Neurons as Detectors of Coherent Sets in Sensory Dynamics

Joshua L. Pughe-Sanford^{1,*} Xuehao Ding^{1,*} Jason J. Moore^{1,2,*} Anirvan M. Sengupta^{1,3} Charles Epstein¹ Philip Greengard¹ Dmitri B. Chklovskii^{1,2}

¹Center for Computational Neuroscience, Flatiron Institute, Simons Foundation, New York, NY, USA

²Neuroscience Institute, NYU Langone Medical Center, New York, NY, USA

³Physics Department, Rutgers University, New Brunswick, NJ, USA

{jpughesanford, xding, cepstein, pgreengard, mitya}@flatironinstitute.org,
jason.moore@nyulangone.org, anirvans.physics@gmail.com

* Equal contribution

Abstract

We model sensory streams as observations from high-dimensional stochastic dynamical systems and conceptualize sensory neurons as self-supervised learners of compact representations of such dynamics. From prior experience, neurons learn coherent sets—regions of stimulus state space whose trajectories evolve cohesively over finite times—and assign membership indices to new stimuli. Coherent sets are identified via spectral clustering of the stochastic Koopman operator (SKO), where the sign pattern of a subdominant singular function partitions the state space into minimally coupled regions. For multivariate Ornstein-Uhlenbeck processes, this singular function reduces to a linear projection onto the dominant singular vector of the whitened state-transition matrix. Encoding this singular vector as a receptive field enables neurons to compute membership indices via the projection sign in a biologically plausible manner. Each neuron detects either a predictive coherent set (stimuli with common futures) or a retrospective coherent set (stimuli with common pasts), suggesting a functional dichotomy among neurons. Since neurons lack access to explicit dynamical equations, the requisite singular vectors must be estimated directly from data, for example, via past-future canonical correlation analysis on lag-vector representations—an approach that naturally extends to nonlinear dynamics. This framework provides a novel account of neuronal temporal filtering, the ubiquity of rectification in neural responses, and known functional dichotomies. Coherent-set clustering thus emerges as a fundamental computation underlying sensory processing and transferable to bio-inspired artificial systems.

Neurons in early sensory areas are traditionally thought to extract from recent inputs low-dimensional latent variables that are maximally informative about the near future [1, 2, 3, 4]. Such extraction exploits statistical regularities acquired over evolutionary, developmental, and behavioral timescales from previously encountered natural stimuli [5, 6, 7]. To formalize this intuition for temporally correlated sensory stimuli, we postulate that they are generated by high-dimensional, potentially nonlinear, stochastic dynamical processes, and conceptualize neurons as self-supervised learners of *coherent sets*—regions of the stimulus state space that evolve cohesively over finite time intervals [8, 9]—thus enabling compact representations of sensory dynamics.

Coherent sets can be uncovered via spectral clustering of the *stochastic Koopman operator (SKO)*—a linear, albeit infinite-dimensional, operator that evolves observables over a finite time interval [9, 10]. The sign of the first non-trivial (subdominant) singular function of the SKO partitions state space into two minimally interacting coherent sets (Fig. 1a). Accordingly, a neuron can compute a membership index of a new input by evaluating the sign of a subdominant singular function. Because singular values and functions remain real even for irreversible dynamics, this approach generalizes metastable set detection beyond the reversible cases that eigenfunction methods require [11].

We demonstrate that for multivariate Ornstein–Uhlenbeck (OU) process [12]—a canonical example of linear stochastic dynamics and a reasonable model of summed input to a neuron [13, 14]—a subdominant singular function of the SKO corresponds to a projection of the input onto a singular vector of the whitened finite-time transition matrix, Fig. 1b. A neuron that stores this singular vector in its synaptic weights and temporal filter can compute the corresponding membership index via the sign of the weighted sum of the inputs. Although the conventional spectral clustering framework assumes stability, our singular function results extend to unstable systems by focusing on singular values closest to one. In addition to projecting on the right singular vector *predicting* near-future inputs (Fig. 1b), projecting on the left one is also possible—*retrospecting* the recent past (Fig. 1c). Because such projections require different synaptic weights and temporal filters they must be implemented by distinct neurons.

Recognizing that the underlying dynamical equations are not available to neurons, a biologically plausible detection of coherent sets requires a data-driven algorithm that can infer them directly from observations. This can be done using past-future canonical correlation analysis (CCA) [15, 10] which can be implemented locally [16]. This algorithm has the additional advantage of being applicable to nonlinear dynamics: it relies on estimating a Galerkin projection of the SKO onto a chosen functional basis via Monte Carlo integration over observed data. If the upstream neurons implement such projection, the post-synaptic neuron could then locally learn a requisite singular vector and compute a membership index via the sign of the weighted sum of the inputs.

Viewing neurons as coherent set detectors sheds light onto several longstanding neurophysiological observations. First, temporal receptive fields of neurons emerge naturally as subdominant singular vectors projecting input lag-vector representations of dynamical states. Second, the ubiquity of response rectification in neurons, exemplified by the well-known ON/OFF segregation in early visual circuits, is interpreted as a principled clustering mechanism. Finally, the theory predicts complementary neuronal classes that predict near future using predictive coherent sets or retrospect recent past using retrospective ones, consistent with known neuronal functional dichotomies such as tufted versus mitral cells in the olfactory bulb or non-lagged versus lagged cells in the lateral geniculate nucleus (LGN). Thus, the detection of coherent sets can serve as a powerful algorithmic primitive for neural computation supporting prediction and retrospection. This offers insights into biological processes and could inspire future artificial neural networks.

The remainder of the paper is organized as follows. Section 1 discusses related work. Section 2 reviews the definition of transfer operators, coherent sets and their connection to spectral clustering through the singular functions of the SKO. In Section 3, we derive the central result: under OU dynamics, the subdominant singular functions of the SKO correspond to dot products between the state and the singular vectors of the whitened transition matrix. Section 4 reviews a data-driven algorithm for identifying subdominant singular functions, which extends naturally to nonlinear dynamics and defines neuronal units rectifying positive or negative parts of the subdominant singular function. Section 5 analyzes and reviews several experimental datasets, interpreting temporal receptive fields through the hypothesis that biological neurons cluster coherent sets.

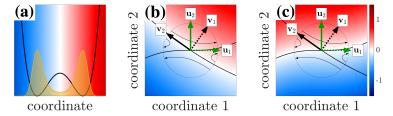


Figure 1: Subdominant SKO singular functions partition states into coherent sets (blue vs red). (a) For a one-dimensional double-well potential (black) the dominant singular function corresponds to a stationary distribution (yellow) and the subdominant singular function (red-blue) partitions the state space to minimize the leakage between the coherent sets. Left (b) and right (c) subdominant singular functions partition the state space near a 2D saddle point based on shared future and past, respectively. Black lines indicate attractive and repulsive invariant manifolds approximated by the linear subspaces in the vicinity of the saddle point (circle). Left (green) and right (black) singular vectors of the finite-time whitened transfer matrix, $[\mathbf{u}_1\mathbf{u}_2]\Lambda[\mathbf{v}_1\mathbf{v}_2]^{\top}$. As the forecast horizon grows, \mathbf{v}_1 becomes orthogonal to the stable subspace and \mathbf{u}_2 becomes orthogonal to the unstable subspace.

1 Related Work

Viewing early sensory processing as efficient or predictive coding of natural stimuli has a long tradition [5, 6, 7, 17, 2]. Closest to our work is the information bottleneck (IB) framework [1, 2, 3, 4], which for Gaussian variables reduces to past–future CCA [18]. Extending IB to dynamical systems and restricting compression to one bit corresponds to encoding the sign of the subdominant singular function of the SKO. Compared to slow feature analysis [19, 20], we incorporate an explicit dynamical-systems formulation and introduce a natural nonlinearity that enables hierarchical architectures without hand-crafted features. Predictive and retrospective neuron classes have appeared in a lattice filter model of the visual pathway [21], though that work was limited to linear processing. Relative to predictive coding models [22, 23, 24, 21], where neurons emit prediction errors, our neurons encode coherent-set memberships. While the ON/OFF division has been attributed to metabolic efficiency [25], we propose a computational explanation. Our account of lagged and non-lagged LGN cells differs from earlier models [26] by not relying on nonlinearity, suggesting that these types can emerge alongside ON/OFF segregation.

Prediction and retrospection in saddle-point dynamics have been discussed in the context of unstable periodic orbits of chaotic attractors using the dominant mode of the local SKO (and its adjoint) [27, 28], whereas we focus on a subdominant singular function defining a coherent set pair [29, 9]. Saddle point analysis based on eigenfunctions instead of singular functions predicted neuronal filters orthogonal to growing exponentials, yielding predictive neurons without considering retrospection [30].

Clustering has been previously proposed as a model of static neuronal computation on temporally uncorrelated inputs [31], capturing rectification and sparsity but not temporal receptive fields or sensory dynamics. Probabilistic coding frameworks such as the Bayesian brain hypothesis [32] have been widely studied although we are not aware of the suggestion that single neurons represent eigenfunctions or singular functions of the SKO. Self-supervised learning has been applied to visual networks [33], but these models typically omit analysis on the neuronal level. Koopman-based objectives have recently been incorporated into deep architectures to extract predictive features [34]. In contrast, we use SKO singular functions to define coherent sets [29, 9], providing a direct explanation for rectification and a principled division into predictive and retrospective neurons.

Estimation of singular functions from data originated in molecular dynamics as the VAMP framework [11]. VAMPnets [35] learn such functions with deep ReLU networks trained by backpropagation. In our model, features arise from rectified projections onto singular vectors within each neuron and can be hierarchically composed without backpropagation.

2 Discovering coherent sets via spectral clustering of transfer operators

We begin with the concept of coherent sets, which are regions of state space that tend to move as a whole under the stochastic dynamics over finite time intervals [29, 9]. Formally, a pair of sets (A, B) is coherent if a state that begins in A at time t is very likely to be found in B at time $t + \tau$:

$$\mathbb{P}[\mathbf{X}(t+\tau) \in \mathcal{B} \mid \mathbf{X}(t) \in \mathcal{A}] \approx 1,$$

where $\mathbb{P}[X|Y]$ is a probability of X given Y, $\mathbf{X}(t) \in \mathcal{X}$ denotes the random state of the system at time t. If in addition the probabilities of visiting \mathcal{A} at time t and \mathcal{B} at time $t + \tau$ are equal, then this relation works in both directions: observing the system in \mathcal{B} at time $t + \tau$ makes it likely that it was in \mathcal{A} at time t. Thus, membership in coherent sets, \mathcal{A} and \mathcal{B} , can be used to predict the future and retrospect the past, respectively, of a stochastic process.

To compute coherent sets we utilize the stochastic Koopman operator (SKO), which encodes how *observables* evolve in expectation [29, 9, 10], see the Supplement, Section 1 for a concise introduction. An observable is any scalar function of \mathbf{X} with a finite expectation value over the stochastic process. For instance, from the neuron's perspective, each synaptic input can be thought of as such an observable of stimuli. The SKO acts on an observable f

$$\mathcal{K}_{\tau}f(\mathbf{x}) = \mathbb{E}[f(\mathbf{X}(t+\tau))|\mathbf{X}(t) = \mathbf{x}],\tag{1}$$

i.e. it maps the present observable to its expected value τ units of time later. Linearity of the expectation makes \mathcal{K}_{τ} a linear operator.

Let $\mu(\mathbf{x})$ and $\nu(\mathbf{x})$ denote the probability densities of the system's state at times t and $t + \tau$, respectively [29, 9, 10]. For any measurable region $\mathcal{A} \subseteq \mathcal{X}$,

$$\mathbb{P}[\mathbf{X}(t) \in \mathcal{A}] = \int_{\mathcal{A}} \mu(\mathbf{x}) \, d\mathbf{x}, \qquad \mathbb{P}[\mathbf{X}(t+\tau) \in \mathcal{A}] = \int_{\mathcal{A}} \nu(\mathbf{x}) \, d\mathbf{x}. \tag{2}$$

These densities define inner products between functions f and g as

$$\langle f, g \rangle_{\mu} = \int_{\mathcal{X}} f(\mathbf{x})^* g(\mathbf{x}) \,\mu(\mathbf{x}) \,d\mathbf{x}, \qquad \langle f, g \rangle_{\nu} = \int_{\mathcal{X}} f(\mathbf{x})^* g(\mathbf{x}) \,\nu(\mathbf{x}) \,d\mathbf{x}.$$
 (3)

When the state space is divided into two pairs of regions, identifying the most coherent ones amounts to maximizing the following objective [9, 10]:

$$\max_{\mathcal{A},\mathcal{B}} \left\{ \mathbb{P} \left[\mathbf{X}(t+\tau) \in \mathcal{B} \, \middle| \, \mathbf{X}(t) \in \mathcal{A} \right] + \mathbb{P} \left[\mathbf{X}(t+\tau) \in \mathcal{B}^{c} \, \middle| \, \mathbf{X}(t) \in \mathcal{A}^{c} \right] \right\} \\
= \max_{\mathcal{A},\mathcal{B}} \left\{ \frac{\langle \mathbf{1}_{\mathcal{A}}, \mathcal{K}_{\tau} \mathbf{1}_{\mathcal{B}} \rangle_{\mu}}{\langle \mathbf{1}_{\mathcal{A}}, \mathbf{1}_{\mathcal{A}} \rangle_{\mu}} + \frac{\langle \mathbf{1}_{\mathcal{A}^{c}}, \mathcal{K}_{\tau} \mathbf{1}_{\mathcal{B}^{c}} \rangle_{\mu}}{\langle \mathbf{1}_{\mathcal{A}^{c}}, \mathbf{1}_{\mathcal{A}^{c}} \rangle_{\mu}} \right\}, \tag{4}$$
subject to
$$\int_{\mathcal{A}} \mu(\mathbf{x}) \, d\mathbf{x} = \mathbb{P}[\mathbf{X}(t) \in \mathcal{A}] = \mathbb{P}[\mathbf{X}(t+\tau) \in \mathcal{B}] = \int_{\mathcal{B}} \nu(\mathbf{x}) \, d\mathbf{x},$$

where $(\mathcal{A}, \mathcal{B})$ are measurable subsets of the state space \mathcal{X} , and $(\mathcal{A}^c, \mathcal{B}^c)$ denote their complements. The indicator function $\mathbf{1}_{\mathcal{A}}(\mathbf{x})$ equals 1 if $\mathbf{x} \in \mathcal{A}$ and 0 otherwise. Large values of the objective correspond to pairs of regions that remain coherent under the dynamics—that is, regions that are least dispersive over the time interval τ .

In practice, the maximizers of this quotient are well-approximated by certain singular functions of the SKO [29, 9, 10]. To see this, note that \mathcal{K}_{τ} has an adjoint operator $\mathcal{K}_{\tau}^{\dagger}$, defined so that

$$\langle f, \mathcal{K}_{\tau} g \rangle_{\mu} = \langle \mathcal{K}_{\tau}^{\dagger} f, g \rangle_{\nu}. \tag{5}$$

The finite-time forward-backward and backward-forward operators

$$\mathcal{F}_{\tau} = \mathcal{K}_{\tau} \mathcal{K}_{\tau}^{\dagger}, \qquad \qquad \mathcal{B}_{\tau} = \mathcal{K}_{\tau}^{\dagger} \mathcal{K}_{\tau} \tag{6}$$

are then self-adjoint and, under mild assumptions, compact [10, 15]. As such, each admits a countable spectral decomposition in terms of orthonormal basis

$$\mathcal{F}_{\tau} = \sum_{i=0}^{D} \lambda_i^2 \, v_i \, \langle v_i, \cdot \rangle_{\mu}, \qquad \qquad \mathcal{B}_{\tau} = \sum_{i=0}^{D} \lambda_i^2 \, u_i \, \langle u_i, \cdot \rangle_{\nu}, \qquad (7)$$

where D may be infinite. The functions $u_i(\mathbf{x})$ and $v_i(\mathbf{x})$ are the singular functions of \mathcal{K}_{τ} , satisfying

$$[\mathcal{K}_{\tau}u_i](\mathbf{x}) = \lambda_i v_i(\mathbf{x}), \qquad [\mathcal{K}_{\tau}^{\dagger}v_i](\mathbf{x}) = \lambda_i u_i(\mathbf{x}). \tag{8}$$

Pairs of singular functions with large λ_i are approximate maximizers of (4). Moreover, the sets

$$\mathcal{V}_i^{\pm} = \{ \mathbf{x} : \pm v_i(\mathbf{x}) > 0 \}, \qquad \qquad \mathcal{U}_i^{\pm} = \{ \mathbf{x} : \pm u_i(\mathbf{x}) > 0 \}, \qquad (9)$$

approximate coherent sets, where \mathcal{U}_i^{\pm} is approximately the image of \mathcal{V}_i^{\pm} under the dynamics.

The principal singular functions of the SKO are trivial, $u_0(\mathbf{x}) = v_0(\mathbf{x}) = 1$: they define the trivial coherent set, $\mathcal{U}_0 = \mathcal{V}_0 = \mathcal{X}$. The subdominant singular functions define the least dispersive non-trivial coherent set pair of the dynamics, $(\mathcal{V}_i^\pm, \mathcal{U}_i^\pm)$. We propose that the neurons compute membership indices by evaluating the signs of such singular functions. Left singular functions look forward in time predicting the future. Right singular functions look backward in time allowing to retrospect past events. Both types of measurements are important in ascertaining the state of partially observed dynamical systems, and we predict the existence of both predictive and retrospective neurons.

3 Koopman singular functions for Ornstein-Uhlenbeck processes

Identifying coherent sets through the subdominant singular functions of the SKO provides a rigorous theoretical framework, whose application requires computing these singular functions. In this section, we attempt to provide an intuitive picture by considering a linear stochastic dynamical system — OU

process — for which subdominant SKO singular functions can be found in the closed form providing much needed intuition. The temporal dynamics of neuronal input currents elicited by sensory stimuli has long been approximated by the OU process [13, 14]:

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{A}\mathbf{x}(t) + \xi(t),\tag{10}$$

where $\langle \xi(t)\xi(t')^{\top}\rangle = \mathbf{D}\delta(t-t')$. Here, we assume that \mathbf{A} is a real matrix with real eigenvalues and eigenvectors. We are particularly interested in saddle-point OU for the following reasons. If a critical point is purely repulsive, it will not be visited by the autonomous dynamics and, hence, will be physically irrelevant. Applying this approach to nearly isotropic attractive critical points gives partitions that are not 'distinguished' [9], i.e. do not represent a genuine clustering of states which could describe qualitatively different parts of the phase space.

The probability density of the state variable \mathbf{x} , $p(\mathbf{x},t)$, evolves according to the forward Kolmogorov equation (aka the Fokker-Planck equation) utilizing the forward Kolmogorov operator \mathcal{L} [36]:

$$\frac{\partial p(\mathbf{x},t)}{\partial t} = \mathcal{L} p(\mathbf{x},t) \equiv -\nabla \cdot (\mathbf{A} \mathbf{x} p(\mathbf{x},t)) + \frac{1}{2} \nabla \cdot (\mathbf{D} \nabla p(\mathbf{x},t)). \tag{11}$$

The dynamics of measurement expectation $g(\mathbf{x},t) = \mathbb{E}[g(\mathbf{X}(t))|\mathbf{X}(0) = x]$ are given by the backward Kolmogorov equation:

$$\frac{\partial g(\mathbf{x},t)}{\partial t} = \mathcal{L}^{\dagger} g(\mathbf{x},t) \equiv (\mathbf{A}\mathbf{x}) \cdot \nabla g(\mathbf{x},t) + \frac{1}{2} \nabla \cdot (\mathbf{D} \nabla g(\mathbf{x},t)), \tag{12}$$

where \mathcal{L}^{\dagger} is the adjoint of the forward Kolmogorov operator with respect to the standard Euclidean inner product, serving as the generator of the SKO:

$$\mathcal{K}_{\tau} = \exp\left(\mathcal{L}^{\dagger}\tau\right). \tag{13}$$

The stationary distribution of probability density under \mathcal{L} , satisfying $\mathcal{L}\rho_0=0$, is derived in the Supplement, Section 2:

$$\rho_0(\mathbf{x}) \sim \exp\left(-\frac{1}{2}\mathbf{x}^{\top}\mathbf{\Sigma}^{-1}\mathbf{x}\right),$$
(14)

where $\mathbf{x} \in \mathbb{R}^n$ and Σ is a solution of the Lyapunov equation,

$$\mathbf{A}\boldsymbol{\Sigma} + \boldsymbol{\Sigma}\mathbf{A}^{\top} = -\mathbf{D}.\tag{15}$$

For attractive dynamics, where all eigenvalues of \mathbf{A} have negative real parts, Eq. (15) has a unique positive-definite solution corresponding to the covariance matrix, $\mathbf{\Sigma} = \mathbb{E}\left[\mathbf{x}\mathbf{x}^{\top}\right]$. For repulsive or saddle-point dynamics, Eq. (15) is a Sylvester equation, which has a unique solution if and only if \mathbf{A} has no eigenvalues related by sign reversal [37]. Thus, for a generic choice of \mathbf{A} , Eq. (15) has a unique solution that is symmetric, real, and invertible. For repulsive directions, corresponding to eigenvalues of \mathbf{A} with positive real parts, Eq. (14) characterizes how fast the density grows away from the fixed point [38].

Under stationarity, $\mu(\mathbf{x}) = \nu(\mathbf{x}) = \rho_0(\mathbf{x})$, the adjoint of \mathcal{K}_{τ} with respect to $\rho_0(\mathbf{x})$ is given by

$$\mathcal{K}_{\tau}^{\dagger} = \operatorname{diag}(\rho_0)^{-1} \exp\left(\mathcal{L}\tau\right) \operatorname{diag}(\rho_0), \tag{16}$$

which is termed the reweighted Perron-Fronbenius operator and can also be interpreted as the SKO of the reverse-time dynamics. We then search for the eigenfunctions of the operators \mathcal{F}_{τ} and \mathcal{B}_{τ} , Eq. (6), which are well defined on short time scales.

We find (see the Supplement, Section 3) a family of eigenfunctions of \mathcal{F}_{τ} , \mathcal{B}_{τ} that can be expressed in terms of the linear projection of the state vector, \mathbf{x} :

$$\mathcal{F}_{\tau}(\mathbf{v}_{i}^{\top}\mathbf{x}) = \lambda_{i}^{2}(\mathbf{v}_{i}^{\top}\mathbf{x}), \qquad \mathcal{B}_{\tau}(\mathbf{u}_{i}^{\top}\mathbf{x}) = \lambda_{i}^{2}(\mathbf{u}_{i}^{\top}\mathbf{x}),$$

$$(17)$$

where \mathbf{u}_i and \mathbf{v}_i satisfy the following matrix eigenvector equations (see the Supplement, Section 3):

$$e^{\mathbf{A}^{\top} \tau} \Sigma^{-1} e^{\mathbf{A} \tau} \Sigma \mathbf{v}_i = \lambda_i^2 \mathbf{v}_i, \qquad \Sigma^{-1} e^{\mathbf{A} \tau} \Sigma e^{\mathbf{A}^{\top} \tau} \mathbf{u}_i = \lambda_i^2 \mathbf{u}_i.$$
 (18)

Eigenvalues of transfer operators for expanding maps could be greater than one [39]. Since ideal coherent sets with zero leakage correspond to a unity eigenvalue, in this paper a subdominant

eigenfunction refers to the non-trivial eigenfunction associated with an eigenvalue closest to one. That is, the subdominant eigenfunction for attractive (repulsive) dynamics is associated with the second largest (smallest) eigenvalue. We proved that the subdominant eigenfunctions belong to the linear family Eq. (17) for attractive and repulsive dynamics. Other eigenfunctions of \mathcal{F}_{τ} , \mathcal{B}_{τ} are either a constant, associated with the eigenvalue one, that does not change sign, or represent higher-order polynomials in \mathbf{x} , associated with less dominant eigenvalues (see the Supplement, Section 3, which, however, does not consider saddle points). Therefore, to identify the least dispersive coherent sets, it is sufficient to focus on the eigenfunctions in Eqs. (17). The pair of minimally leaking coherent sets is given by the solution of Eqs. (18) with λ_i^2 closest to unity. To obtain the coherent set membership indices, we recover indicator functions, Eq. (4), by taking the sign of these singular functions.

Eigenvectors \mathbf{v}_i and \mathbf{u}_i play a complementary role in prediction and retrospection. As the operator \mathcal{F}_{τ} propagates observables forward then backward in time, neurons projecting inputs onto \mathbf{v}_i are predictive (Fig. 1b). As the operator \mathcal{B}_{τ} propagates the observable backward then forward in time, neurons projecting inputs onto \mathbf{u}_i are retrospective (Fig. 1c). In the case of 2D saddle-point OU (Fig. 1b,c), when the shared past and future refer to other fixed points, the coherent sets of interest are the expanding coherent set for prediction (\mathbf{v}_1) and the contracting coherent set for retrospection (\mathbf{u}_2) . In Section 4 of the Supplement, we analytically prove that, as the forecast horizon goes to infinity, singular vectors \mathbf{v}_1 and \mathbf{u}_2 converge to the top and bottom left eigenvector of \mathbf{A} which are orthogonal to the stable and unstable invariant subspaces, respectively.

So far we considered fully observable OU processes. However, in reality, the OU process is often only partially observed for example via a linear projection of the state onto a scalar variable,

$$y(t) = C\mathbf{x}(t). \tag{19}$$

Neurally, such projection can be an input to a single-synapse neuron or a summed total synaptic current. Assuming observability, the state can be represented by an *n*-dimensional lag vector [40]:

$$\hat{\mathbf{x}}(t) = [y(t), y(t-1), \dots, y(t-n+1)]^{\top}, \tag{20}$$

evolving via a companion matrix, equivalent to the following auto-regressive model:

$$y(t+1) = a_1 y(t) + a_2 y(t-1) + \dots + a_n y(t-n+1) + \xi(t).$$
(21)

4 Coherent sets from data

The infinite-dimensional stochastic Koopman operator (SKO) can be approximated by a Galerkin projection onto a finite set of basis functions (or features), $\{\phi_i(\mathbf{x})\}_{i=1}^d$ [41, 10]. In the data-driven formulation, the SKO matrix representation under this basis is [15, 10] (see Supplement, Section 5):

$$\mathbf{K}_{\tau} = \mathbf{\Sigma}_{0}^{-1} \mathbf{\Sigma}_{\tau},\tag{22}$$

where Σ_0 and Σ_{τ} are the covariance matrices:

$$[\mathbf{\Sigma}_0]_{ij} \approx \frac{1}{S} \sum_{s=1}^S \phi_i(\mathbf{X}(t_s)) \phi_j(\mathbf{X}(t_s)), \qquad [\mathbf{\Sigma}_\tau]_{ij} \approx \frac{1}{S} \sum_{s=1}^S \phi_i(\mathbf{X}(t_s)) \phi_j(\mathbf{X}(t_s+\tau)), \qquad (23)$$

and describe the correlation between measurement functions i and j measured instantaneously, or with delay τ , respectively.

Such matrices are well defined if the process is stationary which may not be true for neurons. A common workaround is to assume local stationarity, that is, the distribution of inputs changes slowly compared to the timescale of the dynamics. In this case, the neuron can continuously update the estimate of the dynamics using exponential forgetting [42]. Our method can accommodate this type of non-stationarity by computing Eqs. (23) locally in time, using a temporal filter that discounts older observations. This allows the estimated covariances to track gradual changes in the underlying distribution without assuming global stationarity.

As a direct consequence of Eq. (5), the adjoint of \mathbf{K}_{τ} is $\mathbf{K}_{\tau}^{\dagger} = \mathbf{\Sigma}_{0}^{-1} \mathbf{K}_{\tau}^{\top} \mathbf{\Sigma}_{0}$. The forward and backward operators have matrix representations [15, 10]:

$$\mathbf{F}_{\tau} = \mathbf{K}_{\tau} \mathbf{\Sigma}_{0}^{-1} \mathbf{K}_{\tau}^{\top} \mathbf{\Sigma}_{0} = \mathbf{\Sigma}_{0}^{-1} \mathbf{\Sigma}_{\tau} \mathbf{\Sigma}_{0}^{-1} \mathbf{\Sigma}_{-\tau}, \quad \mathbf{B}_{\tau} = \mathbf{\Sigma}_{0}^{-1} \mathbf{K}_{\tau}^{\top} \mathbf{\Sigma}_{0} \mathbf{K}_{\tau} = \mathbf{\Sigma}_{0}^{-1} \mathbf{\Sigma}_{-\tau} \mathbf{\Sigma}_{0}^{-1} \mathbf{\Sigma}_{\tau}, \quad (24)$$

which reduce to Eq. (18) for attractive OU processes (see the Supplement, Section 6). The singular functions of the SKO within the Galerkin projection are given by

$$v_i(\mathbf{x}) = \vec{v}_i \cdot \vec{\phi}(\mathbf{x}),$$
 $u_i(\mathbf{x}) = \vec{u}_i \cdot \vec{\phi}(\mathbf{x}),$ (25)

where \vec{v}_i and \vec{u}_i are eigenvectors of the matrices \mathbf{F}_{τ} and \mathbf{B}_{τ} , respectively. Notice that \vec{v}_i and \vec{u}_i are precisely the solution of past-future CCA.

By representing these eigenvectors in synaptic weights and temporal filters, a neuron may compute a leading singular function of the SKO restricted to the linear span of its input features. By indicating the sign of the singular function, the neuron then computes the membership index for each input. For the sake of simplicity, below we only consider the temporal component of the singular vector. Even if the neuron has access to only a single synapse, it can retain a history of that synapse's activity over d consecutive time steps. This construction—known as delay embedding or delay coordinates [43]—yields also a Galerkin approximation. We take $\mathbf{x} \in \mathbb{R}^d$ to be a lag vector of the input current to the neuron over the previous d intervals of time, such that $\mathbf{X}(t) = (I(t), I(t - \Delta t), \dots, I(t - (d-1)\Delta t))$, where I(t) is the observed current at time t. In this regime, the basis functions are given by

$$\phi_i(\mathbf{x}) = x_i, \tag{26}$$

i.e. $\vec{\phi}(\mathbf{x})$ is the identity function. Singular vectors act as temporal filters over the delayed signal, to

$$r_{pre}^{\pm,i}(t) = H\left(\pm v_i(\mathbf{X}(t))\right) = H\left(\pm \sum_{k=0}^{d-1} [\vec{v}_i]_k I(t - k\Delta t)\right),$$

$$r_{ret}^{\pm,i}(t) = H\left(\pm u_i(\mathbf{X}(t))\right) = H\left(\pm \sum_{k=0}^{d-1} [\vec{u}_i]_k I(t - k\Delta t)\right),$$
(27)

where \vec{u}_i and \vec{v}_i are singular vectors (cf. Eq. (25)), and the Heaviside function, H(x) = 1 for $x \ge 0$, and H(x) = 0 for x < 0.

Therefore, sensory streams can be clustered into coherent sets via data-driven Galerkin approximations of the SKO. By learning the singular vectors from the features represented by the activity of the upstream neurons and encoding them in the synaptic weight and temporal filters, biological neurons can cluster inputs using integrate-and-fire dynamics.

5 Coherent set clustering perspective on neurophysiology

Here we apply the coherent set clustering framework to biological neurons focusing on three observations: temporal receptive fields, neuronal rectification and predictive/retrospective properties. We restrict our consideration to the processing of a scalar time series viewed as a one-dimensional projection of a multidimensional state-space dynamics, Eq. (27). Such a scalar time series could represent an input to a single-input neuron or the total synaptic current into a multi-input neuron. Here, inspired by our result for OU processes that the relevant singular functions lie in the span of a lag vector basis, we choose Eq. (26) even in the data-driven, non-linear setting. Even such simple basis choice produces tangible results.

5.1 Temporal receptive fields

We consider early visual processing where a natural stimulus can be generated by emulating the movement of the retinal image due to self-motion or saccades by scanning a natural image or its model. In turn, natural images are commonly modeled by the "dead leaves" model partitioning the space into patches of different but uniform luminance with sharp transitions between them [44]. Therefore, the resulting input time series is a set of plateaus at different levels with sharp transitions between them, Fig. 2a. Such scalar time series models input to post-photoreceptor neurons for invertebrates such as *Drosophila*, where initial processing is segregated between adjacent "pixels" or total current in a vertebrate bipolar cell.

We view this scalar time series as a linear projection of a high-dimensional dynamical system. Thus, we perform past-future CCA on the lag-vectors formed from the scalar time series. The canonical correlations reveal a spectral gap following the top two. The lag-vector space is partitioned into a pair of coherent sets in a canonical direction. The top right singular vector amounts to a

low-pass filter (Fig. 2b) similar to sustained bipolar cells of the vertebrate retina or L3 neuron in *Drosophila* [45]. The second canonical direction acts as a smoothed temporal derivative (Fig. 2b) in general agreement with experimentally reported filters of transient bipolars in vertebrates, Fig. 2c or L1 and L2 cells in *Drosophila* [22, 45]. Because the stimulus is symmetric with respect to time-reversal, the complementary left singular vectors are obtained by simply inverting the time axis (and inverting the sign for the second singular vector).

We further characterize the experimentally measured temporal responses by interpreting them from the perspective of predictive and retrospective coherent sets. To make this connection, we take advantage of the observation that a practical predictive filter must be significantly aligned with an unstable eigendirection and a practical retrospective filter must be significantly aligned with a stable eigendirection, Fig. 1b,c. Moreover, for a two-dimensional (hyperbolic) saddle point, the linear temporal filter of the predictive neuron must be orthogonal to the attractive subspace, while the linear temporal filter of the retrospective neuron must be orthogonal to the repulsive subspace, Fig. 1b,c, see Section 3 and the Supplement, Section 4. Repulsive modes in the lag-vector space are expanding exponentials in real time and attractive modes are contracting exponentials, suggesting a simple way to interpret linear temporal receptive fields by computing the cosine similarity, S, with exponentials or, equivalently, computing normalized Laplace transforms, Fig. 2c,d. The experimentally measured temporal filter is orthogonal to a contracting exponential and aligned with the expanding exponential (Fig. 2c,d), indicating a predictive neuron. For details, see the Supplement, Section 7.

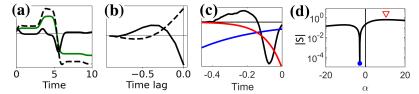


Figure 2: Data-driven temporal filters learned on the scalar time series from the natural image model compared with that of a bipolar cell. (a) Stimulus produced by the "dead leaves" model of natural images [44] (green), stimulus filtered through first (dashed black) and second (solid black) right singular vectors. (b) First (dashed) and second (solid) right singular vectors from Eq. (24) corresponding to predictive neurons. (c) Black: experimentally measured flash response of the salamander retinal bipolar cell (D.B. Kastner & S.A. Baccus, personal communication) on the inverted time axis approximates its linear filter. Compare this with with the second singular vector (solid) in b. Note that the decay of the filter at zero time delay (absent in the theoretical result without additional constraints) is a consequence of causality and continuity of the filter implementation by a biological system. The maximally aligned (red) and the orthogonal (blue) exponentials are overlaid. (d) Cosine similarities of the temporal filters with a battery of exponentials. Blue dot indicates orthogonality and red triangle indicates maximum alignment. Closer alignment with expanding exponentials and orthogonality to a contracting exponential indicates the neuron's predictive nature.

5.2 Neuronal rectification exemplified by ON and OFF cells

Our coherent set clustering perspective suggests that neurons determine membership indices by applying a Heaviside step-function to the SKO singular function (or its sign-inverse), Eq.(27). Such non-linearity could be naturally implemented by the spiking mechanism and correspond to the ON and OFF ganglion cells of the vertebrate retina, named this way because they respond to luminance increments and decrements, respectively. However, neurons post-synaptic to photoreceptors, bipolar cells in vertebrates and large monopolar cells in flies, are non-spiking. Neuronal activity is represented by continuously varying graded potentials that determine a non-negative synaptic vesicle release rate. Many such neurons thus rectify their input and are also classified as either ON or OFF cell. Such rectified response to luminance variation is graded and could be viewed as soft clustering by using not only the sign but also the magnitude of the singular function, Fig. 1, resulting in a rectified linear unit (ReLU)-like operation.

5.3 Predictive and retrospective properties of biological neurons

In this Subsection, we analyze and review large datasets of experimentally measured temporal receptive fields from the perspective of predictive and retrospective coherent sets.

Tufted and mitral cells of the mammalian olfactory bulb. In mammals, information from olfactory sensory neurons is relayed to the rest of the brain by two neuron classes: tufted cells (TCs) and mitral cells (MCs). We analyzed a dataset of the temporal receptive fields of 204 TCs and MCs recorded blindly from the rat olfactory bulb [46] from the coherent set clustering perspective. As described above (5.1, Temporal Receptive Fields), we computed the cosine similarity of experimentally measured temporal filters with a battery of growing and decaying exponentials. In addition to identifying the orthogonal exponential, we also identified the sign of the exponent which has the highest cosine similarity with the temporal filter: positive exponents correspond to predictive neurons, and negative to retrospective. As a result, we found a mixture of predictive and retrospective properties, Fig. 3. Approximately half of the recorded cells had temporal filters orthogonal to decaying exponentials and, therefore, were likely predictive. About 4% were orthogonal to growing exponentials and, hence, were likely retrospective. About one-third were orthogonal to both growing and decaying exponentials suggesting that they analyze a higher- than 2-dimensional dynamics and could be either retrospective or predictive. See the Supplement, Section 7 for details.

We speculate that TCs and MCs are mostly predictive and retrospective, respectively, based on the existing literature. First, they respond at different phases of the sniff cycle as monitored by air flow [47]: TCs are preferentially active during the exponential growth phase and silent during the exponential decay phase. Conversely, MCs are preferentially active during the exponential decay phase and silent during the exponential growth phase. Second, MC responses lag relative to TCs suggesting they play a different role [47, 48, 49, 50], but see [51, 52]. Finally, TCs and MCs differ in their responses to synaptic inputs from the electrophysiologically stimulated olfactory sensory neurons [53]: TCs receive direct inputs and respond with an immediate biphasic profile typical of prediction, while MCs are activated indirectly with a delay characteristic of retrospection.

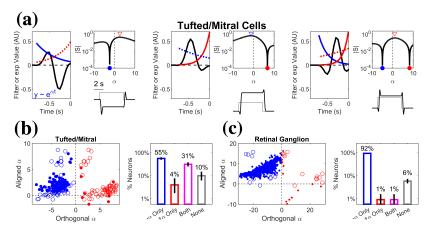


Figure 3: Temporal receptive fields of olfactory bulb and retinal neurons interpreted through the lens of coherent sets exhibit predictive and retrospective properties. (a) Linear temporal filters (black) from three rat olfactory bulb tufted/mitral neurons [46]. Solid colored lines are orthogonal exponentials; dotted colored lines are aligned exponentials. Top insets show the cosine similarity, S, of the temporal filter with exponentials. Orthogonal and aligned values are marked by dots and triangles, respectively. Bottom insets show convolution of the filter with a step pulse. The first neuron is orthogonal to a decaying and aligned with expanding exponentials and interpreted as predictive. The second neuron is orthogonal to an expanding and aligned with decaying exponentials and interpreted as retrospective. The third neuron is orthogonal to both a decaying and expanding exponential but aligned to an expanding one. (b) Left, for each tufted/mitral cell, the α value yielding an orthogonal exponential is plotted against the α that is maximally aligned. Filled dots indicate neurons orthogonal to a single α . Open dots indicate neurons orthogonal to two α values. Right, The proportion of cells orthogonal to a negative exponential only, a positive exponential only, both, or neither is shown in the bar graph. 95% confidence intervals of the mean obtained through a binomial fit are plotted as black lines. (c) Same as **b** but for retinal ganglion cells (RGCs). See Fig. S1 for example temporal receptive fields. Note that very few RGCs are orthogonal to expanding exponentials.

Retinal ganglion cells (RGCs). To explore whether the distributions of predictive and retrospective neurons vary across sensory modalities, we analyzed a large dataset of temporal receptive fields of 1345 RGCs from dissociated vertebrate retina [54], Figs. S1, 3C. In vertebrates, signals from

differentially stimulated photoreceptors are combined in the retina and, therefore, an RGC's receptive field has a spatial component. Here, we only focus on the temporal component obtained by a rank-1 decomposition of the spatio-temporal receptive field. The majority of RGCs' receptive fields were orthogonal to decaying exponentials and aligned with expanding exponentials, hence predictive, though 1% of cells were retrospective, Fig. 3C. See the Supplement, Section 7 for details.

Lagged and non-lagged cells in the Lateral Geniculate Nucleus. RGCs project to the Lateral Geniculate Nucleus (LGN) [55], in which two classes of relay cells have been described in cat [56, 57, 58], monkey [59], and mouse [60]: one with the temporal receptive fields of predictive cells and one with retrospective properties, Fig. S2. Because the firing of the latter class of neurons substantially lags a step stimulus onset, they were termed "lagged cells". Comparing the response profiles of our dataset of tufted/mitral cells to a step stimulus provides insight into how lagged cells may manifest, Fig. 3A. Note that the middle, retrospective, neuron is initially inhibited by the positive step response, but overcomes this inhibition due to the delayed positive filter lobe. Thus, we identify non-lagged and lagged neurons with predictive and retrospective units, respectively.

6 Discussion

We propose that neurons act as *coherent–set detectors*: each unit computes a linear projection of its input onto a singular vector of a whitened finite–time transition matrix and then takes the sign or rectifies its positive or negative part. We derive this mechanism analytically for saddle–point Ornstein–Uhlenbeck (OU) dynamics and suggest that it extends to nonlinear systems via Galerkin projection on a chosen feature basis. In practice, a leading singular function can be learned directly from data, providing a path toward biologically plausible implementations (e.g., local weight updates). This perspective captures several physiological regularities, including rectification and selective sensitivity to temporally extended (finite–time) structure in stimuli as well as predictive and retrospective neuron types.

Open issues:

- (i) Unstable modes and reference measures. Section 3 characterizes SKO singular functions for the OU model, but linking them to the spectral clustering framework of Section 2 under instability requires a reference measure that decays away from the saddle neighborhood. Practically, this suggests finite—time localization (windowing), exponential discounting, or committor/Doob—transformed weights so that the inner product prioritizes trajectories that remain near the saddle. Interpreting the OU model as a local description then would render the spectral objects well posed on that neighborhood.
- (ii) Data-driven estimation along unstable directions. The Galerkin procedure of Section 4 matches the OU analysis for stable modes; extending it to unstable modes likely requires modified forward-backward compositions over short horizons so that finite-time expansion (singular values (>1)) is preserved rather than suppressed by whitening. Establishing a rigorous equivalence in this regime remains open.
- (iii) Beyond 2D saddles. The correspondence between finite-time singular vectors and eigenvectors shown for 2D saddles may not extend to higher dimensions, especially with multiple unstable directions. Systematic analysis in higher-dimensional settings is needed.
- (iv) Circuit feedback loops. While many early sensory pathways are predominantly feedforward and fall within our framework, feedback is ubiquitous in deeper circuits. Incorporating explicit closed—loop interactions—for example, via the controller perspective of [61]—could unify coherent—set detection with action selection and behavior generation.
- (v) Intra-neuronal feedback. In our current account, rectification follows linear projection and does not influence learning of synaptic weights or temporal filters. In reality, synaptic plasticity may depend on rectified output via dendritic backpropagation. Incorporating such output into projection learning could provide the needed windowing for local learning mechanisms.

Societal impact. By linking biological neural computation to finite-time dynamical structure, this work advances our understanding of brain function and may inform approaches to mental health and neurological disorders. It could also inspire biologically grounded, self-supervised architectures for artificial neural networks operating in dynamical settings.

Acknowledgments. We are grateful to T. Ahamed, R. Cosentino, A. Genkin, B. Gutkin, J. M. Heninger, P. Kidd, A. Koulakov, I. Nemenman, S.E. Palmer, S. Qin, D. Rinberg, E. Schneidman, D.

Schwab, E. Simoncelli, S. Solla, P. Thomas and, especially, P. Cvitanovic for insightful discussions, to P. Gupta, D.F. Albeanu, D.B. Kastner, and S.A. Baccus for sharing their datasets. This research was supported in part by the National Science Foundation (NSF) EFRI BRAID: Scalable-Learning Neuromorphics, award No. 2318152. Some work was initiated and performed at the Kavli Institute for Theoretical Physics (KITP) supported by grant NSF PHY-2309135 and at the Aspen Center for Physics, which is supported by NSF grant PHY-2210452.

References

- [1] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- [2] William Bialek, Rob R De Ruyter Van Steveninck, and Naftali Tishby. Efficient representation as a design principle for neural coding and computation. In 2006 IEEE international symposium on information theory, pages 659–663. IEEE, 2006.
- [3] Stephanie E Palmer, Olivier Marre, Michael J Berry, and William Bialek. Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22):6908–6913, 2015.
- [4] Matthew Chalk, Olivier Marre, and Gašper Tkačik. Toward a unified theory of efficient, predictive, and sparse coding. *Proceedings of the National Academy of Sciences*, 115(1):186–191, 2018.
- [5] Fred Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183, 1954.
- [6] Horace B Barlow et al. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1(01):217–233, 1961.
- [7] Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001.
- [8] Michael Dellnitz and Oliver Junge. On the approximation of complicated dynamical behavior. *SIAM Journal on Numerical Analysis*, 36(2):491–515, 1999.
- [9] Gary Froyland. An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems. *Physica D: Nonlinear Phenomena*, 250:1–19, 2013.
- [10] Stefan Klus and Nataša Djurdjevac Conrad. Dynamical systems and complex networks: A koopman operator perspective. *Journal of Physics: Complexity*, 5(4):041001, 2024.
- [11] Hao Wu and Frank Noé. Variational approach for learning markov processes from time series data. *Journal of Nonlinear Science*, 30(1):23–66, 2020.
- [12] Wikipedia contributors. Ornstein–uhlenbeck process Wikipedia, the free encyclopedia, 2025. [Online; accessed 15-May-2025].
- [13] Henry C Tuckwell. Stochastic partial differential equations in neurobiology: linear and nonlinear models for spiking neurons. In *Stochastic biomathematical models: with applications to neuronal modeling*, pages 149–173. Springer, 2012.
- [14] Alain Destexhe, M Rudolph, J M Fellous, and T J Sejnowski. Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24, 2001.
- [15] Péter Koltai, Hao Wu, Frank Noé, and Christof Schütte. Optimal data-driven estimation of generalized markov state models for non-equilibrium dynamics. *Computation*, 6(1):22, February 2018.
- [16] David Lipshutz, Yanis Bahroun, Siavash Golkar, Anirvan M Sengupta, Dmitri B Chklovskii, and others many. A biologically plausible neural network for multichannel canonical correlation analysis. *Neural Computation*, 33(9):2309–2352, 2021.

- [17] Byron H Price and Jeffrey P Gavornik. Efficient temporal coding in the early visual system: existing evidence and future directions. *Frontiers in Computational Neuroscience*, 16:929348, 2022.
- [18] Gal Chechik, Amir Globerson, Naftali Tishby, and Yair Weiss. Information bottleneck for gaussian variables. *Advances in Neural Information Processing Systems*, 16, 2003.
- [19] Laurenz Wiskott and Terrence J Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.
- [20] David Lipshutz, Charles Windolf, Siavash Golkar, and Dmitri B Chklovskii. A biologically plausible neural network for slow feature analysis. *Advances in Neural Information Processing Systems*, 33:14986–14996, 2020.
- [21] Karol Gregor and Dmitri Chklovskii. A lattice filter model of the visual pathway. *Advances in Neural Information Processing Systems*, 25, 2012.
- [22] Mandyam Veerambudi Srinivasan, Simon Barry Laughlin, and Andreas Dubs. Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1205):427–459, 1982.
- [23] Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.
- [24] Shaul Druckmann, Tao Hu, and Dmitri Chklovskii. A mechanistic model of early sensory processing based on subtracting sparse representations. *Advances in Neural Information Processing Systems*, 25, 2012.
- [25] Julijana Gjorgjieva, Haim Sompolinsky, and Markus Meister. Benefits of pathway splitting in sensory coding. *Journal of Neuroscience*, 34(36):12127–12144, 2014.
- [26] Dawei W Dong and Joseph J Atick. Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in Neural Systems*, 6(2):159– 178, 1995.
- [27] Predrag Cvitanović and Domenico Lippolis. Knowing when to stop: How noise frees us from determinism. In *AIP Conference Proceedings*, volume 1468, pages 82–126. American Institute of Physics, 2012.
- [28] Jeffrey M Heninger, Domenico Lippolis, and Predrag Cvitanović. Neighborhoods of periodic orbits and the stationary distribution of a noisy chaotic system. *Physical Review E*, 92(6):062922, 2015.
- [29] Gary Froyland, Naratip Santitissadeekorn, and Adam Monahan. Transport in time-dependent dynamical systems: Finite-time coherent sets. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(4), 2010.
- [30] Siavash Golkar, Jules Berman, David Lipshutz, Robert Mihai Haret, Tim Gollisch, and Dmitri B Chklovskii. Neuronal temporal filters as normal mode extractors. *Physical Review Research*, 6(1):013111, 2024.
- [31] Cengiz Pehlevan, Alexander Genkin, and Dmitri B Chklovskii. A clustering neural network model of insect olfaction. In 2017 51st Asilomar Conference on Signals, Systems, and Computers, page 593–600. IEEE, IEEE, 2017.
- [32] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11):1432–1438, 2006.
- [33] Yosef Singer, Yayoi Teramoto, Ben DB Willmore, Jan WH Schnupp, Andrew J King, and Nicol S Harper. Sensory cortex is optimized for prediction of future input. *Elife*, 7:e31557, 2018.

- [34] Jinho Choi, Sivaram Krishnan, and Jihong Park. Koopman autoencoder via singular value decomposition for data-driven long-term prediction. In 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP), pages 1–6. IEEE, 2024.
- [35] Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. Vampnets for deep learning of molecular kinetics. *Nature communications*, 9(1):5, 2018.
- [36] Grigorios A Pavliotis. Stochastic processes and applications. Texts in applied mathematics, 60, 2014.
- [37] Rajendra Bhatia and Peter Rosenthal. How and why to solve the operator equation ax- xb= y. *Bulletin of the London Mathematical Society*, 29(1):1–21, 1997.
- [38] Chulan Kwon, Ping Ao, and David J Thouless. Structure of stochastic dynamics near fixed points. Proceedings of the National Academy of Sciences, 102(37):13029–13033, 2005.
- [39] Oscar F Bandtlow, Wolfram Just, and Julia Slipantschuk. Spectral structure of transfer operators for expanding circle maps. In *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, volume 34, pages 31–43. Elsevier, 2017.
- [40] Tohru Katayama. Subspace methods for system identification. Springer, 2005.
- [41] Matthew O. Williams, Clarence W. Rowley, and Ioannis G. Kevrekidis. A kernel-based method for data-driven koopman spectral analysis. *Journal of Computational Dynamics*, 2(2):247–265, 2015.
- [42] Bernard Widrow and Samuel D Stearns. Adaptive Signal Processing. Prentice-Hall, Englewood Cliffs, NJ, 1985.
- [43] Hassan Arbabi and Igor Mezić. Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the koopman operator. *SIAM Journal on Applied Dynamical Systems*, 16(4):2096–2126, January 2017.
- [44] D Jeulin. Dead leaves models: from space tesselation to random functions. In Symposium on the Advances in the Theory and Applications of Random Sets, pages 137–156. World Scientific, 1996
- [45] Madhura D Ketkar, Burak Gür, Sebastian Molina-Obando, Maria Ioannidou, Carlotta Martelli, and Marion Silies. First-order visual interneurons distribute distinct contrast and luminance information across on and off pathways to achieve stable behavior. *Elife*, 11:e74937, 2022.
- [46] Priyanka Gupta, Dinu F Albeanu, and Upinder S Bhalla. Olfactory bulb coding of odors, mixtures and sniffs is a linear sum of odor time profiles. *Nature neuroscience*, 18(2):272–281, 2015.
- [47] Izumi Fukunaga, Manuel Berning, Mihaly Kollo, Anja Schmaltz, and Andreas T Schaefer. Two distinct channels of olfactory bulb output. *Neuron*, 75(2):320–329, 2012.
- [48] Nixon M Abraham, Hartwig Spors, Alan Carleton, Troy W Margrie, Thomas Kuner, and Andreas T Schaefer. Maintaining accuracy at the expense of speed: stimulus similarity defines odor discrimination time in mice. *neuron*, 44(5):865–876, 2004.
- [49] Naoshige Uchida and Zachary F Mainen. Speed and accuracy of olfactory discrimination in the rat. *Nature neuroscience*, 6(11):1224–1229, 2003.
- [50] Dmitry Rinberg, Alexei Koulakov, and Alan Gelperin. Speed-accuracy tradeoff in olfaction. Neuron, 51(3):351–358, 2006.
- [51] Marta Díaz-Quesada, Isaac A Youngstrom, Yusuke Tsuno, Kyle R Hansen, Michael N Economo, and Matt Wachowiak. Inhalation frequency controls reformatting of mitral/tufted cell odor representations in the olfactory bulb. *Journal of Neuroscience*, 38(9):2189–2206, 2018.
- [52] Shaina M Short and Matt Wachowiak. Temporal dynamics of inhalation-linked activity across defined subpopulations of mouse olfactory bulb neurons imaged in vivo. *eneuro*, 6(3), 2019.

- [53] David H Gire, Kevin M Franks, Joseph D Zak, Kenji F Tanaka, Jennifer D Whitesell, Abigail A Mulligan, René Hen, and Nathan E Schoppa. Mitral cells in the olfactory bulb are mainly excited through a multistep signaling path. *Journal of Neuroscience*, 32(9):2964–2975, 2012.
- [54] David B Kastner and Stephen A Baccus. Coordinated dynamic encoding in the retina using opposing forms of plasticity. *Nature Neuroscience*, 14(10):1317–1322, 2011.
- [55] Daqing Cai, Gregory C Deangelis, and Ralph D Freeman. Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *Journal of neurophysiology*, 78(2):1045–1061, 1997.
- [56] David N Mastronarde. Two classes of single-input x-cells in cat lateral geniculate nucleus. i. receptive-field properties and classification of cells. *Journal of neurophysiology*, 57(2):357–380, 1987.
- [57] Jonathan Wolfe and Larry A Palmer. Temporal diversity in the lateral geniculate nucleus of cat. *Visual neuroscience*, 15(4):653–675, 1998.
- [58] Alan B Saul and AL Humphrey. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *Journal of neurophysiology*, 64(1):206–224, 1990.
- [59] Alan B Saul. Lagged cells in alert monkey lateral geniculate nucleus. *Visual neuroscience*, 25(5-6):647–659, 2008.
- [60] Denise M Piscopo, Rana N El-Danaf, Andrew D Huberman, and Cristopher M Niell. Diverse visual features encoded in mouse lateral geniculate nucleus. *Journal of Neuroscience*, 33(11):4642–4656, 2013.
- [61] Jason J Moore, Alexander Genkin, Magnus Tournoy, Joshua L Pughe-Sanford, Rob R de Ruyter van Steveninck, and Dmitri B Chklovskii. The neuron as a direct data-driven controller. *Proceedings of the National Academy of Sciences*, 121(27):e2311893121, 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction closely and accurately describe the content of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: While conceptualizing neurons through coherent set clustering offers a novel perspective on neural computation, this framework does not yet account for the role of feedback. In photoreceptor neurons, rectification output appears not to influence synaptic learning, but in other neuron types, learning may depend critically on dendritic backpropagation. Furthermore, the widespread presence of feedback loops in brain circuits highlights the need to incorporate such mechanisms into the broader dynamical systems framework.

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was
 only tested on a few datasets or with a few runs. In general, empirical results often
 depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach.
 For example, a facial recognition algorithm may perform poorly when image resolution
 is low or images are taken in low lighting. Or a speech-to-text system might not be
 used reliably to provide closed captions for online lectures because it fails to handle
 technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Proofs are provided in Section 3 and the methods.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented
 by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Simulation parameters and pseudocode for Figure 2 are provided in the supplemental material. The procedure to arrive at results in Figure 3 is detailed in the supplemental material.

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility.

In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: The simulations and analyses performed can be described with words, and are described in the supplemental material.

Guidelines

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be
 possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not
 including code, unless this is central to the contribution (e.g., for a new open-source
 benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new
 proposed method and baselines. If only a subset of experiments are reproducible, they
 should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: No training or testing was performed.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail
 that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars in Figure 3C indicate 95% confidence intervals of the mean, obtained through a binomial fit. No statistical tests were performed.

Guidelines:

The answer NA means that the paper does not include experiments.

- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how
 they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The modest compute requirements are described in the supplemental material. Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Our work conforms in every respect with the NeurIPS Code of Ethics

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Discussed at the end of the paper. The work is foundational research that can facilitate our understanding of the brain and facilitate designing biologically inspired neural networks.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal
 impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Datasets from two previously published studies were used to create Figure 3. These papers are appropriately cited, but the datasets were obtained directly from their creators, and are not publicly licensed assets.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
 - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
 - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

705

706

707

708

709

710

711

712

713

714

715

716

717

718

720

721 722

723

724

725

726

727

728

729

730

731

732

733

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not crowdsource nor use human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not crowdsource nor use human subjects.

Guidelines:

 The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

734

735

736

737

738

739

740

741

742

743

745

746

747

748 749

750

751

753

Justification: The core development in this research does not involve LLMs in any aspect.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.