for subsequent analysis. (C) We recover the results of the equipartition theorem in the reconstructed state space, which illustrates that complete phase space information is contained in the reconstructed state space. We use a 6th-order finite difference approximation to estimate the velocity $\hat{v} = X_{K^*} \cdot \hat{d}$ from the reconstructed state X_{K^*} , where \hat{d} represent the 6th-order finite difference stencil. We then compute the average kinetic energy and show that it closely matches the expected value of $(k_B T)/2$ from equipartition. We use $K^* = 7$ frames across temperatures, as the behavior of the entropy rate with K and N exhibits is similar across the range of temperatures studied here, Fig. S1(A), S5(A). Error bars represent 95% bootstrapped confidence intervals of the mean velocity bootstrapped over 500 s time points.

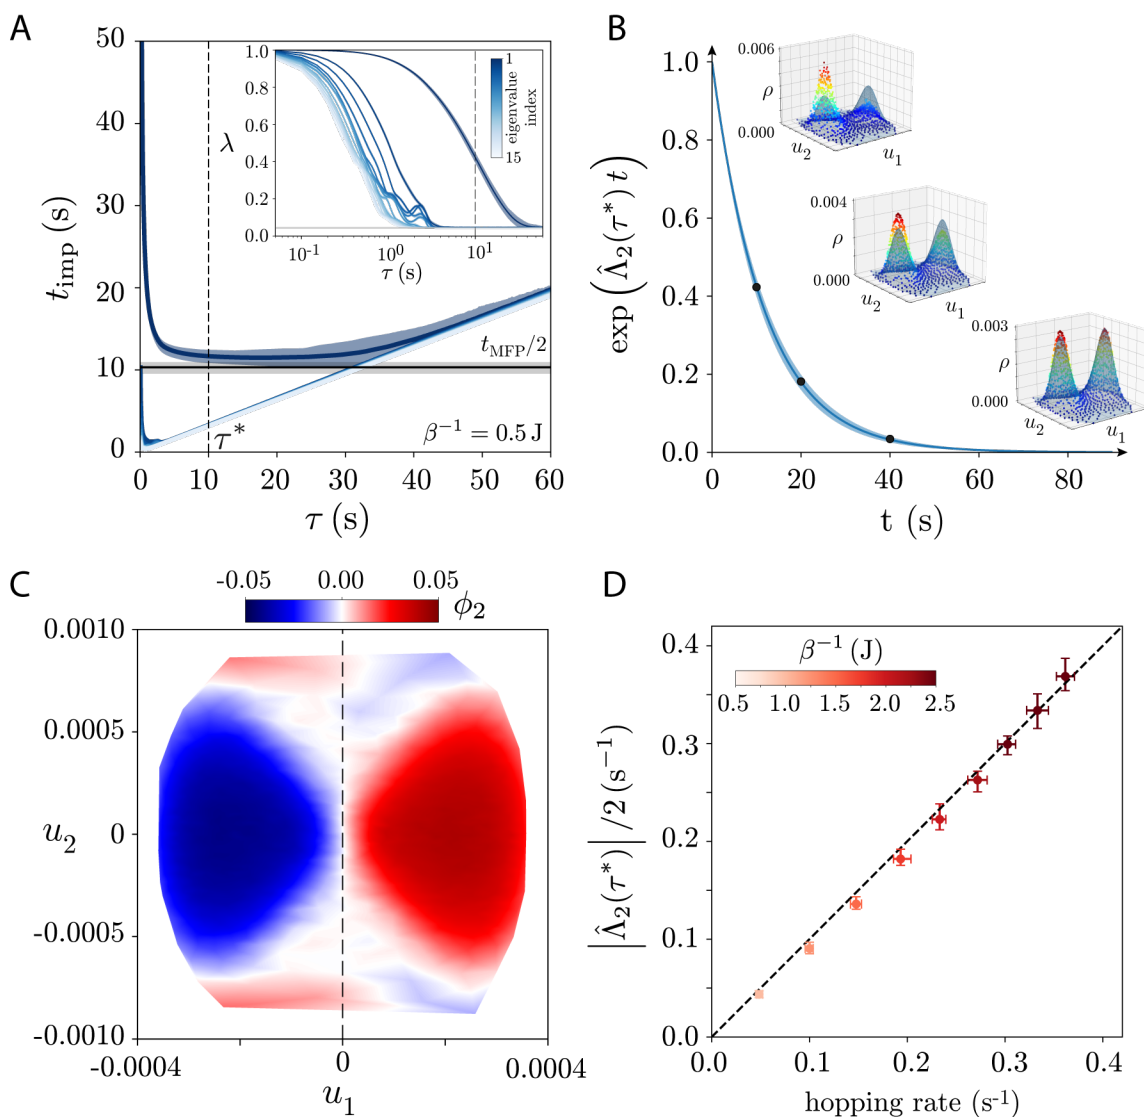


FIG. 3. Ensemble dynamics from the reconstructed state space of a Brownian particle in a double-well potential with viscous damping. We take $x(t)$ from the Langevin dynamics of Eq. 12, reconstruct the state space with $K^* = 7$ delays and $N = 1000$ partitions and construct the finite-time ensemble evolution operator $P(\tau)$ and its reversibilized counterpart $P_r(\tau)$. (A) We plot the eigenvalues $\lambda(\tau)$ of $P_r(\tau)$ (inset) and corresponding implied timescale $t_{\text{imp}}(\tau)$ for $\beta^{-1} = k_B T = 0.5 \text{ J}$. We choose a transition timescale τ^* such that there is a clear time scale separation; Markov dynamics in the first non-trivial eigenfunction ϕ_2 are evident when t_{imp} is constant (dark blue). The short- τ transient results from a nearly diagonal transition matrix reflected in the accumulation of eigenvalues near 1. For large τ the eigenvalues stop changing as the sequence probabilities reach the noise floor (gray horizontal line) and become independent: in this regime t_{imp} grows linearly. The longest relaxation time is approximately constant after $\tau \sim 5 \text{ s}$, converging to half the mean first passage time $t_{\text{MFP}}/2$ (black horizontal line) as expected from theory. We choose $\tau^* = 10 \text{ s}$ for subsequent analysis. (B) Ensemble dynamics captured by ϕ_2 projected onto the largest two singular vectors (u_1, u_2) of the reconstructed state space, which match the position and velocity degrees of freedom, Fig. S1(B). The decay captures the slowest evolution towards the equilibrium density. We start from an initial ensemble sharply concentrated on a single partition on the bottom of the left well, and propagate the density using P with $\beta^{-1} = 0.5 \text{ J}$. (C) Contour plot of ϕ_2 projected onto the (u_1, u_2) space for $\beta^{-1} = 0.5 \text{ J}$. The sign of ϕ_2 effectively splits the reconstructed state space into the two wells: $\phi_2 < 0$ corresponds to the right well, while $\phi_2 > 0$ corresponds to the left well. Note also that the high velocity regions $|u_2|$ correspond to transition regions. (D) $|\hat{\Lambda}_2(\tau^*)|/2$ predicts the hopping rate across a range of temperatures. We choose τ^* separately for each temperature to reflect the different dynamics. The behavior of the longest implied time scale is qualitatively similar across temperatures, reaching a minimum at the plateau region, Fig. S2(A). We thus choose τ^* as the time scale that minimizes the longest t_{imp} . Error bars represents 95% confidence intervals over 200 non-overlapping trajectory segments of length 5000 s.

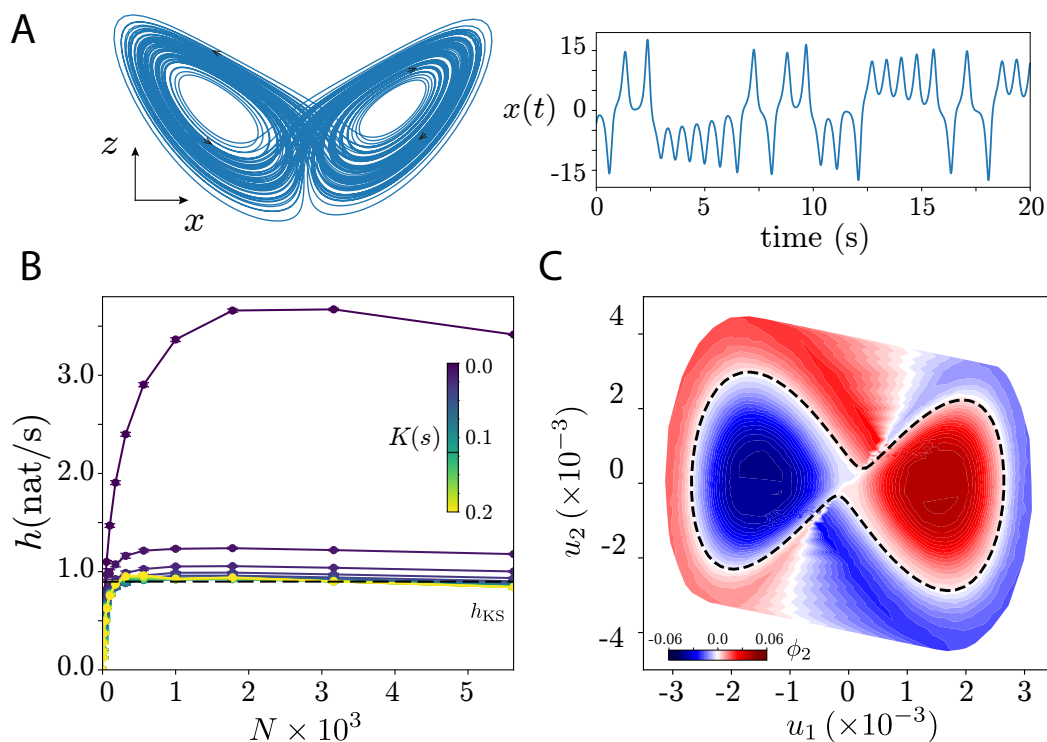


FIG. 4. Maximally predictive ensemble dynamics in the Lorenz system. (A) We simulate the Lorenz equations, Eq. (14), in the standard chaotic regime and extract measurements of $x(t)$, sampled at $\delta t = 0.01$ s for 5×10^5 s. An example 100 s $\{x, y\}$ trajectory is shown at left, and a sample 20 s measurement of x is shown at right. (B) The entropy rate h as a function of the number of partitions N for increasing number of delays K (color coded). From $K \sim 10\delta t = 0.1$ s the entropy rate is approximately constant and we use $K^* = 0.12$ s for subsequent analysis. The estimated entropy rates converges to the maximum Lyapunov exponent $h_{KS} = \sum_{\lambda > 0} \lambda$: dashed line is the reference value of the largest exponent $h = 0.90563$ (nat/s) [127]. (C) Contour plot of the first non-trivial transfer eigenvector of the inferred reversibilized transfer operator, ϕ_2 , projected along the largest two SVD modes of the reconstructed state space (u_1, u_2) , computed with transition time $\tau^* = 0.1$ s and $N = 10^{3.5}$ partitions. The sign of ϕ_2 divides the state space into its almost invariant sets, which are partially split along the shortest unstable periodic orbit [67], here identified directly from the data using recurrences (see Methods). Error estimates on the slowest relaxation time are 95% confidence intervals on the estimates obtained from 250 non-overlapping trajectory segments of duration 2000 s.

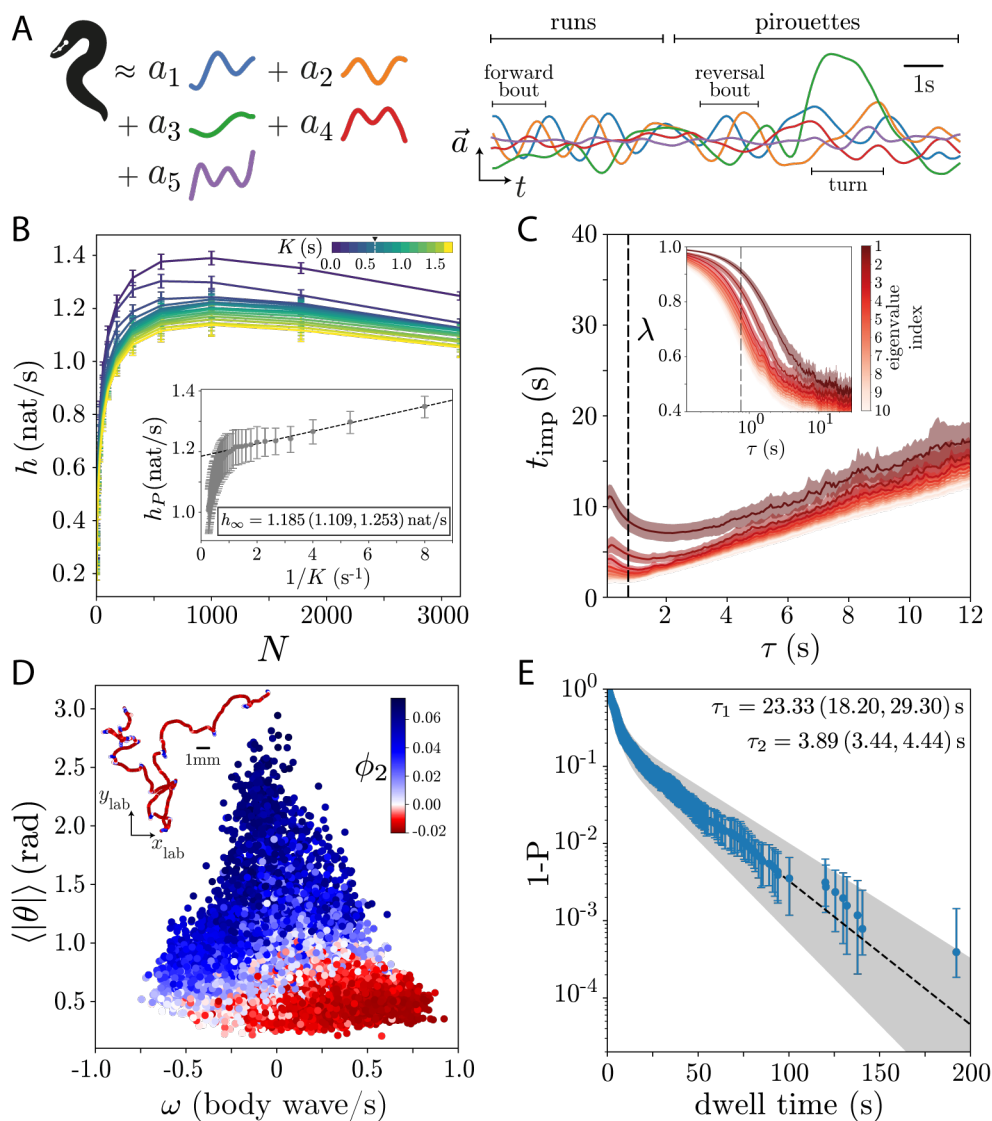


FIG. 5. Deriving “run-and-pirouette” navigation from posture dynamics in *C. elegans*. (A-left) We represent the centerline posture dynamics of the worm’s body by a 5D (eigenworm) time series [74]. In food-free conditions, the short time scale behavior is roughly composed of forward bouts, reversals and sharp turns. On longer time scales, worms (like other animals) combine basic behaviors into longer term strategies such as navigation, here exemplified by relatively straight “runs” interrupted by abrupt “pirouettes” [77]. (B) Entropy rate as a function of the number of delays K and partitions N . The curves collapse after $K \sim 0.5$ s, indicating that the entropy rate is approximately constant. We choose $K^* = 11$ frames = 0.6875 s for subsequent analysis. For $N \gtrsim 1000$ partitions finite-size effects are apparent, and we extract the $K \rightarrow \infty$ limit of the entropy rate by extrapolation $h_\infty = 1.19 \pm (1.11, 1.25)$ nat/s (inset). Error bars are 95% confidence intervals of the mean, bootstrapped across worms. (C) The largest 10 implied time scales as a function of transition time τ . We choose $\tau^* = 0.75$ s (dashed line) for subsequent analysis, $|\hat{\Lambda}_2^{-1}|(\tau^*) = 8.42$ (7.40, 9.52) s. Error bars are 95% confidence intervals bootstrapped across worms. (inset) Eigenvalues of the transfer operator as a function of τ . As with previous examples, at very small τ or very large τ the eigenvalues collapse. We observe the largest separation for intermediate τ . (D) We use projections along ϕ_2 to color the body wave phase velocities ω and mean body curvature $\langle |\theta| \rangle$, averaged in $2\tau = 1.5$ s windows. Positive values align with negative phase velocities and large turning amplitudes, indicative of “pirouettes”, while negative values correspond to positive phase velocities and low turning amplitudes, indicative of “forward runs”. (inset) An example 10 minute long centroid trajectory color coded by ϕ_2 . Negative projections occur during “runs”, while positive values are found when worms abruptly change the locomotion direction. (E) We cluster ϕ_2 into two coarse-grained states and find that the cumulative distribution of resulting run lengths $1 - P(t_{\text{state}} < t)$ is characterized by two time scales, which we extract by fitting a sum of exponential functions, and in agreement with previous phenomenological observations [77]. The timescale errors and error bars are 95% confidence intervals bootstrapped across events. The shaded area represents the 95% confidence intervals bootstrapped across worms, and the deviation at longer dwell-times is discussed in the context of Fig. 7.

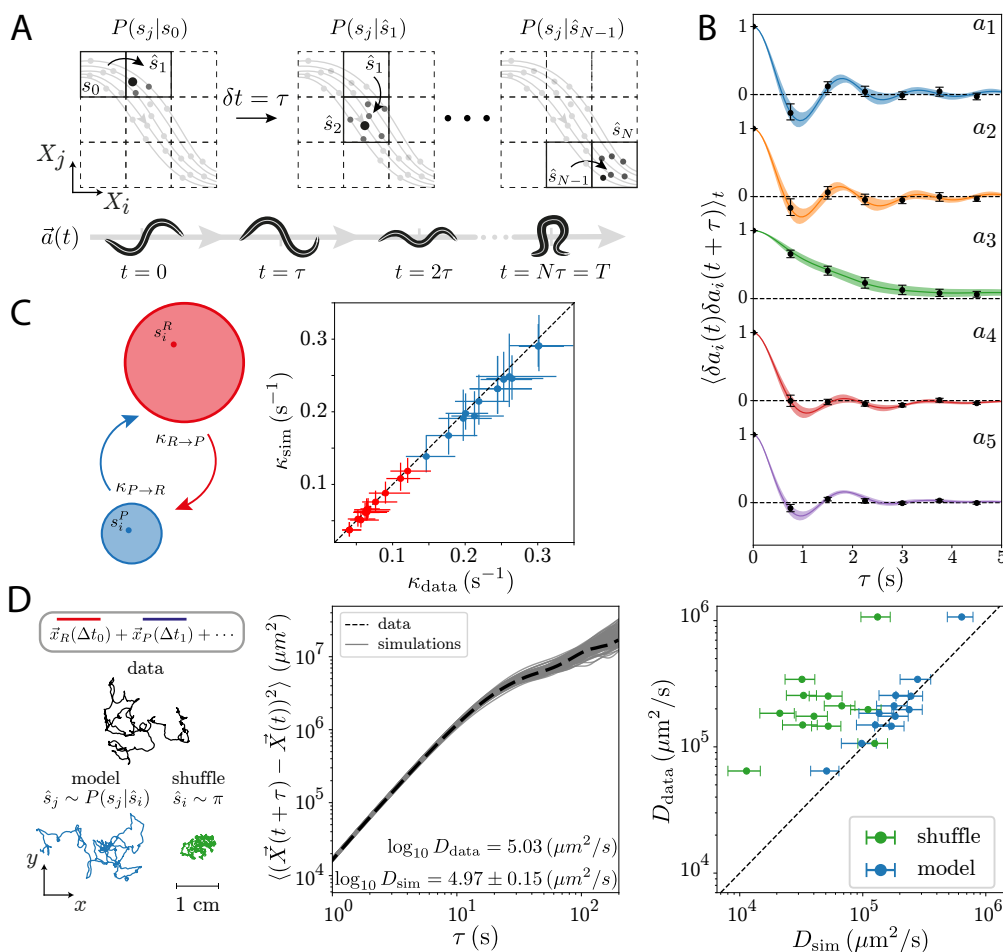


FIG. 6. Transfer operator derived kinetics predict fine scale posture dynamics as well as large scale diffusive properties of foraging trajectories. (A) Schematic of the simulation method. Starting from the initial microstate of each worm s_0 , the next microstate is obtained by sampling from $P(s_j|s_0)$ constructed with $\tau = \tau^* = 0.75$ s. Each new microstate is added in the same fashion resulting in a microstate symbolic sequence of length $N = T\delta t/\tau^*$. Within each microstate, we randomly choose a point in the associated reconstructed state space, X , and use this point to define the body posture \vec{a} . Through this effective diffusion in state space we obtain time series of postures through a sequence of “eigenworm” coefficients $\vec{a}(t)$ sampled at $\delta t = \tau^*$. (B) Operator derived posture dynamics accurately predicts the autocorrelation function of the “eigenworm” coefficient time series. Colors denote the different modes, while black dots correspond to the simulations. Error bars and shaded regions correspond to 95% confidence intervals on the estimate of the mean autocorrelation function bootstrapped across worms. (C) We coarse-grain the symbolic sequence into run and pirouette states, as in Fig. 5, and show the mean transition rates from the run to the pirouette state $\kappa_{R \rightarrow P}$ (red) and from the pirouette to the run state $\kappa_{P \rightarrow R}$ (blue) obtained for each of the 12 recorded worms by inverting the mean dwell times in each state. The operator derived simulations exhibit kinetic rates κ_{sim} that agree with the kinetic rates obtained from the data κ_{data} . Data error bars are 95% confidence intervals of the mean bootstrapped across run and pirouette segments, simulation error bars are 95% confidence intervals of the mean bootstrapped across 1000 simulations. (D-left) We construct simulated centroid trajectory by iteratively drawing from the space of actual trajectories, choosing a centroid run trajectory X_R (red) or a centroid pirouette trajectory X_P (blue) with duration Δt which matches the dwell time of the microstate operator dynamics. We then append each new centroid segment to the end of the previous segment and align the direction of the worm’s movement across the boundary. We show the trajectory of an example worm, as well as simulated trajectories generated from the operator dynamics $\hat{s}_j \sim P(s_j|\hat{s}_i)$ (blue), and from a shuffle which obeys the same steady-state distribution $\hat{s}_i \sim \pi$ (green). (D-middle) Mean square displacement of centroid trajectories obtained from 1000 simulations (gray) and the data (black dashed line) for an example worm undergoing a ballistic-to-diffusive transition. Errors of the diffusion coefficient are standard deviations across simulations. (D-right) Estimated effective diffusion coefficients obtained from simulations D_{sim} and the data D_{data} . We estimate D by fitting the slope of the mean square displacement curves in the interval $\tau \in [60, 100]$ s and find operator-predicted values (blue) closely match those obtained from the data. Symbolic sequences that preserve the steady-state distribution but are otherwise random (green) fail to capture the diffusive properties of the centroid trajectories. Errors are standard deviations of the estimated diffusion coefficients across simulations.

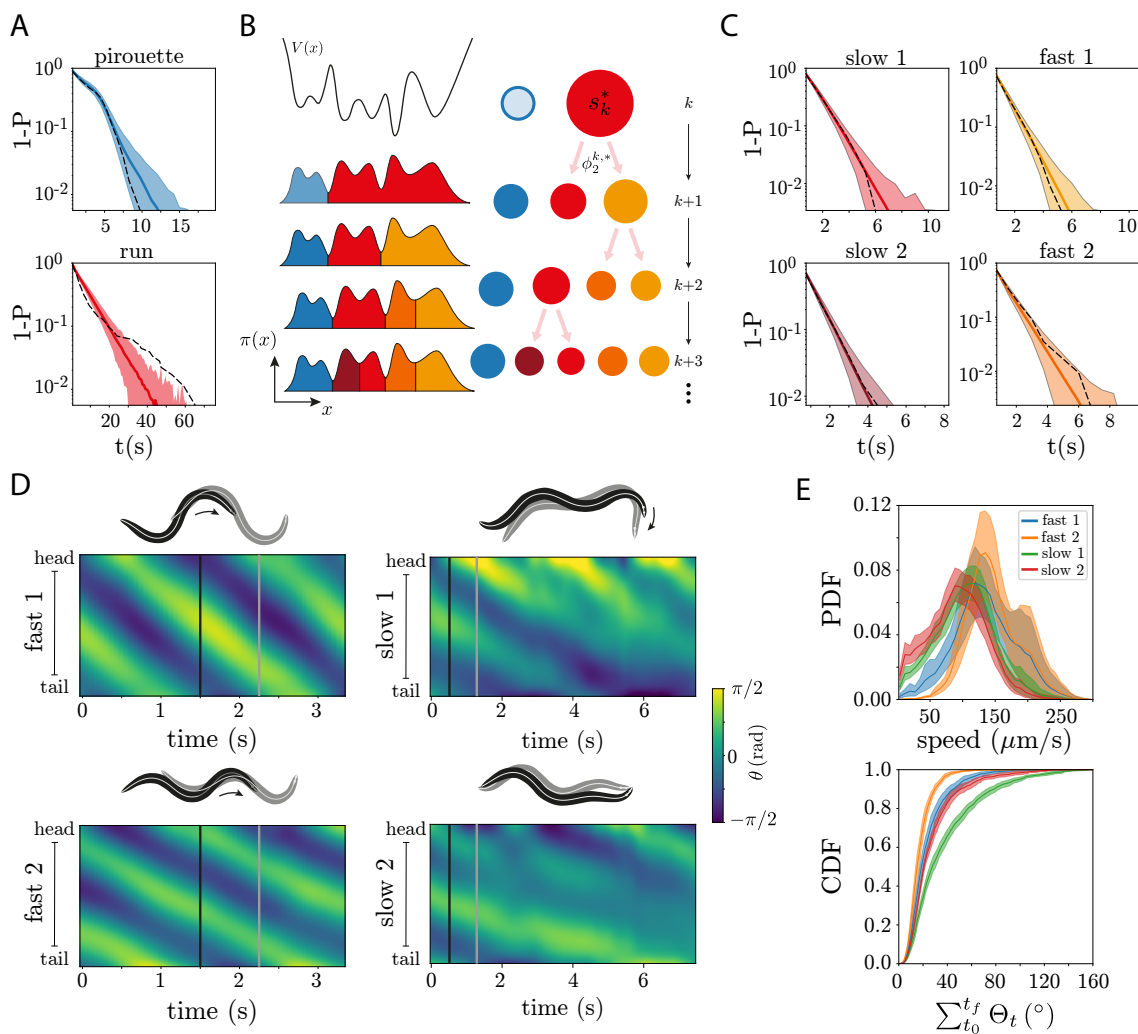


FIG. 7. An operator-based subdivision of “run-and-pirouette” states reveals fine-scale behaviors underlying the dwell time distributions. (A) Dwell time distributions of the first two metastable states for simulations of the operator dynamics (blue - pirouette, red - run) and the data (black dashed line) for an example worm. The run state exhibits a heavy tail not accurately captured in our simulations. Shaded errors are 95% confidence intervals across 1000 simulations. (B) Operator-based state space subdivision. At each iteration k , we subdivide the highest measure metastable state, s_k^* : we ignore the rest of the state space (shaded blue region) and partition along the first non-trivial eigenfunction obtained from the transition probability matrix conditioned only on the states within this metastable set, $\phi_2^{k,*}$ [83] (Methods). In principle, subdividing can be repeated as long as transitions between microstates are observed. (C) Dwell time distributions in the 4-state subdivision of the run state, for an example worm. Subdividing the run state into slower and faster forward movements better captures the dwell time distributions, see also Fig. S8(A). Shaded errors are 95% confidence intervals across 1000 simulations. (D-left) The two fast states correspond to distinct gaits. We show example body curvature θ kymographs with example postures taken $\tau = 0.75$ s apart. (D-right) Example kymographs for the two slow states: “slow 1” exhibits head-casts [128], while “slow 2” has only minor movements, effectively a “pause”. (E) - Probability Distribution Function (PDF) of the centroid speeds for the four run states. Differences between fast states (blue and orange) and between slow states (green and red) are subtle from a centroid perspective, with the speed distributions nearly overlapping. (E-bottom) - Cumulative distribution function (CDF) of the summed angle between subsequent centroid velocity vectors measuring the overall curvature of the path. The higher wavelength observed in “fast 1” states results in a gradual “weathervaning” [129, 130] of the trajectory when compared to the trajectories of the “fast 2” state which are comparably straighter. Similarly, the “head-casts” in the “slow 1” state result in the gradual slow reorientation of the worm, when compared to the “slow 2” state for which trajectories are less curved. Shaded areas correspond to 95% confidence intervals bootstrapped across worms. To avoid spurious large angles due to reversals, we discard the 10% largest magnitude angles in all states.

SUPPLEMENTARY MATERIAL




FIG. S1. **Details of the state space reconstruction of an underdamped particle in a double well potential.** (A) Entropy rate as a function of delay K and number of partitions N across temperatures. The qualitative behavior of the entropy rate is notably conserved across temperatures. (B) Singular value decomposition of the reconstructed state space X_K^* with $K^* = 7$ frames. The first two SVD modes capture nearly 100% of the variance (top). In addition, the resulting singular vectors capture the position and the velocity degrees of freedom (bottom): u_1 is essentially a weighted sum of the position, while u_2 is a derivative filter of the position, yielding an estimate of the local velocity.

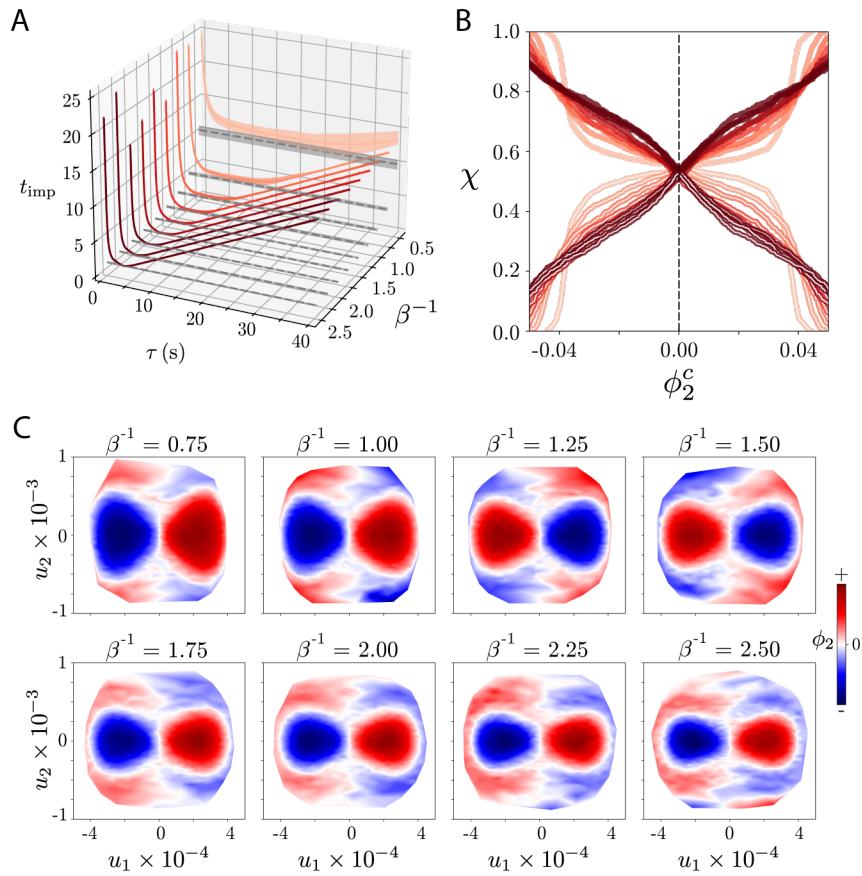


FIG. S2. Details of the identification of metastable states in the Double Well potential across temperatures.

(A) Inferred longest implied time scales t_{imp} as a function of transition time τ and effective temperature $\beta^{-1} = k_B T$. The qualitative behavior of t_{imp} with τ is conserved across temperatures: early transient decay to the time scale of the hopping rate $t_{\text{MFP}}/2$ (gray dashed line) and long time linear increase with τ indicating independence of the states. Increasing the temperature has the expected effect of reducing the duration of the plateau: higher temperatures result in faster mixing, reducing the time it takes before states become independent and t_{imp} grows linearly. Accordingly, we observe a decrease in t_{imp} with temperature. We leverage the conserved qualitative behavior of the t_{imp} curves across temperatures to choose a consistent τ^* as the time scale that minimizes t_{imp} . This results in a lower τ^* for higher temperatures, which reflects the faster nature of the dynamics: the fast modes relax faster, enabling for a larger time scale separation at shorter τ . (B) Coherence measure χ , as defined in Eq. (11) as a function of the inferred first non-trivial eigenfunction of \mathcal{P}_τ , ϕ_2 , for different temperatures. For each temperature, we show both $\chi_{\mu,\tau}(S^+)$ and $\chi_{\mu,\tau}(S^-)$ as well as the minimum between them, χ , in a gray dashed line. Changing τ^* to reflect the nature of the dynamics yields a relatively constant maximum χ across temperatures. Notably, at τ^* the coherence measure indicates that nearly half the measure has escaped the metastable states. (C) ϕ_2 projected onto the first two SVD modes of the reconstructed state X_K across temperatures: the slow eigenfunction mainly splits the state space along the position coordinate and the additional high velocity $|u_2|$ transition regions.

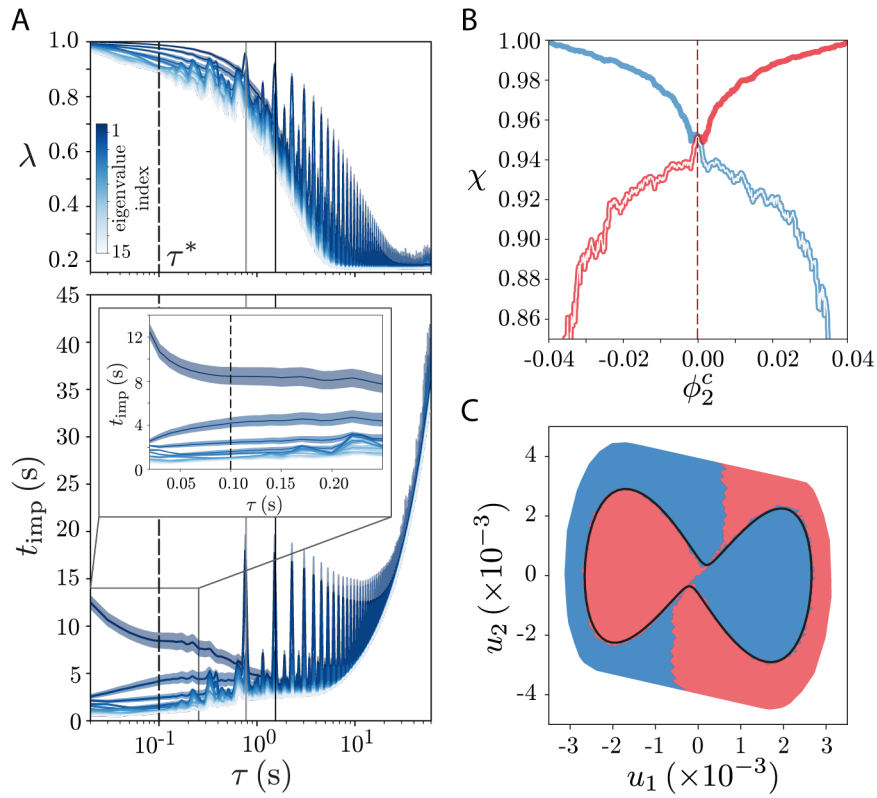


FIG. S3. Details of the transfer operator analysis of the Lorenz system (A) Behavior of the implied time scales t_{imp} with transition time τ for the Lorenz system. (A-top) 15 largest magnitude non-trivial eigenvalues as a function of τ . For very small τ the eigenvalues are nearly degenerate around 1. After this initial transient, the eigenvalues spread, but unlike the Double Well example the time scale separation is much less evident since we are in presence of deterministic chaos. For $\tau \sim 0.75$ s, we observe an increase in the inferred eigenvalues resulting from the periodicity of the Lorenz attractor: the black solid line represents the period of the shortest identified unstable periodic orbit (UPO) $T_{\text{UPO}} = 1.55$ s [122] (see Methods), while the gray solid line represents half of this period. The period of the UPO sets the onset of the increase in the estimated eigenvalues: the periodicity induces an aliasing effect; the transfer operator becomes again nearly identity and we observe the same behavior at integer multiples of $T_{\text{UPO}}/2$. For very large τ , the transfer operator is composed of near copies of the stationary distribution, resulting in the collapse of the eigenvalues, as observed for the Double Well potential, Fig. 3(A-inset). (A-bottom) Implied relaxation times as a function of τ . As expected from the behavior of the eigenvalues (top), there is an initial short transient before the longest relaxation time settles into a near constant value (inset). For larger τ , the effects of the periodicity of the Lorenz attractor come into play: the time scale separation shortens and we observe peaks and integer multiples of $T_{\text{UPO}}/2$. Despite the cyclic pattern, for large τ we observe a linear growth of the implied relaxation time scales as in the Double Well example (here seen as exponential growth on a logarithmic scale), Fig. 3(A). We choose $\tau^* = 0.1$ s yielding an upper bound on the mixing time of $|\hat{\Lambda}_2^{-1}|(\tau^*) = 8.45$ (7.70, 9.26) s. (B) Degree of almost invariance as measure by χ , Eq.(11). Colors represent the contribution of each almost invariant set ($\chi_{\mu,\tau}(S^+)$ and $\chi_{\mu,\tau}(S^-)$) while the white line is the minimum between sets, χ . We split the state space along the global maximum of $\chi \sim 0$. (C) Resulting almost invariant sets, projected onto the first two singular vectors of the trajectory matrix, X_K and plot of the shortest UPO identified through recurrences (see Methods). The boundary between almost invariant sets is partially defined by the stable manifold of the UPO [67], of which we show only the orbit itself for clarity. Error bars represent 95% confidence intervals from 250 trajectories of duration 2000 s.

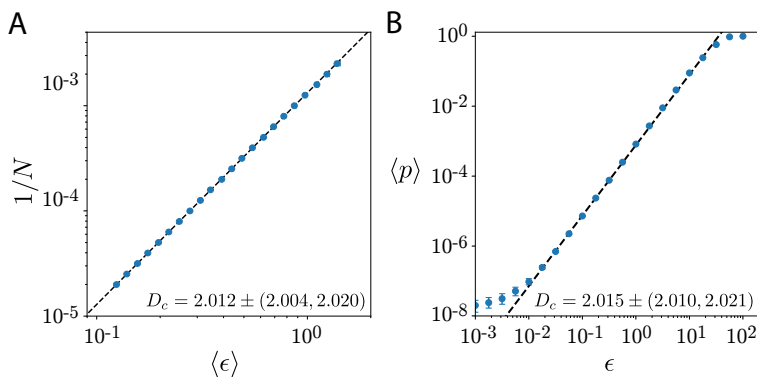


FIG. S4. **Estimation of the information dimension in a partitioned state space.** The information dimension can be obtained by studying the scaling of the probability distribution with length scale. The average probabilities $\langle p \rangle$ of finding a neighbor within a length scale ϵ scales as $\langle p \rangle = \epsilon^{D_c}$, in which D_c is the correlation dimension (or pointwise dimension [43]), which generally bounds the information dimension from above $D_c \geq D_I$. For a typical chaotic attractor, and in the limit $\epsilon \rightarrow 0$, the correlation dimension is indistinguishable from the information dimension, and thus $\langle p \rangle \sim \epsilon^{D_I}$. Estimates of information dimension are typically done using Chebyshev distances: $d(X, Y) = \max_i^D |X_i - Y_i|$ for $X, Y \in \mathbb{R}^D$ and $|\cdot|$ the absolute value. (A) Information dimension estimates in a partitioned state space. The probability of belonging to a partition scales with the inverse of the number of partitions, and thus $p \propto 1/N \propto \epsilon^{D_I}$. Therefore, we can directly estimate the information dimension by measuring how the within cluster distances scale with increasing N . We estimate D_I by linear regression of $1/N \propto D_I \log(\epsilon)$, where $\langle \cdot \rangle$ represents the median. Error estimates for D_I are obtained by bootstrapping the median within-cluster distances. Error bars in $\langle \epsilon \rangle$ are generally smaller than the marker size. (B) Geometric estimation of the correlation dimension using the Cohen-Proccaccia (CP) method [40] in our Lorenz system simulation. The CP method assesses the scaling of the probability distribution with length scale, by measuring the correlation function of a collection of state space points. We sample 5000 points from the attractor, and count the number of neighbors within a distance ϵ from the sampled points using a Chebyshev distance. D_I is then estimated by linear regression $\log \langle p \rangle \propto D_I \log \epsilon$. Error bars represent 95% confidence intervals of the estimate done 100 times with different 5000 point samples.

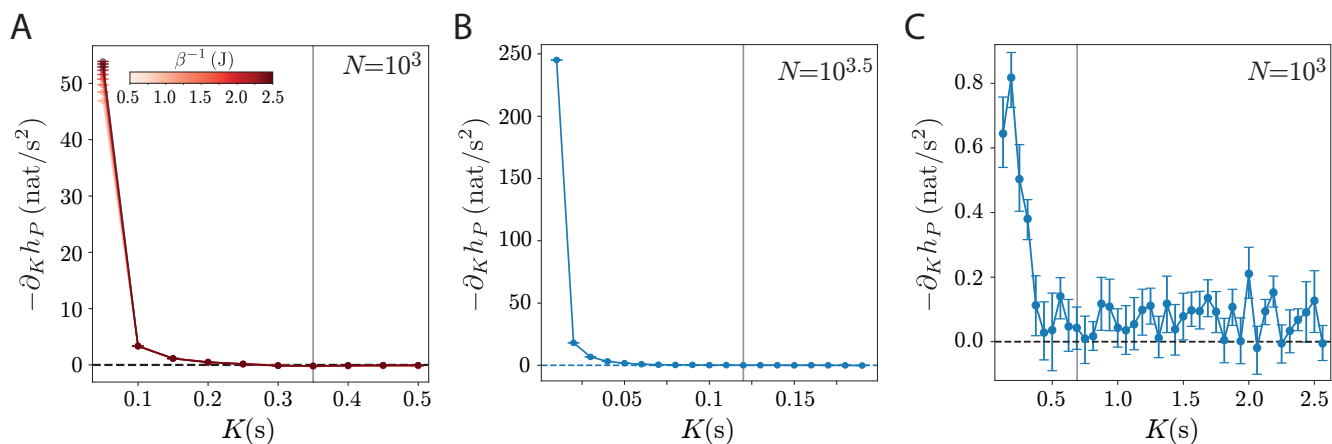


FIG. S5. **Change in entropy rates as a function of delays K for the double well example, the Lorenz system and *C. elegans* posture data.** (A) Double well dynamics for $N = 10^3$ partitions. The entropy rate reaches a plateau after $K \gtrsim 6$ frames = 0.3s and we choose $K^* = 7$ frames = 0.35s. Notably, the change in entropy rate is still significantly nonzero at $K = 2$ frames, reflecting the non-Markovian effects of the Euler integration scheme as discussed in the main text. Error bars represent 95% confidence intervals bootstrapped over 5000s trajectory segments. (B) Lorenz system for $N = 10^{3.5}$ partitions. The entropy rates reaches a plateau after $K \gtrsim 0.05$ s and we choose $K^* = 0.12$ s. Error bars represent 95% confidence intervals bootstrapped over 2000s trajectory segments. (C) *C. elegans* posture dynamics for $N = 10^3$ partitions. The entropy rates reaches a plateau after $K \gtrsim 0.5$ s and we choose $K^* = 0.6875$ s. Error bars represent 95% confidence intervals bootstrapped across worms.

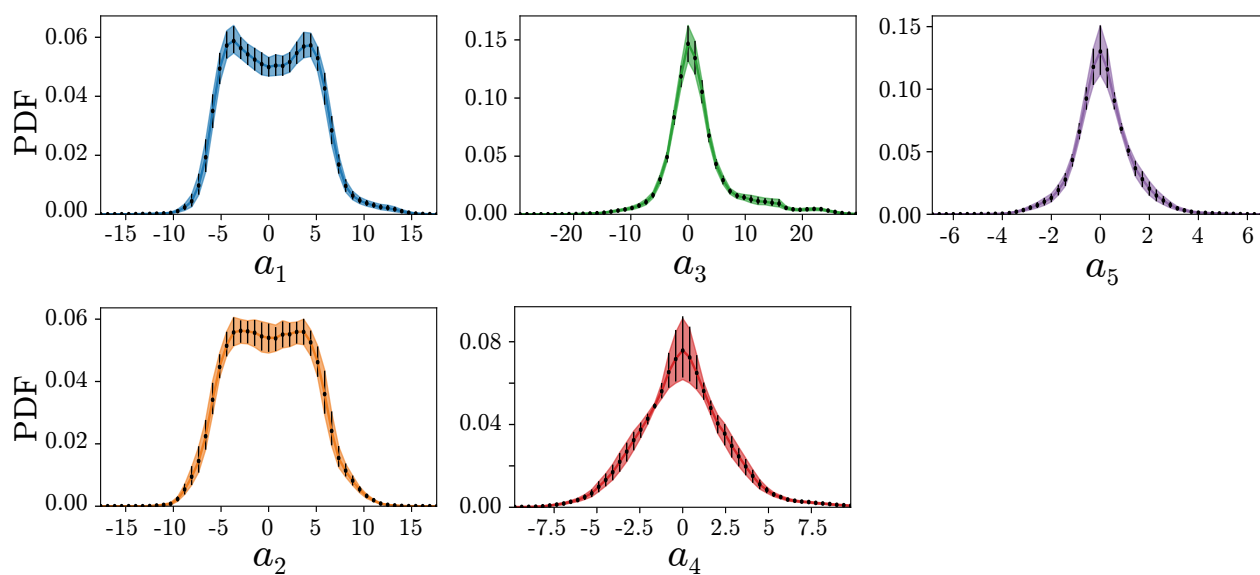


FIG. S6. Probability Density Function (PDF) of the “eigenworm” coefficients show a tight agreement between the data (colors) and simulations (black error bars). Error bars correspond to 95% confidence intervals bootstrapped across worms.

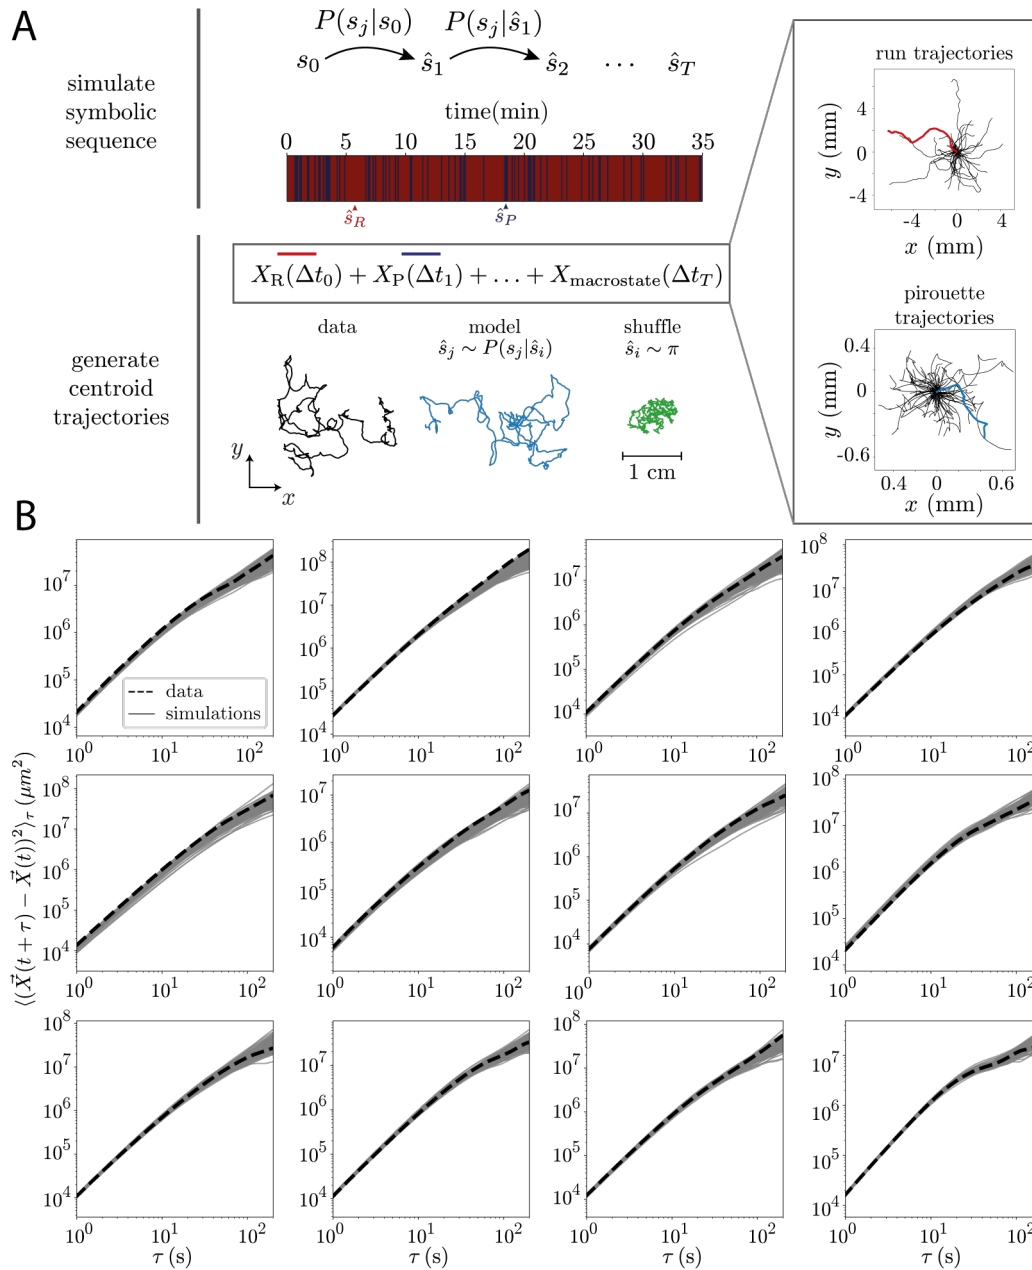


FIG. S7. Details of the centroid trajectory simulations. (A) Illustration of the simulation method. (top) Starting from the initial condition s_0 , a partition on the reconstructed state space, the next state is sampled using the conditional probability $P(s_j|s_0)$. We iterate until we get a symbolic microstate sequence of the same length as the worm recordings. (middle) We associate each microstate with a “run” (red) or “pirouette” (blue) macrostate in which the microstate is contained, resulting in a binary sequence of macrostates with duration Δt_i . (right) We construct the simulated centroid trajectory by iteratively drawing from the space of actual trajectories, choosing a centroid run trajectory X_R (red) or a centroid pirouette trajectory X_P (blue) with duration which matches the dwell time resulting from the microstate operator dynamics. (bottom) We append each new centroid segment to the end of the previous segment and align the direction of the worm’s movement across the boundary. We thus retain the trajectory characteristics within each macrostate but destroy any correlations across them. We show the trajectory of an example worm, as well as simulated trajectories generated from the operator dynamics $\hat{s}_j \sim P(s_j|\hat{s}_i)$ (blue), and from a shuffle which obeys the same steady-state distribution $\hat{s}_i \sim \pi$ (green). (B) We show the mean square displacements for the data (dashed line) of each worm, as well as for 100 centroid trajectory simulations generated from symbolic sequences simulated with the Markov model (gray). By fitting a linear function in the interval $\tau \in [60, 100]$ s we obtained the effective diffusivity estimate of Fig. 6(C-right).

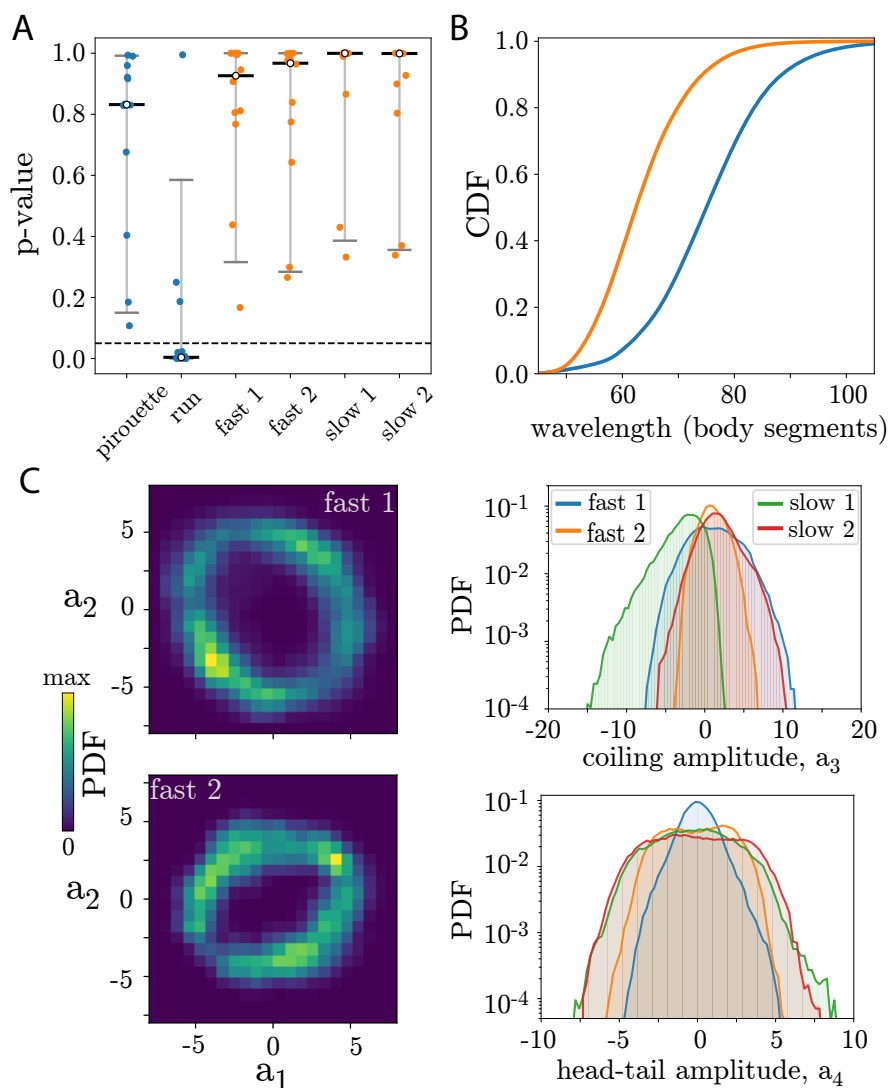


FIG. S8. **Details of the operator-guided subdivision of *C. elegans* postural dynamics.** (A) Scatter plots of the p -values of a Kolmogorov-Smirnov test comparing the dwell time distributions obtained from simulations and the data for a 2-state (blue) and a 5-state (orange) coarse-graining, across worms. In the two state partition, a large fraction of the dwell time distributions in the run state are not accurately captured by the simulations ($p_{\text{value}} < 0.05$, dashed line), whereas subdividing the run state yields dwell time distributions that are statistically indistinguishable from the data. Vertical bars represent the median, 5-th and 95-th percentiles and we add horizontal Gaussian jitter to the scatter plots to aid visualization. (B) Cumulative Distribution Function (CDF) of the body wavelength in the two fast states, obtained as the peak in the spatial power spectrum (obtained using Welch’s method). The first fast state exhibits larger wavelength, spanning almost the entire body, whereas the second slow state exhibits shorter wavelengths. (C-left) The difference in wavelength in the two fast states is translated into different cyclic patterns in the (a_1, a_2) plane, which captures the worms’ body wave. Such subtle differences likely become more evident through the state space reconstruction. (C-right) Probability distribution of the coiling amplitude (a_3 , top) and head-tail amplitude (a_4 , bottom), across run states. The state “fast 1” exhibits a higher coiling amplitude and a lower head-tail amplitude, while “fast 2”, which displays straighter centroid paths, exhibits a lower coiling amplitude and higher modulation of the head-tail amplitude. The two slow states exhibit distinct biases in coiling amplitude: “slow 1” has a clear dorsal bias ($a_3 < 0$) while “slow 2” has a slight ventral bias ($a_3 > 0$). Both also exhibit more activity in the head-tail mode compared to the fast states.




FIG. S9. Markovianity of the irreversible Markov chain as measured through the inferred relaxation times. The reversibilized transition matrix provides an optimal partition into almost invariant sets, but the resulting kinetics does not necessarily capture the underlying dynamics. To directly probe the Markovianity of the underlying slow dynamics, we have to estimate the relaxation times for the non-reversibilized transition matrix, which should not change with τ when the dynamics is Markovian. However P generally presents complex eigenvalues and so the relaxation times are not directly attainable through spectral analysis. To solve this problem in systems which display a time-scale separation, we approximate the slow relaxation dynamics by using the metastable states to build a two-state, coarse-grained Markov chain P_c , which necessarily has only real eigenvalues $\lambda_c \in \mathbb{R}$ (details in Methods). The corresponding relaxation time is then obtained through $|\Lambda_c^{-1}| = -\tau / \log \lambda_c(\tau)$. In general, we find that the regime in which $|\hat{\Lambda}_2^{-1}|$ from P_r is constant overlaps with regime in which $|\Lambda_c^{-1}|$ is also constant, although generally $|\Lambda_c^{-1}|$ exhibits a larger Markovian regime than $|\hat{\Lambda}_2^{-1}|$. (A) Double Well example across temperatures, β^{-1} . The time scales of P converge to the hopping time scale even at $\tau = \delta t$, unlike $|\hat{\Lambda}_2^{-1}|$ from P_r . The underdamped nature of the dynamics creates microscopic fluxes that are erased by the reversibilization of Eq. (10), resulting in the overestimation of the relaxation times measured through P_r as the system lingers for longer in each metastable state. At τ^* however, the dynamics becomes effectively overdamped, and the time scales of P and P_r coincide as the system is macroscopically reversible. Dashed horizontal line represents half the mean first passage time estimated from the data as in Fig. 3(A). Error bars represents 95% confidence intervals over 200 non-overlapping trajectory segments of length 5000 s. (B) Lorenz system. The time scales from P_c are shorter than the reversibilized time scales of Fig. S3(A) across τ , reflecting the dissipative irreversible nature of the dynamics. Nonetheless, choosing $\tau^* \gtrsim 0.1$ s for robust almost invariant sets also yields nearly Markovian slow kinetics. In addition, the periodic structure of the Lorenz system is reflected in $|\Lambda_c^{-1}|$, as it was for $|\hat{\Lambda}_2^{-1}|$, Fig. S3(A). Error bars represent 95% confidence intervals over 200 non-overlapping trajectories segments of duration 2000 s. (C) *C. elegans* foraging data. The inferred time scales of P_c reach a plateau after about $\tau = 0.75$ s as in Fig. 5(C); dashed horizontal line highlights $|\Lambda_c^{-1}|(\tau^*)$. As in the Lorenz system, $|\Lambda_c^{-1}|$ are shorter than $|\hat{\Lambda}_2^{-1}|$ across τ , reflecting the large scale irreversibility of the dynamics. In fact, while $|\hat{\Lambda}_2^{-1}|(\tau^*)$ from P_r overestimates the expected $|\Lambda_2^{-1}|$ from “run” and “pirouette” transition rates, the timescales obtained from P , $|\Lambda_c^{-1}|(\tau^*) = 3.48 \pm (3.03, 3.86)$ s, accurately predict the hopping dynamics. Error bars are 95% confidence intervals bootstrapped across worms.

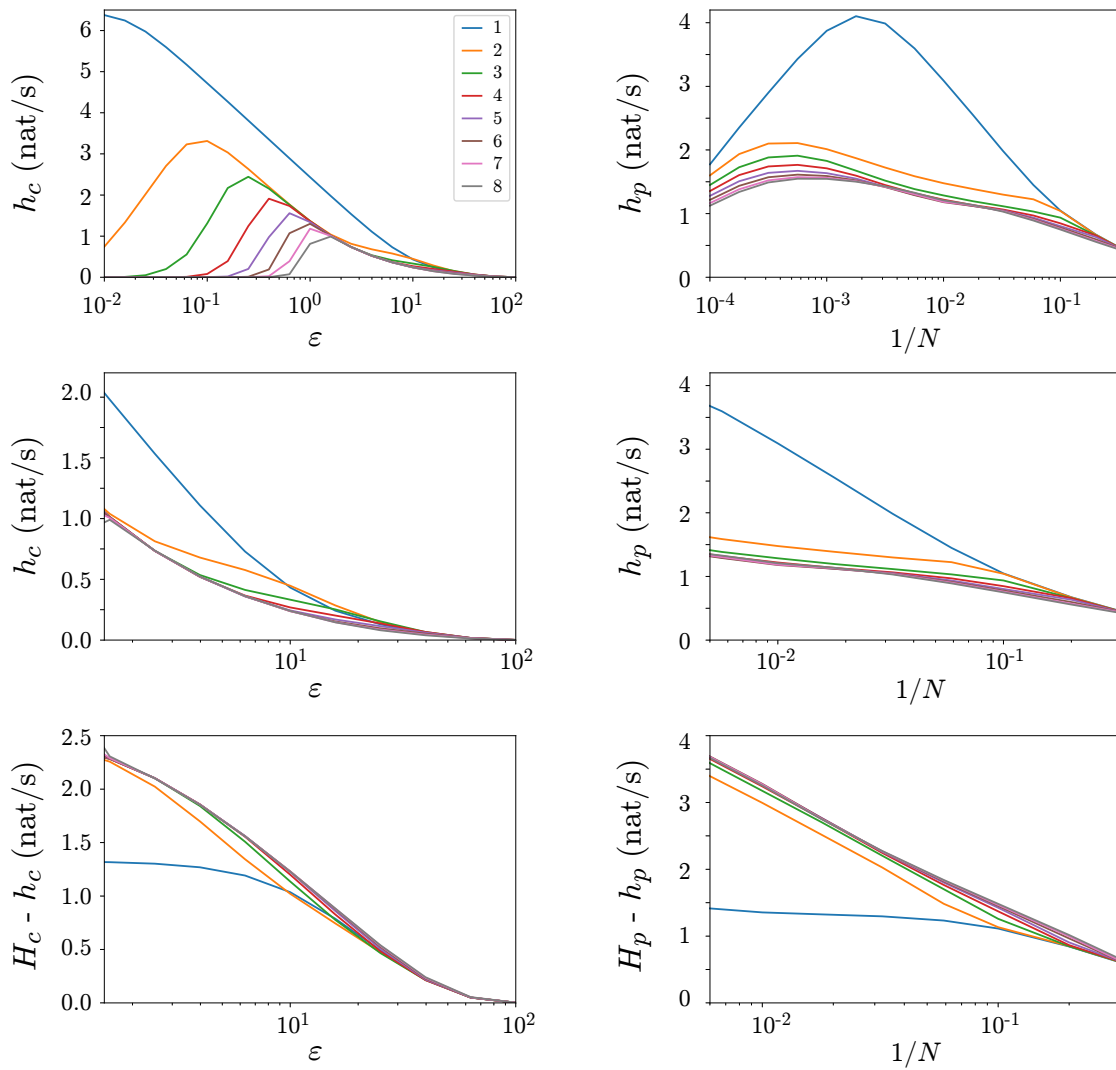


FIG. S10. Comparison between the Cohen and Procaccia method [40] (left) and our partition based approach (right) to estimate entropy rates and predictive information as a function of the number of delays K for an AR(2) model. (Top) Entropy rate as a function of length scale ϵ computed using Eq. (19) (left) and as a function of the inverse of number of partitions $1/N$ computed using Eq. (5). Colors indicate the growing number of delays K . The geometric based approach suffers from finite size effects with the growing dimensionality of the space (from adding K into y^K), but the entropy rate is generally higher for $K = 1$ than for $K > 1$. Similarly, h_p exhibits finite size effects for large numbers of partitions, and the estimates of the entropy rate drop. (Middle) Same as the top plot but now focused on a length scale regime at which the finite size effects of the estimates of entropy are not present. Note that the growth of the entropy with decreasing length scale (increasing number of partitions) is much faster for $K = 1$. (Bottom) Predictive information as a function of length scale ϵ (left) and the inverse of the number of partitions $1/N$ (right). For $K > 1$ the reconstructed state space maximizes the predictability.