

Implications of hierarchical Markov models of behavior: on irreversibility, predictability, and dimensionality

John J. Vastola and Kanaka Rajan

Department of Neurobiology, Harvard Medical School
Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University
{john_vastola, kanaka_rajan}@hms.harvard.edu

Abstract

The maturation of quantitative tools for studying the high-level structure of animal behavior, and especially tools which represent spontaneous behavior as a sequence of stereotyped and neurally well-defined ‘syllables’, demands that the field revisit a fundamental theoretical question: if the coarse structure of behavior can be accurately described by Markov models, what do these models really tell us about behavior? In this work, we explore the theoretical implications of these models and discuss how they allow us to quantitatively formulate questions about the sequence-like nature and effective dimensionality of behavior. One important insight is that the eigenvalues and eigenvectors of various model-associated matrices furnish interpretable time scales and modifications of behavior that occur on those time scales. We illustrate our points using both toy examples and Markov models fit to real data. By analyzing the consequences of Markov representations, we clarify the theoretical meaning of progress in quantifying behavior.

Introduction

Understanding the high-level structure of animal behavior—particularly the spontaneous and innate behavior associated with naturalistic environments (Tinbergen, 1951)—has only recently become quantitatively tractable. Techniques for the large-scale quantification of behavior, including algorithms for tracking animal positions and poses for at least tens of hours (Mathis et al., 2018; Pereira et al., 2020, 2022), enable us to precisely formulate and answer questions about how behavior evolves across contexts. One revelation from these quantitative ethology studies is that behavior is perhaps simpler than anyone might have expected: there is good evidence that behavior is not infinitely variable, but consists of fairly stereotyped sequences of discrete *syllables* (Weinreb et al., 2024; Wiltchko et al., 2015). Simple syllables, like ‘walking’ or ‘sniffing’, can be chained together to form more complex behaviors, like investigating objects or exploring an environment.

In the syllable-based view, behavior is quantitatively characterized via Markov chain models with finitely many states. Animals behave *as if* there is a constant probability per unit time of transitioning to a different discrete

behavior, in a fashion generally modulated by both the current behavior and some kind of latent state (e.g., hunger). Older efforts to construct Markov models of behavior emphasized the poor quality of their fit to data (or equivalently, signatures of ‘non-Markovianity’), which was interpreted as a sign that unobservable latent states acted as confounds (Berman et al., 2016; Overman et al., 2022; Schwarz et al., 2015). But more recent efforts have shown that Markov models provide an accurate description of the behavior of *C. elegans* (Costa et al., 2024), larval zebrafish (Sridhar et al., 2024), and mice (Weinreb et al., 2026), provided one explicitly accounts for latent states. The work of Weinreb et al. (2026) is particularly noteworthy, because in addition to identifying behavioral syllables and latent states which modulate them, it shows that these syllables and latent states have neural correlates in the dorsomedial prefrontal cortex. When considered in light of the earlier finding that the dorsolateral striatum encodes behavioral syllables (Markowitz et al., 2023), one finds strong support for the claim that Markov models capture biologically meaningful structure rather than arbitrary abstractions.

It stands to reason that these techniques will improve and lead to even better Markov models of behavior. These improvements highlight a *theoretical* question more subtle than questions concerning fit quality: what do these models *actually tell us* about the high-level structure of behavior? In this work, we argue that drawing upon the rich mathematical literature associated with Markov models yields interesting tentative answers to fundamental questions about behavior. Here, we focus especially on questions regarding (i) the sequence-like character of behavior, and (ii) the effective dimensionality of behavior, but our point is in principle broader. A key insight is that spectral decompositions associated with Markov chains can provide interpretable time scales and modes of behavior, which are ‘global’ in the sense that they depend on the high-level structure of the Markov chain rather than on just a small number of syllable transition rates. But there is not just one informative decomposition—there are at least two, and each is naturally associated with a somewhat different question.



We will begin by reviewing the mathematics of Markov models and their application to behavior. In the following sections, we consider the sequence-like character of behavior, the dimensionality of behavior, and the consequences of hierarchical structure. Throughout, we illustrate our points using both carefully chosen toy models and a model fit to data from Weinreb et al. (2026).

Markov models of behavior

The discreteness of behavior

The primitive insight that suggests Markov models of behavior is that, although behavior certainly involves many approximately *continuous* degrees of freedom—which implies that the number of possible behaviors is effectively infinite—there is on the other hand a great deal of stereotypy. This is especially apparent when behavior is suitably coarse-grained, so that slight variations of one movement (e.g., a reach) are treated as identical. Following work on MoSeq (Weinreb et al., 2024; Wiltchko et al., 2015), we will call discrete units of behavior *syllables*, but the reader should note that essentially the same concept has been used with other names, like ‘eigenworms’ (Stephens et al., 2008) and ‘bout types’ (Johnson et al., 2020; Marques et al., 2018).

Motivated by this insight, in what follows we ignore continuous structure and treat behavior as discrete. We consider models with $N_B \geq 1$ possible behaviors, and use an integer $b \in \{1, \dots, N_B\}$ to index each behavior.

From bags to bigrams to Markov models

The simplest discrete model of behavior is probably a ‘bag of behaviors’ model akin to ‘bag of words’ models of language (Wallach, 2006). One can construct such a model by observing an animal for a long period of time T , and then determining the total amount of time T_b it spent performing each behavior b . Assuming behavioral syllables are defined in a way that prevents the animal from doing more than one at a time, and that the set of N_B behaviors covers all possibilities, the **bag model** defines the probability of performing b as $\pi_b := T_b/T$. This model decides what to do next by ‘drawing from the bag’ according to the probabilities $\pi := (\pi_1, \dots, \pi_{N_B})^T$, regardless of the current behavior.

Although this model has some use in naturalistic settings (Altmann, 1974), particularly since it requires less data to fit than more sophisticated models (there are only $N_B - 1$ free parameters), it is a bad model of behavior. It is bad principally because it neglects *transition* structure (Fig. 1a): for example, blowflies alternate between front leg grooming and anterior grooming movements in a stereotyped fashion (Dawkins & Dawkins, 1976).

The most obvious way to improve the bag model is to instead compute **bigrams**: how often behavior i follows behavior j . To measure these, divide the total amount of time T an animal behaves into time bins, each of size

$\sim \Delta t$. This ‘mesoscopic’ time scale Δt must be chosen small enough that a single behavior dominates each bin. One then obtains the transition probabilities p_{ij} by counting the number of times an i bin follows a j bin.

Markov models of behavior assume bigram structure completely characterizes dynamics. We can construct a continuous-time model from bigram statistics by assuming, for small enough Δt , that the probability p_{ij} of a j to i transition scales linearly with Δt , i.e., $p_{ij} = k_{ij}\Delta t + \mathcal{O}((\Delta t)^2) \approx k_{ij}\Delta t$. We can then construct

$$H_{ij} := \begin{cases} -\sum_{b \neq i} k_{bi} & \text{for } i = j \\ k_{ij} & \text{for } i \neq j, \end{cases} \quad (1)$$

which we call a **transition matrix**. H is *infinitesimal stochastic* since (i) its off-diagonal entries, which represent transition rates, are nonnegative; and (ii) its columns sum to zero (Baez & Biamonte, 2018).

This matrix defines a continuous-time Markov chain model of behavior (Fig. 1b, c) via the **master equation**

$$\frac{d}{dt}\mathbf{p}(t) = \mathbf{H}\mathbf{p}(t) \quad \text{i.e.,} \quad \frac{d}{dt}p_b(t) = \sum_{a=1}^{N_B} H_{ba}p_a(t) \quad (2)$$

where $p_b(t)$ is the probability that the animal is performing behavior $b \in \{1, \dots, N_B\}$ at time t , and $\mathbf{p}(t) := (p_1(t), \dots, p_{N_B}(t))^T$. Note that, because the columns of H sum to zero, probability is conserved, since $\mathbf{1}^T H = \mathbf{0}^T$ implies $\frac{d}{dt}\mathbf{1}^T \mathbf{p}(t) = \mathbf{1}^T \mathbf{H}\mathbf{p}(t) = \mathbf{0}^T \mathbf{p}(t) = 0$.

Parenthetically, we note that using a continuous-time (rather than discrete-time) formulation has a number of advantages for studying behavior. In continuous-time formulations, it is natural to think in terms of bin-size-independent rates and time scales, whereas in discrete-time formulations model properties depend on Δt .

What about non-Markovianity?

It has been pointed out numerous times that bigram statistics are generally insufficient for describing real animal behavior (Berman et al., 2016; Schwarz et al., 2015), which is equivalent to the statement that Markov models are a poor description. But as previously mentioned, various authors have shown that the situation changes if latent states are included in the model’s notion of ‘state’ (Costa et al., 2024; Sridhar et al., 2024; Weinreb et al., 2026). In this work, we take the fairly uncontroversial view that non-Markovianity is merely a sign that the state space ought to be enlarged.

Technical assumptions restricting models

We do not allow all possible models of behavior (or equivalently, all possible infinitesimal stochastic matrices H), because some are degenerate in one or more senses. In brief (see Appendix A for more discussion), we assume: (i) there are no states that it is impossible to enter;




Figure 1: **Markov chain models of behavior.** (a) Transitions between different discrete behaviors (here, labeled 1, 2, and 3) are not random, but structured. Different colors represent different stochastic simulations. (b) Markov models of behavior can be visualized as graphs whose nodes represent behaviors and behavioral states, and whose arrows reflect transition rates. In the pictured 3-state model, it is impossible to transition directly between 1 and 3. (c) The transition matrix H defines a dynamics in probability space (here: the space of distributions of 3 behaviors). All Markov models we consider eventually converge to a unique steady state distribution π (which here is uniform).

(ii) there are no states it is impossible to leave; (iii) all states are connected by some directed path through the transition graph; and (iv) H is diagonalizable. These assumptions ensure that there exists a well-defined **steady state distribution** $\pi := \lim_{t \rightarrow \infty} p(t)$ which Eq. 2 always asymptotically approaches (Fig. 1c). Up to scaling, π is also the unique (right) eigenvector of H with eigenvalue zero. In some sense, using the bag model amounts to throwing away H and keeping only π .

Irreversibility and action sequences

A fundamental feature of animal behavior is that it involves sequences of actions with a fairly stereotyped ordering. For example, in fruit fly courtship a male typically orients toward a female, taps her abdomen with his forelegs, extends one wing, and then begins a courtship song (Vilella & Hall, 2008). If we take Markov models seriously as a description of behavior, it is interesting to note that not *all* possible models generate predictable sequences of actions; rather, different models do this more or less, in a way that can be quantified.

Sequences reflect irreversibility

One way of mathematically formalizing the idea that behavior consists of stereotyped sequences is via the concept of **irreversibility**: if we watch a movie of a behaving animal, it ought to look strange when played backwards. In the context of Markov models, we can define a precise notion of ‘playing the movie backwards’ by defining the time-reverse of a behavioral Markov chain. If $k_{ji}\Delta t$ is the probability of moving from state j to state i in a small amount of time Δt , then by Bayes’ rule the time-reversed chain has a small-time transition probability

$$\tilde{p}_{ij} := \frac{p_{ji}\pi_i}{\pi_j} = k_{ji} \frac{\pi_i}{\pi_j} \Delta t, \quad (3)$$

where we fix the prior by assuming the chain is stationary. Let D denote the diagonal matrix with $D_{ii} = \pi_i$ for all i .

These probabilities can be associated with the transition matrix $\tilde{H} := D H^T D^{-1}$, called the *adjoint* of H , which is infinitesimal stochastic and hence defines its own Markov chain. This chain’s graph looks just like the graph of the original chain, but with all arrows reversed (Fig. 2a). As one might expect, for chains where transitions tend to be bidirectional (which one might call ‘diffusive’), the time-reversed chain tends to look similar to the original chain; meanwhile, chains which involve directed cycles tend to look quite different from their time-reverse.

Detailed balance as a null hypothesis

What is the *least* sequence-like Markov model of behavior? If we accept that irreversibility is a reasonable way to quantify sequence-like-ness, then we can define a ‘null hypothesis’ via the concept of **detailed balance**. Models satisfying detailed balance obey the condition

$$H_{ij}\pi_j = H_{ji}\pi_i \quad (4)$$

for all pairs of behaviors i and j , which is equivalent to asking that $H = \tilde{H}$ (Kelly & Yudovina, 2014). This condition says that, when the chain is at steady state, any transition one observes is equally likely to have been observed in the opposite order. One cannot tell whether a movie of behavior was played ‘forwards’ or ‘backwards’, because there is no meaningful difference.

The concept of detailed balance will be useful to us for a variety of reasons, even beyond this discussion of irreversibility; we will see in the following sections that it makes it easier to define the dimensionality of behavior.

How irreversible is behavior?

We will now consider two questions. First, how can we measure the extent to which behavior is irreversible given a Markov chain description of it? Second, how irreversible is real animal behavior? Since irreversibility reflects the extent to which H differs from its adjoint $\tilde{H}$, a




Figure 2: **Quantifying irreversibility.** (a) One way to visualize the time-reverse of a Markov model is to reverse the arrows of the corresponding graph. However, note that $\tilde{H}_{ij} \neq H_{ji}$ in general. Pictured: the cyclic model with $N = 3$ states. (b) The entropy production rate (EPR) of the cyclic model equals $(k_f - k_r) \log_2(k_f/k_r)$. It increases as circulation around the ring becomes more directed. (c) The EPR differs for different latent states in a Markov model fit to open-field mouse behavior, and is highest in the ‘social engagement’ state.

simple way to quantify it is to measure the fluxes

$$J_{ij} = H_{ij}\pi_j - H_{ji}\pi_i = \pi_j(H_{ij} - \tilde{H}_{ij}) \quad (5)$$

for all pairs of behaviors i and j . (Since $J_{ij} = -J_{ji}$, ordering does not matter.) In simple cases, we can compute J_{ij} explicitly. For the **cyclic model**, which has N states arranged along a ring so that ‘forward’ transitions occur at a rate $k_f > 0$ and ‘reverse’ transitions occur at a rate $k_r \geq 0$ (see Appendix B), $|J_{ij}|$ equals $|k_f - k_r|/N$ for neighboring states and zero otherwise. But while measuring fluxes has the advantage that it localizes irreversibility to specific sets of transitions, there are at least two problems: the J_{ij} can be small because they scale with N , and it is not obvious how to compare them across systems. Alternatively, we can define an aggregate measure of irreversibility, with the **entropy production rate** (EPR) (Lynn et al., 2021; Seifert, 2005)

$$\dot{S} := \sum_{i \neq j} H_{ij}\pi_j \log_2 \left(\frac{H_{ij}\pi_j}{H_{ji}\pi_i} \right) \quad (6)$$

a canonical choice. For chains satisfying detailed balance, all fluxes J_{ij} are zero and the EPR is exactly zero. For the cyclic model $\dot{S} = (k_f - k_r) \log_2(k_f/k_r)$ (Fig. 2b), so it increases as k_f and k_r are made increasingly different.

To briefly address the irreversibility of real behavior, we quantified the EPR associated with a hierarchical Markov model fit to open-field mouse behavior from Weinreb et al. (2026) (see Appendix C). This model has five high-level latent states which modulate behavior, labeled ‘grooming’, ‘exploratory locomotion 1’, ‘exploratory locomotion 2’, ‘local investigation’, and ‘social engagement’. We find that the H associated with each latent has a nonzero EPR, with the ‘social engagement’ state having around double the EPR of the other states (Fig. 2c). This suggests the unsurprising conclusion that real behavior is indeed at least somewhat irreversible, and can be more or less irreversible in different contexts.

Dynamics and dimensionality

Another fundamental feature of behavior is that it tends to be somewhat low-dimensional, in the sense that the actions and action sequences actually executed by animals occupy an extremely restricted region of the space of all possible actions and action sequences. Does a Markov model description of behavior tell us anything about this? While it cannot tell us about dimensionality in terms of the *number* of actions animals can access, since by construction it only involves the actions an animal actually performed, it can say something interesting about the dimensionality of *transitions* between behaviors, and hence provide a nontrivial notion of the dimensionality of behavior. We elaborate on this idea below.

Intuition: distribution dimensionality

Perhaps the most naive way to define the dimensionality of behavior is to consider grouping or coarse-graining behaviors themselves; in our Markov model picture, this corresponds to grouping syllables with similar connectivity, for example via graph clustering algorithms. On the other hand, in moving from continuous behavior to discrete syllables, some such coarse-graining (e.g., treating many different locomotion patterns as “walk” or “run”) has already been done. It is not clear what useful information additional syllable-level coarse-graining tells us.

A different and arguably more useful notion of dimensionality concerns the behavior distribution $\mathbf{p}(t)$. Because $\mathbf{p}(t)$ is a (normalized) vector in $\mathbb{R}^{N_B}$, how $\mathbf{p}(t)$ evolves in time can be visualized as a trajectory in $\mathbb{R}^{N_B-1}$. Any change in context—for example, some change in the environment, or change in the animal’s latent state—will cause $\mathbf{p}(t)$ to be different than π , so that the Markov chain is out of equilibrium. As time passes, it will approach π along a path which depends on its initial condition. One way to pose the question of dimensionality is: for many different initial conditions, what is the dimensionality of these trajectories?




Figure 3: **Dimensionality via dynamics.** From left to right: model schematic, eigenvalues λ_k of $\mathbf{H}$, time scales $\tau_k := -1/\text{Re}(\lambda_k)$ ($\lambda_k \neq 0$) ordered from greatest to least, and effective (dynamical) dimensionality as a function of the time scale τ . Each row depicts these for a different model. **(a)** The bag model ($N = 20$ used) has just two distinct eigenvalues and one time scale, and its dimensionality is either maximal or zero. **(b)** The cyclic model ($N = 20$ used) has many distinct eigenvalues, which are real if $k_f = k_r$ (gray) but can be complex-valued if $k_f \neq k_r$ (red). The top two nonzero modes provide a reasonable approximation to the system for many choices of τ , so the system is roughly two-dimensional. **(c)** Mice in a ‘grooming’ state exhibit behavior involving a range of sub-second time scales, and dimensionality is not well summarized by a small number of modes. **(d)** In the ‘exploratory locomotion 1’ state, the imaginary parts of some eigenvalues are larger, which suggests behavior is more irreversible.

Since $\mathbf{p}(0)$ can in principle be anything, one may suspect the answer is always $N_B - 1$, the dimensionality of the whole space. But this is *not* true if we *additionally* coarse-grain over time, and only answer our question above a time scale $\tau \geq 0$. On long enough time scales, the space of trajectories generally looks lower-dimensional—and in the $\tau \rightarrow \infty$ limit, the space collapses to the point at π . One way to visualize this is to imagine cutting out the beginning of many movies of behavior (i.e., frames before τ), and then considering only the remaining frames of the trajectories.

For some moderate value of τ , when do we expect the dimensionality to be lower than maximal? This can happen when there is a separation of time scales in behavioral transitions, which can cause some adjustments to the behavior distribution to happen quickly (or at least, faster than τ) and others to happen more slowly. A 3-

state model of sleep provides a trivial albeit informative example: animals may transition between ‘more awake’ and ‘less awake’ states relatively quickly (see Appendix B), but between these and a ‘sleep’ state slowly. Beyond the fast time scale, the model ‘looks’ like it has only two states (‘awake’ and ‘sleep’), and $\mathbf{p}(t)$ trajectories are effectively one-dimensional.

Making this intuition rigorous

We can formalize this intuition by invoking a mathematical result known as the Perron-Frobenius theorem. This theorem says that given certain technical assumptions that rule out degenerate cases (see Appendix A), the behavior distribution $\mathbf{p}(t)$ has a spectral decomposition

$$\mathbf{p}(t) = \pi - \sum_{k=1}^{N_B-1} c_k \mathbf{v}_k e^{\lambda_k t} = \pi - \sum_{k=1}^{N_B-1} c_k \mathbf{v}_k e^{-t/\tau_k + i\omega_k t} \quad (7)$$

where $\lambda_k = -1/\tau_k + \omega_k i$ is the k -th eigenvalue of $\mathbf{H}$, $\mathbf{v}_k$ is the k -th right eigenvector of $\mathbf{H}$, and the c_k are initial-condition-dependent coefficients. This bears on trajectories through distribution space because the $\mathbf{v}_k$ form a basis for that space, and $\mathbf{v}_k$ only meaningfully contributes to $\mathbf{p}(t)$ on time scales shorter than τ_k . Hence, we will say that the **dynamical dimensionality** of a model is r on a time scale τ if there are $0 \leq r \leq N_B - 1$ eigenvectors whose time scale is at least τ .

Behavioral modes and time scales

This result also provides an interesting formalization of behavioral modes and time scales. The components of each $\mathbf{v}_k$ sum to zero, so each eigenvector represents adjustments to the behavior distribution (i.e., do some behaviors more and others less) which happen on a time scale τ_k . Consider, for example, the 3-state sleep-wake Markov chain: one eigenvector corresponds to equilibration of the awake states, and the other corresponds to movement between the awake and asleep states (see Appendix B). One interesting feature of these summaries is that they are ‘global’: they do not depend only on specific behavioral transitions, but on the structure of transitions writ large. Generally, multiple transition rates k_{ij} factor into any eigenvalue or eigenvector, as one can see by studying simple examples.

A Markov chain version of the **bag model** (see Appendix B) has maximal dimensionality until τ exceeds a critical value (Fig. 3a). The cyclic model, whose transitions are much more structured, is effectively two-dimensional for a broad range of τ (Fig. 3b). The model of open-field mouse behavior from Weinreb et al. (2026) involves a wider range of sub-second time scales than these toy models, and is not obviously low-dimensional unless τ is at least a second (Fig. 3c, d).

Predictability and dimensionality

Unfortunately, the notion of dimensionality defined in the previous section is not without problems. It does not necessarily answer a related question: how many *orthogonal* vectors are required to characterize changes in the behavior distribution beyond a time scale τ ? If detailed balance is violated and a model is highly irreversible, it is possible for its eigenvectors to be linearly independent but nearly aligned. This is essentially another manifestation of the well-known issue that non-normal dynamics are more complex than normal dynamics (Ganguli et al., 2008). This issue motivates a different notion of dimensionality associated with predictability, which we call a model’s **predictive dimensionality**.

Predictability yields a PCA analogue

Following Bialek (Bialek, 2022; Bialek et al., 2001), a principled information-theoretic approach to our question requires that we quantify the mutual information between

the past and future. But a key difficulty with this approach is that mutual information and associated quantities can be difficult to actually compute. An alternative is to approximate mutual information as a simpler quadratic form; one can show (see Appendix D) that the mutual information between the system’s state at one time, and its state after τ additional time, is roughly

$$I(\tau) \approx \frac{1}{2} \text{tr} \mathbf{K}(\tau) - \frac{1}{2}, \quad (8)$$

where $\mathbf{K}(\tau) := \mathbf{D}^{1/2} e^{\mathbf{H}^T \tau} \mathbf{D}^{-1} e^{\mathbf{H} \tau} \mathbf{D}^{1/2}$ is the past-future kernel. Unlike $\mathbf{H}$, this matrix is always symmetric and positive definite, so its eigenvalues ν_k are all real and positive. This is important because $\mathbf{K}(\tau)$ can essentially play the role that the covariance matrix does in principal component analysis (PCA): we can approximate $\mathbf{p}(t)$ using eigenvectors of $\mathbf{K}(\tau)$ to capture as much ‘variance’ as possible. Supposing eigenvalues are ordered from greatest to least, and that $\nu_k = 1$ is excluded, we say that the top r eigenvectors explain the fraction

$$\text{ve}(\tau) := \frac{\nu_1 + \dots + \nu_r}{\sum_{k=1}^{N_B-1} \nu_k} \leq 1 \quad (9)$$

of the overall variance. We say a model has **predictive dimensionality** r on a time scale τ if r is the minimum number of modes required for $\text{ve}(\tau)$ to exceed a desired threshold (e.g., 90%). Interestingly, predictive and dynamical dimensionality look similar for systems which obey detailed balance. Since $\mathbf{K}(\tau) = \mathbf{D}^{-1/2} e^{\mathbf{H} \tau} e^{\mathbf{H}^T \tau} \mathbf{D}^{1/2}$, detailed balance implies that $\nu_k = e^{2\lambda_k \tau}$, and that the eigenvectors of $\mathbf{K}(\tau)$ are those of $\mathbf{H}$ up to a $\mathbf{D}^{-1/2}$ rescaling. Away from detailed balance, $\mathbf{K}(\tau)$ can be viewed as producing a certain ‘reversibilization’ of the original model, which approximates it with a model that *does* satisfy detailed balance (Choi & Wolfer, 2024).

Behavioral modes and time scales

The eigenvalues and eigenvectors of $\mathbf{K}(\tau)$ are all real, and hence easier to interpret: they are time scales on which a given orthogonal behavioral mode is relevant for predicting changes in the behavior distribution $\mathbf{p}(t)$, in the sense of maximizing predictability. Dynamical and predictive dimensionality look very similar for the toy models we consider (Fig. 4a, b), but predictive dimensionality tends to be significantly larger than dynamical dimensionality for models fit to open-field mouse behavior (Fig. 4c, d). We speculate that this is because, when $\mathbf{H}$ is highly irreversible and many of its eigenvectors are aligned, one must use more orthogonal vectors to capture the same amount of predictability.

Consequences of hierarchy

Thus far we have ignored a major impetus for technical advances: non-Markovianity in observable-behavior-only




Figure 4: **Dimensionality via predictability.** From left to right: model schematic, ‘variance’ explained as a function of the number of modes used (different colors correspond to different values of the imposed time scale τ), time scales $\tau_k := -2\tau / \log(\nu_k)$ ($\nu_k \neq 1$) ordered from greatest to least (red dashed line: time scales of $\mathbf{H}$), and predictive dimensionality vs τ (here, number of modes required to capture $\geq 90\%$ variance; red dashed line: dynamical dimensionality vs τ). Each row depicts these for a different model. **(a)** Each mode of the bag model ($N = 20$ used) explains an equal amount of variance, and dimensionality is near-maximal. **(b)** A small number of modes explains most variance for the cyclic model ($N = 20$ used), especially for larger τ . **(c-d)** Few modes explain most variance when τ is more than a few seconds, but predictive dimensionality tends to be larger than dynamical dimensionality.

models (Berman et al., 2016), which can be addressed by fitting a hierarchical model with latent states (Weinreb et al., 2026). The addition of latent states only really changes our definition of ‘state’—it now means a specific behavior *and* latent state. On the other hand, because latent state transitions occur on a time scale somewhat slower than behavioral transitions, these larger Markov chains are endowed with special structure. We briefly discuss the consequences of this structure below.

Whole properties versus part properties

One striking consequence of combining a latent-state-only chain with a behavior-only chain is that, even if both chains are reversible, the larger chain may not be—and hence, overall behavior can be sequence-like. To understand this point, consider a simple example in which the animal possesses only two possible behaviors and two possible latent states. The latent-state-only and behavior-

only chains must satisfy detailed balance, because all two-state chains do. But the overall chain can be cyclic, and hence arbitrarily irreversible. For example, consider the four-state model with transitions

$$(s, b) \rightarrow (s, b') \rightarrow (s', b') \rightarrow (s', b) \rightarrow (s, b) \rightarrow \dots$$

For similar reasons, it is also possible to obtain a reversible model by combining irreversible chains. Hence, qualitatively new properties can appear at the level of the full model. This also means that it is not a priori obvious whether enlarging the state space will increase or decrease effective dimensionality, since highly irreversible models (like the cyclic model) can be low-dimensional.

Spectral structure of hierarchical chains

If latent state and behavioral transitions were completely independent, the spectral structure of the larger chain would be extremely simple: the eigenvalues of the full




Figure 5: **Hierarchical spectral structure in behavior.** The first five columns depict the eigenvalues of $\mathbf{H}$ (top row) and associated time scales (bottom row, ordered from greatest to least) in different latent states. The sixth column shows all eigenvalues overlaid, which closely matches the eigenvalues of the full latent-state-and-behavior model (last column, top row). But the longest time scales of the full model (bottom, gray arrow) come not from any behavior-only $\mathbf{H}$, but from slow dynamics associated with transitions between latent states.

$\mathbf{H}$ would be a *sum* of the eigenvalues of the state-only and behavior-only matrices, and its eigenvectors would be a tensor product of each matrix’s eigenvectors. Interestingly, something similar is true when latent states and behavior can influence one another, provided that each type of transition occurs on very different time scales: the eigenvalues of the bigger chain are approximately the union of the eigenvalues associated with each latent state’s behavior-only chain, and the eigenvectors of the bigger chain correspond to fast behavioral-change modes and slow state-change modes. This result follows from singular perturbation theory (Yin & Zhang, 2012), and may not be accurate if the time scale separation is not sufficiently large. This structure appears in the hierarchical model from Weinreb et al. (2026): the bulk of the spectrum of the full open-field mouse behavior model (Fig. 5, last column, top) looks like a union of the spectra associated with each latent state (Fig. 5, top row). But the slowest time scales, which are not present in any of the behavior-only models (Fig. 5, bottom row), are instead inherited from latent state dynamics.

Discussion

We explored various consequences of Markov models of behavior, with a particular focus on quantifying its sequence-like character and effective dimensionality. The main contribution of this work is not a new model of behavior, but a clarification of what *must* be true if Markov models are a good description of it.

Several of the consequences we identify do not immediately follow from the observation that the structure of $\mathbf{H}$ reflects the high-level structure of behavior. For example, our analysis shows that there are multiple coherent ways to define the effective dimensionality of behavior, with

one notion reflecting dynamical structure and another relating to predictability. This distinction becomes especially important for highly irreversible models, since these notions can become substantially different. More generally, our results suggest that taking Markov models seriously as descriptions of behavior can help uncover conceptual issues that are easy to overlook when such models are treated purely as descriptive summaries.

Objections and limitations

At this stage, it is worth explicitly discussing possible objections to the approach we have outlined here. First, by treating behavior as the primary object of study, are we making the same mistakes as the behaviorists (Chomsky, 1959; Skinner, 1950)? The answer is no. That our approach works at all depends on taking one of the major lessons of cognitive science and neuroscience—the idea that behavior ought to be understood in terms of internal states—seriously, since Markov models of behavior *without* latent states tend not to fit (Costa et al., 2024; Sridhar et al., 2024; Weinreb et al., 2026).

Second, what do we learn from fitting and analyzing Markov models? Why quantify time scales, irreversibility, and dimensionality when quantifying them is not the primary objective? We think there are two reasons. Doing so allows us to make more precise claims; an animal’s behavior is not just low-dimensional, but (e.g.) ten times lower-dimensional than when the animal is in some other state. These quantities also furnish what might fairly be called ‘statistics’ of behavior that can be compared across conditions. When the environment changes, we can describe associated behavioral changes not just qualitatively, but in terms of quantitative changes in irreversibility, dimensionality, and so on.

Third, how sensitive are these analyses to the choice of granularity? Most complex syllables can be decomposed into smaller movements (e.g., walking consists of many left-foot and right-foot steps), and many types of movements (e.g., a reach) can be further divided into subtypes (e.g., a reach to the left or right), so it seems unlikely that we will learn anything biologically meaningful if our results depend strongly on how we have chosen to parcellate behavior. In this paper, we have chosen to neglect the question of whether one model of behavior is more appropriate than another, and have concentrated on the simpler question of what a given model implies. But this objection to our approach is significant, because it is almost trivially true that one's results depend strongly on the parcellation of behavior used. For example, when it is possible to coarse-grain behavior so that there are only two states—'active' and 'inactive', say—our analyses would conclude that the system is reversible and low-dimensional. Meanwhile, if we treated a very large number of behavioral states as distinct, we would likely come to the opposite conclusion.

This is a fundamental limitation of our approach: what we learn depends on the assumed model, which includes the details of how behavior is coarse-grained. But this is not necessarily a fatal flaw for two reasons. Principles of parsimony may suggest that a certain coarse-graining balances goodness of fit against model complexity. Also, if multiple levels of description are of interest, one can simply analyze each of the corresponding Markov models and get numbers (e.g., dimensionalities) associated with each scale. This idea is in the spirit of analyzing 'barcodes' (collections of some statistic across scales), a strategy used in topological data analysis (Chaudhuri et al., 2019; Curto & Sanderson, 2025).

Finally, an important limitation of the analyses we have discussed is that they may become practically unfeasible when the number of states and behaviors becomes too large. This may happen in a few ways: behavior may not have been sufficiently coarse-grained (e.g., too many different reaches are treated as distinct); explaining data may require introducing too many latent states; or the environment is too complicated. Of these, environmental complexity is perhaps the easiest way to make state space prohibitively large. Outside of dark open-field contexts where not much is happening, the dynamics of environmental transitions ought to be included in a Markov description of an animal's behavioral transitions if they significantly modulate behavior. In the worst case, the overall Markov model's size may scale like $\approx (N_E \cdot N_S \cdot N_B)^2$, where N_E is the number of distinct environmental states and N_S is the number of distinct latent states. In addition to making certain analyses (e.g., of eigenvalues and eigenvectors) slower, large increases in the complexity of the environment may make the inference of a Markov model from data extremely difficult.

Hence, while Markov models may be indispensable for studying spontaneous behavior in simple environments, the analyses discussed here may be less useful in richer environments.

RL and going beyond animal behavior

At first glance, Markov models appear to provide a slightly different picture of animal behavior than reinforcement learning (RL) (Sutton & Barto, 2018). First, in RL behavioral policies are typically conditioned on the animal's current state rather than their current or previous action; hence, accommodating Markov models in an RL framework requires that the notion of state include information about previous actions. This is not unprecedented, as treating recent actions as part of what we mean by 'state' can be beneficial in robotics (Radosavovic et al., 2024) and for efficiently learning temporally-extended action sequences (Bacon et al., 2017).

Second, the (current) Markov model picture is descriptive rather than normative, in the sense that it does not say why an animal's H matrix looks the way it does. One way to connect our picture with RL, especially in the absence of an experimenter-defined task and reward structure, is to assume that the animal acts in a way that maximizes intrinsic rewards (Gershman, 2025). For example, if information is intrinsically rewarding (Bromberg-Martin et al., 2024; Bussell et al., 2024), animals may be driven to investigate novel objects.

Finally, we note that our prescription for studying behavior is somewhat agnostic to the choice of animal, the time scale on which their behavior is monitored, and the extent to which their behavior is abstracted, since we only require that a Markov model provides a good description. This approach can hence be applied not just to model organisms like mice, but to more exotic possibilities, including artificial agents trained to perform tasks like foraging (Simmons-Edler et al., 2025). Furthermore, some of the metrics we have discussed (e.g., entropy production rate) can in principle be computed and fairly compared across wildly different animals and contexts.

Artificial agents are a particularly interesting possible object of study, since although it is easy to access and track their moment-to-moment action sequence (e.g., move one tile left, move one tile right), we may be more interested in emergent behavior better described in terms of high-level actions like 'search for food' and 'flee from predator'. Such a higher-level description of behavior could plausibly be quantitatively treated via a Markov model, which implies that it may be fruitful to study performant artificial agents just like animals.

Acknowledgments

KR was funded by the NIH (RF1DA056403, U01NS136507), James S. McDonnell Foundation (220020466), Simons Foundation (Pilot Extension-00003332-02), McKnight Endowment Fund, CIFAR Azrieli Global Scholar Program, NSF (2046583), a Harvard Medical School Neurobiology Lefler Small Grant Award, and a Harvard Medical School Dean's Innovation Award. This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University.

Code availability

Code for reproducing this paper's figures is available at <https://github.com/john-vastola/markov-ccn26>.

References

- Altmann, J. (1974). Observational Study of Behavior: Sampling Methods. *Behaviour*, 49(3-4), 227–266. <https://doi.org/10.1163/156853974X00534>.
- Bacon, P.-L., Harb, J., & Precup, D. (2017). The option-critic architecture. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 1726–1734.
- Baez, J., & Biamonte, J. D. (2018). *Quantum techniques in stochastic mechanics*. World Scientific. <https://doi.org/10.1142/10623>.
- Berman, G. J., Bialek, W., & Shaevitz, J. W. (2016). Predictability and hierarchy in *Drosophila* behavior. *Proceedings of the National Academy of Sciences*, 113(42), 11943–11948. <https://doi.org/10.1073/pnas.1607601113>.
- Bialek, W. (2022). On the dimensionality of behavior. *Proceedings of the National Academy of Sciences*, 119(18), e2021860119. <https://doi.org/10.1073/pnas.2021860119>.
- Bialek, W., Nemenman, I., & Tishby, N. (2001). Predictability, Complexity, and Learning. *Neural Computation*, 13(11), 2409–2463. <https://doi.org/10.1162/089976601753195969>.
- Bromberg-Martin, E. S., Feng, Y.-Y., Ogasawara, T., White, J. K., Zhang, K., & Monosov, I. E. (2024). A neural mechanism for conserved value computations integrating information and rewards. *Nature Neuroscience*, 27(1), 159–175. <https://doi.org/10.1038/s41593-023-01511-4>.
- Bussell, J. J., Badman, R. P., Márton, C. D., Bromberg-Martin, E. S., Abbott, L., Rajan, K., & Axel, R. (2024). Representations of the intrinsic value of information in mouse orbitofrontal cortex. *bioRxiv*. <https://doi.org/10.1101/2023.10.13.562291>.
- Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., & Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9), 1512–1520. <https://doi.org/10.1038/s41593-019-0460-x>.
- Choi, M. C. H., & Wolfer, G. (2024). Systematic approaches to generate reversibilizations of markov chains. *IEEE Transactions on Information Theory*, 70(5), 3145–3161. <https://doi.org/10.1109/TIT.2023.3304685>.
- Chomsky, N. (1959). *Language*, 35(1), 26–58. Retrieved June 11, 2026, from <http://www.jstor.org/stable/411334>
- Costa, A. C., Ahamed, T., Jordan, D., & Stephens, G. J. (2024). A markovian dynamics for *Caenorhabditis elegans* behavior across scales. *Proceedings of the National Academy of Sciences*, 121(32), e2318805121. <https://doi.org/10.1073/pnas.2318805121>.
- Curto, C., & Sanderson, N. (2025). Topological Neuroscience: Linking Circuits to Function. *Annual Review of Neuroscience*, 48(Volume 48, 2025), 491–518. <https://doi.org/https://doi.org/10.1146/annurev-neuro-112723-034315>.
- Dawkins, R., & Dawkins, M. (1976). Hierarchical organization and postural facilitation: Rules for grooming in flies. *Animal Behaviour*, 24(4), 739–755. [https://doi.org/10.1016/S0003-3472\(76\)80003-6](https://doi.org/10.1016/S0003-3472(76)80003-6).
- Ganguli, S., Huh, D., & Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences*, 105(48), 18970–18975. <https://doi.org/10.1073/pnas.0804451105>.
- Gershman, S. J. (2025). Subjective functions. *arXiv e-prints*, Article arXiv:2512.15948, arXiv:2512.15948. <https://doi.org/10.48550/arXiv.2512.15948>.
- Hairer, M. (2006). Ergodic properties of markov processes. *Lecture notes*.
- Johnson, R. E., Linderman, S., Panier, T., Wee, C. L., Song, E., Herrera, K. J., Miller, A., & Engert, F. (2020). Probabilistic Models of Larval Zebrafish Behavior Reveal Structure on Many Scales. *Current Biology*, 30(1), 70–82.e4. <https://doi.org/10.1016/j.cub.2019.11.026>.
- Kelly, F., & Yudovina, E. (2014). *Stochastic networks* (Vol. 2). Cambridge University Press.
- Lynn, C. W., Cornblath, E. J., Papadopoulos, L., Bertolero, M. A., & Bassett, D. S. (2021). Broken detailed balance and entropy production in the human brain. *Proceedings of the National Academy of Sciences*, 118(47), e2109889118. <https://doi.org/10.1073/pnas.2109889118>.
- Markowitz, J. E., Gillis, W. F., Jay, M., Wood, J., Harris, R. W., Cieszkowski, R., Scott, R., Brann, D., Koveal, D., Kula, T., Weinreb, C., Osman, M. A. M., Pinto, S. R., Uchida, N., Linderman, S. W., Sabatini, B. L., & Datta, S. R. (2023). Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature*,

- 614(7946), 108–117. <https://doi.org/10.1038/s41586-022-05611-2>.
- Marques, J. C., Lackner, S., Félix, R., & Orger, M. B. (2018). Structure of the Zebrafish Locomotor Repertoire Revealed with Unsupervised Behavioral Clustering. *Current Biology*, 28(2), 181–195.e5. <https://doi.org/10.1016/j.cub.2017.12.002>.
- Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9), 1281–1289. <https://doi.org/10.1038/s41593-018-0209-y>.
- Overman, K. E., Choi, D. M., Leung, K., Shaevitz, J. W., & Berman, G. J. (2022). Measuring the repertoire of age-related behavioral changes in *Drosophila melanogaster*. *PLOS Computational Biology*, 18(2), 1–19. <https://doi.org/10.1371/journal.pcbi.1009867>.
- Pereira, T. D., Shaevitz, J. W., & Murthy, M. (2020). Quantifying behavior to understand the brain. *Nature Neuroscience*, 23(12), 1537–1549. <https://doi.org/10.1038/s41593-020-00734-z>.
- Pereira, T. D., Tabris, N., Matsliah, A., Turner, D. M., Li, J., Ravindranath, S., Papadoyannis, E. S., Normand, E., Deutsch, D. S., Wang, Z. Y., McKenzie-Smith, G. C., Mitelut, C. C., Castro, M. D., D’Uva, J., Kislin, M., Sanes, D. H., Kocher, S. D., Wang, S. S.-H., Falkner, A. L., ... Murthy, M. (2022). SLEAP: A deep learning system for multi-animal pose tracking. *Nature Methods*, 19(4), 486–495. <https://doi.org/10.1038/s41592-022-01426-1>.
- Radosavovic, I., Xiao, T., Zhang, B., Darrell, T., Malik, J., & Sreenath, K. (2024). Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 9(89), eadi9579. <https://doi.org/10.1126/scirobotics.adi9579>.
- Schwarz, R. F., Branicky, R., Grundy, L. J., Schafer, W. R., & Brown, A. E. X. (2015). Changes in Postural Syntax Characterize Sensory Modulation and Natural Variation of *C. elegans* Locomotion. *PLOS Computational Biology*, 11(8), 1–16. <https://doi.org/10.1371/journal.pcbi.1004322>.
- Seabrook, E., & Wiskott, L. (2023). A Tutorial on the Spectral Theory of Markov Chains. *Neural Computation*, 35(11), 1713–1796. https://doi.org/10.1162/neco_a_01611.
- Seifert, U. (2005). Entropy production along a stochastic trajectory and an integral fluctuation theorem. *Phys. Rev. Lett.*, 95, 040602. <https://doi.org/10.1103/PhysRevLett.95.040602>.
- Simmons-Edler, R., Badman, R. P., Berg, F. B., Chua, R., Vastola, J., Lunger, J., Qian, W., & Rajan, K. (2025). Deep RL needs deep behavior analysis: Exploring implicit planning by model-free agents in open-ended environments. *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=QD06Qv7O0P>.
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, 57(4), 193–216. <https://doi.org/10.1037/h0054367>.
- Sridhar, G., Vergassola, M., Marques, J. C., Orger, M. B., Costa, A. C., & Wyart, C. (2024). Uncovering multiscale structure in the variability of larval zebrafish navigation. *Proceedings of the National Academy of Sciences*, 121(47), e2410254121. <https://doi.org/10.1073/pnas.2410254121>.
- Stephens, G. J., Johnson-Kerner, B., Bialek, W., & Ryu, W. S. (2008). Dimensionality and Dynamics in the Behavior of *C. elegans*. *PLOS Computational Biology*, 4(4), 1–10. <https://doi.org/10.1371/journal.pcbi.1000028>.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd). MIT Press.
- Tinbergen, N. (1951). *The study of instinct*. Clarendon Press/Oxford University Press.
- Vastola, J. J., Ramdas, T., & Gershman, S. J. (2025). Storing long-lived memories via molecular error correction: A minimal mathematical model of Crick’s memory switch. *bioRxiv*. <https://doi.org/10.1101/2025.11.13.688304>.
- Villella, A., & Hall, J. C. (2008). Chapter 3 Neurogenetics of Courtship and Mating in *Drosophila*. Academic Press. [https://doi.org/https://doi.org/10.1016/S0065-2660\(08\)00603-2](https://doi.org/https://doi.org/10.1016/S0065-2660(08)00603-2).
- Wallach, H. M. (2006). Topic modeling: Beyond bag-of-words. *Proceedings of the 23rd International Conference on Machine Learning*, 977–984. <https://doi.org/10.1145/1143844.1143967>.
- Weinreb, C., Kannan, L. T., Newman-Boulle, A., Sainburg, T., Gillis, W. F., Plotnikoff, A., Makowska, S., Pearl, J. E., Osman, M. A. M., Linderman, S. W., & Datta, S. R. (2026). Spontaneous behavior is a succession of self-directed tasks. *Neuron*. <https://doi.org/https://doi.org/10.1016/j.neuron.2025.11.021>.
- Weinreb, C., Pearl, J. E., Lin, S., Osman, M. A. M., Zhang, L., Annapragada, S., Conlin, E., Hoffmann, R., Makowska, S., Gillis, W. F., Jay, M., Ye, S., Mathis, A., Mathis, M. W., Pereira, T., Linderman, S. W., & Datta, S. R. (2024). Keypoint-MoSeq: Parsing behavior by linking point tracking to pose dynamics. *Nature Methods*, 21(7), 1329–1339. <https://doi.org/10.1038/s41592-024-02318-2>.
- Wiltschko, A. B., Johnson, M. J., Iurilli, G., Peterson, R. E., Katon, J. M., Pashkovski, S. L., Abaira, V. E., Adams, R. P., & Datta, S. R. (2015). Mapping Sub-Second Structure in Mouse Behavior. *Neuron*, 88(6), 1121–1135. <https://doi.org/10.1016/j.neuron.2015.11.031>.
- Yin, G. G., & Zhang, Q. (2012). *Continuous-time markov chains and applications: A singular perturbation approach* (Vol. 37). Springer.

Appendix A. Technical assumptions restricting permissible models

The set of all possible infinitesimal stochastic matrices $\mathbf{H}$ —that is, all $N \times N$ matrices whose off-diagonal entries are nonnegative and whose columns sum to zero (Baez & Biamonte, 2018)—is quite large, and includes many transition matrices that we do not expect to be good models of behavior. In any case, regardless of their quality, our analysis does not apply to them. We want to exclude

- models which have states that it is impossible to enter (i.e., behaviors the animal will never do again);
- models which have states that it is impossible to leave (i.e., behaviors the animal will never stop doing); and
- models for which the available set of behaviors is initialization-dependent.

If we do not exclude such models, it is trivially possible to have multiple steady state distributions, or for the steady state distribution to contain zeros or ones. To avoid these possibilities, we will assume that $\mathbf{H}$ is *irreducible*. In terms of the graph associated with $\mathbf{H}$, irreducibility is equivalent to demanding that any node is reachable from any other node via some directed path through the graph. Put differently, there are not multiple isolated ‘islands’ of behavior, and there are no source or sink behaviors.

If $\mathbf{H}$ is irreducible, the Perron-Frobenius theorem applies, and tells us that there exists a unique eigenvector of $\mathbf{H}$ that (i) has eigenvalue zero, (ii) has nonnegative entries, and (iii) sums to one. Moreover, this vector (which we denote by π) is the steady state distribution, and it is always true that $\lim_{t \rightarrow \infty} \mathbf{p}(t) = \pi$ regardless of initial condition. We also assume that $\mathbf{H}$ is diagonalizable (which is not necessarily true for a generic irreducible infinitesimal stochastic matrix), but this is only to guarantee that we can sensibly talk about eigenvalues and eigenvectors.

Readers interested in a deeper mathematical understanding of these restrictions, and of Markov chains more generally, should consult some combination of Hairer (2006), Baez and Biamonte (2018), and Seabrook and Wiskott (2023). Our continuous-time formulation most closely follows that of Baez and Biamonte (2018).

Appendix B. Toy models of behavior

In this appendix, we define and compute basic properties of several toy models of behavior.

The bag model

The **bag model** assumes all possible behaviors are equally likely, and has a transition matrix

$$\mathbf{H} := \frac{k}{N} \begin{pmatrix} -(N-1) & 1 & \cdots & 1 \\ 1 & -(N-1) & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & -(N-1) \end{pmatrix} = -k \mathbf{I} + \frac{k}{N} \mathbf{1}\mathbf{1}^T. \quad (10)$$

We can compute its characteristic polynomial (and hence its eigenvalues) using the matrix determinant lemma:

$$\det \left[-(k + \lambda) \mathbf{I} + \frac{k}{N} \mathbf{1}\mathbf{1}^T \right] = [-(k + \lambda)]^N \left[1 - \frac{1}{N} \mathbf{1}^T \mathbf{1} \frac{k}{(k + \lambda)} \right] = (-1)^N (k + \lambda)^{N-1} \lambda. \quad (11)$$

Hence, there are $N - 1$ eigenvalues with $\lambda = -k$ and one with $\lambda = 0$. Any vector orthogonal to $\mathbf{1}$ is an eigenvector with eigenvalue $\lambda = -k$; the dimensionality of this space of vectors is $N - 1$.

The steady state distribution for this model is uniform, since $\mathbf{1}$ is an eigenvector with eigenvalue zero.

The cyclic model

Define the $N \times N$ **shift matrix** $\mathbf{S}$ via

$$S_{ij} = \begin{cases} 1 & \text{if } i = j + 1 \\ 1 & \text{if } i = 1 \text{ and } j = N \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

This matrix represents a cyclic shift of any vector, since

$$\mathbf{S} \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_{N-1} \\ v_N \end{pmatrix} = \begin{pmatrix} v_N \\ v_1 \\ \vdots \\ v_{N-2} \\ v_{N-1} \end{pmatrix}. \quad (13)$$

The transpose of $\mathbf{S}$ represents a cyclic shift in the other direction:

$$\mathbf{S}^{-1} \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_{N-1} \\ v_N \end{pmatrix} = \begin{pmatrix} v_2 \\ v_3 \\ \vdots \\ v_N \\ v_1 \end{pmatrix}. \quad (14)$$

This implies that $\mathbf{S}^T = \mathbf{S}^{-1}$. Define the **cyclic model** via the $N \times N$ transition matrix

$$\mathbf{H} = -(k_f + k_r) \mathbf{I} + k_f \mathbf{S} + k_r \mathbf{S}^{-1}, \quad (15)$$

where the forward rate k_f and reverse rate k_r are nonnegative constants, at least one of which is nonzero. This model abstracts the idea of transitions along a ‘ring’ of N states. If k_f and k_r are the same, then transitions are ‘diffusive’ in the sense that transitions do not tend to have any directionality. Meanwhile, if one of the rates is much larger than the other, transitions tend to take place in a fixed order. For example, if k_f is the significantly larger one, states tend to be visited in the $1 \rightarrow 2 \rightarrow \dots \rightarrow N$ order.

Eigenvectors and eigenvalues of the shift matrix. It is well known that the eigenvectors of the shift matrix correspond to Fourier modes. To see this, let $\omega := e^{\frac{2\pi i}{N}}$, and define the vector $\mathbf{v}_k$ (for $k = 0, \dots, N-1$) via

$$v_{k\ell} := \frac{1}{\sqrt{N}} \omega^{-k(\ell-1)} \quad (16)$$

for all $\ell = 1, \dots, N$. Since

$$\mathbf{S}\mathbf{v}_k = \omega^{-k(N-1)} \mathbf{v}_k = (\omega^{-1})^{-k} \mathbf{v}_k = \omega^k \mathbf{v}_k, \quad (17)$$

$\mathbf{v}_k$ is an eigenvector of $\mathbf{S}$ with eigenvalue ω^k . Since we have precisely N such vectors, these are all of the eigenvectors and eigenvalues of $\mathbf{S}$. Relatedly, observe that

$$\mathbf{S}^T \mathbf{v}_k = \omega^{-k} \mathbf{v}_k, \quad (18)$$

so the eigenvalues of $\mathbf{S}^T$ are the ‘opposite’ of those of $\mathbf{S}$ in their assignments.

Eigenvectors and eigenvalues of the transition matrix. By definition, $\mathbf{H}$ is a linear combination of $\mathbf{I}$, $\mathbf{S}$, and $\mathbf{S}^{-1}$, which we just learned share eigenvectors. Hence, the $\mathbf{v}_k$ are also eigenvectors of $\mathbf{H}$. Its eigenvalues are different:

$$\mathbf{H}\mathbf{v}_k = \left[-(k_f + k_r) + k_f \omega^k + k_r \omega^{-k} \right] \mathbf{v}_k. \quad (19)$$

Simplifying them, we find

$$\lambda_k = -(k_f + k_r) + k_f \omega^k + k_r \omega^{-k} = -(k_f + k_r) \left[1 - \cos\left(\frac{2\pi k}{N}\right) \right] + (k_f - k_r) \sin\left(\frac{2\pi k}{N}\right) i \quad (20)$$

for $k = 0, 1, \dots, N-1$. Note that when $k_f = k_r$, in which case movement along the ring is ‘diffusive’ rather than directed, the imaginary part of these eigenvalues becomes zero. On the other hand, as $|k_f - k_r|$ increases, the imaginary part grows increasingly large.

The two-state model

The **two-state model** is perhaps the simplest possible non-degenerate Markov chain. As its name implies, it has only two states and some rate of transitions between them. Define its transition matrix via

$$\mathbf{H} := \begin{pmatrix} -k_{21} & k_{12} \\ k_{21} & -k_{12} \end{pmatrix}. \quad (21)$$

Steady state distribution. The steady state distribution is

$$\pi = \begin{pmatrix} \frac{k_{12}}{k_{12} + k_{21}} \\ \frac{k_{21}}{k_{12} + k_{21}} \end{pmatrix}. \quad (22)$$

Eigenvalues and eigenvectors. A trivial calculation tells us that the nonzero eigenvalue of this model is $\lambda = -(k_{12} + k_{21})$, and that the corresponding eigenvector is

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}. \quad (23)$$

This eigenvector is particularly easy to interpret: it represents a change in behavior involving doing behavior 1 more often and behavior 2 less often. Since eigenvectors are only defined up to an overall scaling, it also represents doing behavior 1 less often and behavior 2 more often. This is somewhat trivial, since no other changes in behavior are possible for a model with only two states.

The three-state sleep-wake model

In the main text, we referenced a simple model of sleep with three states, which we will label as ‘asleep’, ‘awake’, and ‘awake and alert’. Transitions between ‘asleep’ and ‘awake’ tend to happen slowly, and transitions between ‘awake’ and ‘awake and alert’ tend to happen quickly; direct transitions between ‘asleep’ and ‘awake and alert’ are prohibited.

This model is interesting because it provides a simple and interpretable example of a model of behavior that can possess well-separated time scales. Define its transition matrix via

$$\mathbf{H} = \begin{pmatrix} -2\epsilon & \gamma & 0 \\ 2\epsilon & -(\epsilon + \gamma + k) & 2\gamma \\ 0 & \epsilon + k & -2\gamma \end{pmatrix} \quad (24)$$

where we assume $k \gg \gamma$ (i.e., k is very fast) and $\gamma \gg \epsilon$ (i.e., ϵ is very slow). Since Vastola et al. (2025) studied this model in detail, we will state their results rather than derive them from scratch.

Steady state distribution. The steady state distribution is

$$\pi = \frac{1}{(\epsilon + \gamma)^2 + k\epsilon} \begin{pmatrix} \gamma^2 \\ 2\gamma\epsilon \\ \epsilon^2 + k\epsilon \end{pmatrix}. \quad (25)$$

Eigenvalues and eigenvectors. The eigenvalues and eigenvectors of this model depend on the ‘spectral gap’

$$\Delta := \frac{\sqrt{(\epsilon + \gamma + k)^2 + 4k(\gamma - \epsilon)} - (\epsilon + \gamma + k)}{2} \geq 0. \quad (26)$$

In terms of this quantity, the nonzero eigenvalues of $\mathbf{H}$ are

$$\lambda = -(\epsilon + \gamma) + \Delta, -2(\epsilon + \gamma) - \Delta - k. \quad (27)$$

In the large k limit, these correspond to well-separated time scales, since (using $\Delta \approx \gamma - \epsilon$)

$$\lambda \approx -2\epsilon, -(3\gamma + \epsilon + k). \quad (28)$$

In particular, $\tau_1 \approx 1/(2\epsilon)$ is very long, and $\tau_2 \approx 1/k$ is very short. The corresponding eigenvectors are

$$\mathbf{v}_1 = \frac{1}{\gamma} \begin{pmatrix} \gamma \\ -\gamma + \epsilon + \Delta \\ -\epsilon - \Delta \end{pmatrix} \quad \mathbf{v}_2 = \begin{pmatrix} \frac{1}{2} \left[1 - \frac{\Delta}{\gamma - \epsilon} \right] \\ -1 \\ \frac{1}{2} \left[1 + \frac{\Delta}{\gamma - \epsilon} \right] \end{pmatrix}. \quad (29)$$

These are somewhat interpretable, especially in the large k limit. For very large k ,

$$\mathbf{v}_1 \approx \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \quad \mathbf{v}_2 \approx \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix}. \quad (30)$$

Hence, the fast mode $\mathbf{v}_2$ moves probability between the ‘awake’ and ‘awake and alert’ states only, while the slower $\mathbf{v}_1$ mode moves probability between the ‘asleep’ and ‘awake and alert’ states.

Appendix C. Hierarchical model fit to real data

Weinreb et al. (2026) fit a hierarchical hidden Markov model (using an approach they called *shMoSeq*) to open-field spontaneous mouse behavior, and recovered a transition matrix $\mathbf{P}^T$ that specifies the small-time probability of transitioning from one syllable and latent state to another. Their ‘mesoscopic’ time scale Δt was effectively 1/30 seconds, since behavior was recorded at 30 frames per second.

To convert their $\mathbf{P}^T$ into an infinitesimal stochastic $\mathbf{H}$, we moved from their row-stochastic convention (i.e., rows sum to one) to a column-stochastic convention by transposing their matrix, and then exploiting the relationship

$$\mathbf{P} = e^{\mathbf{H}\Delta t} \approx \mathbf{I} + \mathbf{H}\Delta t, \quad (31)$$

which becomes increasingly accurate as Δt gets smaller. We did this instead of deriving $\mathbf{H}$ via a matrix logarithm, partly because we empirically found that doing so led to numerical instability.

We used *only* their model labeled ‘solitary_and_social_open_field’, and not any of the other available models. This is partly because this model was simplest, since it involves only a small number of latent states, all of which have been annotated. This model has 100 behavioral syllables and 5 latent states. The latent states were assigned the labels ‘grooming’, ‘exploratory locomotion 1’, ‘exploratory locomotion 2’, ‘local investigation’, and ‘social engagement’, and we have retained those labels here.

Appendix D. Motivation for predictive dimensionality

Here, we derive the approximation to mutual information exploited in the main text to define the predictive dimensionality of a model. Let $\mathbf{P}(\tau) := e^{\mathbf{H}\tau}$ denote the matrix controlling the probability of transitioning from one state to another after an amount of time τ . The mutual information between the system's present state (assuming it is stationary) and the system's state after an amount of time τ has passed is defined as

$$I(\tau) := \sum_{a,b} P_{ba} \pi_a \log \left(\frac{P_{ba} \pi_a}{(\sum_c P_{bc} \pi_c) \pi_a} \right) = \sum_a \pi_a \sum_b P_{ba} \log \left(\frac{P_{ba}}{(\sum_c P_{bc} \pi_c)} \right) = \sum_a \pi_a \sum_b P_{ba} \log \left(\frac{P_{ba}}{\pi_b} \right), \quad (32)$$

where we used the fact that π is the steady state distribution in the last step to write $\pi_b = \sum_c P_{bc} \pi_c$. We will approximate this quantity by assuming P_{ba} is not far from steady state, so that $\delta p_{ba} := P_{ba} - \pi_b$ is small. (More precisely, we assume $\delta p_{ba}/\pi_b \ll 1$.) This assumption is useful because it allows us to Taylor expand $I(\tau)$ to second order in δp_{ba} . Doing so, we find

$$\begin{aligned} I(\tau) &= \sum_a \pi_a \sum_b P_{ba} \left(1 + \frac{\delta p_{ba}}{\pi_b} \right) \log \left(1 + \frac{\delta p_{ba}}{\pi_b} \right) \\ &\approx \sum_a \pi_a \sum_b P_{ba} \left(1 + \frac{\delta p_{ba}}{\pi_b} \right) \left(\frac{\delta p_{ba}}{\pi_b} - \frac{1}{2} \frac{\delta p_{ba}^2}{\pi_b^2} \right) \\ &\approx \sum_a \pi_a \sum_b P_{ba} \left(\frac{\delta p_{ba}}{\pi_b} + \frac{1}{2} \frac{\delta p_{ba}^2}{\pi_b^2} \right). \end{aligned} \quad (33)$$

Since $\sum_b \delta p_{ba} = \sum_b (P_{ba} - \pi_b) = 0$, this simplifies to

$$I(\tau) \approx \frac{1}{2} \sum_{a,b} \delta p_{ba}^2 \frac{\pi_a}{\pi_b} = \frac{1}{2} \sum_{a,b} (P_{ba} - \pi_b)^2 \frac{\pi_a}{\pi_b} = \frac{1}{2} \sum_{a,b} P_{ba}^2 \frac{\pi_a}{\pi_b} - \frac{1}{2}. \quad (34)$$

But this means

$$I(\tau) \approx \frac{1}{2} \text{tr} \left[\mathbf{D}^{1/2} \mathbf{P}^T \mathbf{D}^{-1} \mathbf{P} \mathbf{D}^{1/2} \right] - \frac{1}{2} = \frac{1}{2} \text{tr} \mathbf{K}(\tau) - \frac{1}{2}, \quad (35)$$

where we define the *past-future kernel* via

$$\mathbf{K}(\tau) := \mathbf{D}^{1/2} e^{\mathbf{H}^T \tau} \mathbf{D}^{-1} e^{\mathbf{H} \tau} \mathbf{D}^{1/2}. \quad (36)$$

$\mathbf{K}(\tau)$ is clearly real, symmetric, and positive definite; this implies that its eigenvalues ν_k are real and positive. Moreover, the Perron-Frobenius theorem implies that 1 is always an eigenvalue, and that all other eigenvalues are less than 1. Since 1 is an eigenvalue, we can use the $\nu_k = 1$ term to cancel the offset, so that

$$I(\tau) \approx \frac{1}{2} \sum_{\nu_k < 1} \nu_k. \quad (37)$$

This expression for $I(\tau)$ is a sum of positive eigenvalues, each of which we can interpret as ‘explaining’ part of the overall mutual information between the present and future.