

The Neuron as a Direct Data-Driven Controller

Jason Moore^a, Alexander Genkin^b, Magnus Tournoy^b, Joshua Pughe-Sanford^b, Rob R. de Ruyter van Steveninck^c, and Dmitri B. Chklovskii^{a,b,1}

This manuscript was compiled on January 4, 2024

In the quest to model neuronal function amidst gaps in physiological data, a promising strategy is to develop a normative theory that interprets neuronal physiology as optimizing a computational objective. This study extends the current normative models, which primarily optimize prediction, by conceptualizing neurons as optimal feedback controllers. We posit that neurons, especially those beyond early sensory areas, act as controllers, steering their environment towards a specific desired state through their output. This environment comprises both synaptically interlinked neurons and external motor sensory feedback loops, enabling neurons to evaluate the effectiveness of their control via synaptic feedback. Utilizing the novel Direct Data-Driven Control (DD-DC) framework, we model neurons as biologically feasible controllers which implicitly identify loop dynamics, infer latent states and optimize control. Our DD-DC neuron model explains various neurophysiological phenomena: the shift from potentiation to depression in Spike-Timing-Dependent Plasticity (STDP) with its asymmetry, the duration and adaptive nature of feedforward and feedback neuronal filters, the imprecision in spike generation under constant stimulation, and the characteristic operational variability and noise in the brain. Our model presents a significant departure from the traditional, feedforward, instant-response McCulloch-Pitts-Rosenblatt neuron, offering a novel and biologically-informed fundamental unit for constructing neural networks.

Neuron | Control | Dynamics | ...

Despite the wealth of mechanistic insights into neuronal physiology, constructing generalizable models of brain function remains a formidable challenge in neuroscience. This difficulty largely stems from the inherent variability of biological neurons, characterized by an array of challenging-to-quantify parameters like ion channel densities. A promising strategy to overcome this challenge involves developing a normative theory of neuronal function, conceptualizing neuronal physiology as an optimization of a computational objective. Such a normative theory can potentially mitigate the limitations posed by scarce physiological data through a focus on the functional integrity of computational models.

Shining examples of such a normative approach are the efficient coding and predictive information theories. Efficient coding (1–7), by maximizing transmitted information under physical constraints, views spike-triggered averages (STAs) as optimal feedforward filters and rationalizes their adaptation with input statistics. Predictive information theories (8–12), by optimizing the encoding of future-relevant information, have demonstrated quantitative congruence with experimental observations in early sensory areas. These theories apply beyond these areas, as evidenced by the adaptive nature of feedforward filters in other neuronal types (13, 14).

However, this perspective does not fully account for certain physiological attributes of neurons. Our analysis reveals that neurons adapt not only their feedforward filters but also their spike-history-dependent (feedback) filters, suggesting a functional role beyond basic housekeeping operations like sodium channel deactivation during refractory periods. Furthermore, whereas current injections in neurons with identical high-variance waveforms produce consistent spike trains, constant current injections result in more variable outputs (14). Neither feedback filter adaptation, nor

inconsistent response to constant current injections are predicted by efficient coding.

While prediction remains a crucial aspect of neuronal computation beyond early sensory areas, it likely isn't the sole computational objective. Neurons, particularly in motor and pre-motor areas, are tasked with not only forecasting but also influencing future states of the external environment through precise control signals. Additionally, the pervasive presence of feedback loops in the brain (15–18) underscores that neuronal outputs often modulate their own inputs physiologically.

These observations have led us to expand the predictive neuron model, incorporating optimal feedback control into the normative framework. We posit that neurons, especially those beyond early sensory areas, act as feedback controllers, aiming to steer their environment toward a desired state, as depicted in Fig. 1A. The neuronal environment encompasses both the circuits of interconnected neurons and external motor sensory loops, allowing the neuron to assess control efficacy through synaptic feedback.

At first glance, the task of being a feedback controller may seem daunting for a neuron. To begin with, the dynamics of its environment are not known to the neuron a priori, necessitating learning them from data. Traditional system identification methods tackle this by deducing dynamic parameters (e.g., parameters **A**, **b**, and **C** in linear state-space models, as illustrated in Fig. 1A) from historical observations and control signals (19). These parameters form the basis for deriving a control law that optimizes specific objectives, like optimal or robust control (20). When dealing with low-dimensional or noisy observations, the control law needs to be based not on immediate observations, \mathbf{y} , only but on an estimated state, $\hat{\mathbf{x}}$, (20) derived from past data—a process known as output

^aNeuroscience Institute, NYU Langone Medical School; ^bCenter for Computational Neuroscience, Flatiron Institute; ^cPhysics Department, Indiana University DC designed research; DC, JM, AG, MT, JPS, performed research, JM, AG analyzed data, RRRS contributed data; DC, JM wrote the paper

¹To whom correspondence should be addressed. E-mail: chklovskii@gmail.com

Authors declare no competing interests.

feedback control and Kalman filtering (Fig. 1A). Although, for linear dynamics, the above tasks have known solutions (19, 20), they are computationally too demanding for a single neuron to perform or even to represent the dynamic parameters explicitly.

To implement a biologically plausible feedback controller, we adopt the novel DD-DC framework (21–23). The crux of DD-DC is to sidestep the explicit representation of the controlled dynamical system and the explicit inference of the latent state, instead directly mapping observations to control signals. This mapping is learned from historical pairings of observations and control signals. In scenarios where the remainder of the loop is represented by a linear dynamical system of order $n \geq 1$, with scalar input (control signal) and output (observations), this relationship is characterized by an Auto-Regressive Moving Average (ARMA) process, as depicted in Fig. 3A.

Conceptualizing neurons as controllers in general and modeling them as DD-DCs in particular provides insights into multiple seemingly unrelated experimental observations. Firstly, it can explain the potentiation/depression transition in Spike-Timing Dependent Plasticity (STDP) and its asymmetry (Fig. 1D) (24–26). Secondly, it can account for the temporal extent (non-instantaneous nature) of feed-forward (STA) and feedback (spike-history dependent) filters and their adaptation to input statistics (Fig. 3 B-F). Thirdly, it explains the loss of temporal precision in the neuronal spike-generation mechanism under constant input (Fig. 4, Right) (14). Fourthly, the operation of DD-DC in the online setting requires variability and/or noise, which is consistent with many neurophysiological observations (27–33). Finally, viewing neurons as controllers is consistent with the observations of movement related activity throughout much of the brain including traditionally sensory areas (34, 35).

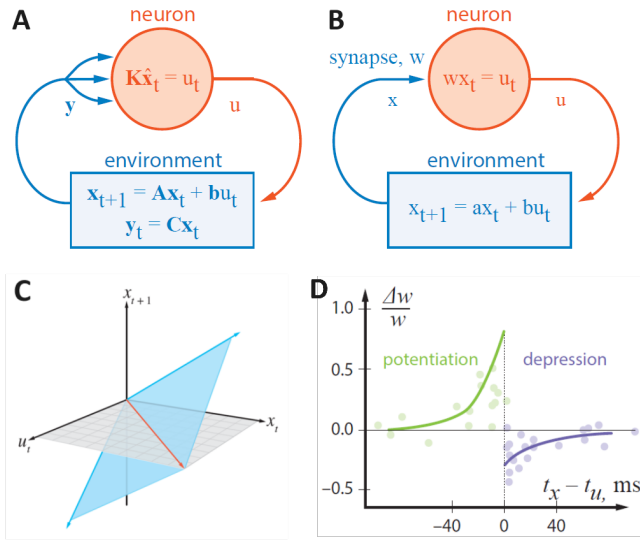


Fig. 1. A: A schematic representation of a neuron as a feedback controller in a closed loop. **B:** A scalar fully observed dynamical system controlled by tuning the weight of a synapse, w , in the control law. **C:** The subspace of valid pairings of observations and controls (blue plane) is spanned by the previously observed states (blue vectors). The intersection of the valid dynamical subspace with the $x_{t+1} = 0$ plane defines the control law (red line). **D:** Spike-timing dependent plasticity (STDP): the relative change in the synaptic weight, $\Delta w/w$, vs. the time interval between the pre- and post-synaptic spikes, $t_x - t_u$, showing the potentiation (causal) and depression (anti-causal) windows (24–26).

Our model applies not only to neuroscience but also to machine learning and artificial intelligence. Current artificial neural networks

are typically based on a neuronal unit inspired by an outdated view of neurons (36, 37). This neuronal unit is overly simplistic in that it lacks internal feedback and temporal dynamics (for more details, see below). Therefore, our proposed DD-DC model of a neuron could serve as an alternative foundational building block for constructing biologically-inspired artificial neural networks.

The Direct Data-Driven Control (DD-DC) framework

In this section we provide an overview of the DD-DC framework (21–23). In our exposition, we use lowercase letters to denote scalar variables, lowercase boldface for column vector variables, and uppercase boldface for matrices and row vectors.

For the sake of clarity, we model the neuronal environment as a linear dynamical system in a discrete-time state-space representation (Fig. 1A):

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{b}u_t, \quad [1]$$

$$\mathbf{y}_t = \mathbf{C}\mathbf{x}_t, \quad [2]$$

where $\mathbf{x}_t \in \mathbb{R}^n$ represents the latent state of the environment at time t , u_t is the control signal from the neuron, and \mathbf{y}_t is the neuron's observation. In the realm of model-based control, the dynamics parameters \mathbf{A} , \mathbf{b} , and \mathbf{C} are typically predefined, an unrealistic presumption in a biological context. We consider the system Eq. (1) to be fully controllable and observable. In some cases, the optimal control signal is linearly related to the estimated latent state variable, $\hat{\mathbf{x}}$,

$$u_t = \mathbf{K}\hat{\mathbf{x}}_t. \quad [3]$$

The DD-DC was developed for scenarios where the parameters \mathbf{A} , \mathbf{b} , and \mathbf{C} are unknown to the controller. It generates control signals directly from the observations, bypassing an explicit representation of the system dynamics and the latent state. This mapping is learned from historical pairings of observations and control signals. The intuition behind DD-DC is that any valid observation-control pairing belongs to the subspace delineated by Eq. (1) spanned by historical pairings (Fig. 1C). This intuition is formalized by Willems' fundamental lemma (21), which posits that each observation-control pairing can be expressed as a linear combination of k historical pairings ($1, 2, \dots, k < t$),

$$\begin{bmatrix} \hat{\mathbf{x}}_{t+1} \\ \hat{\mathbf{x}}_t \\ u_t \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{x}}_{1+1} & \dots & \hat{\mathbf{x}}_{k+1} \\ \hat{\mathbf{x}}_1 & \dots & \hat{\mathbf{x}}_k \\ u_1 & \dots & u_k \end{bmatrix} \mathbf{g}, \quad \mathbf{g} \in \mathbb{R}^k, \quad [4]$$

assuming the matrix composed of $\hat{\mathbf{x}}$ and u rows on the right has full row-rank (linear independence), a condition known as persistent excitation.

The DD-DC computes the control signal directly from historical pairings by solving Eq. (4) for u_t , thereby obviating the need for an explicit representation of the dynamical system (\mathbf{A} , \mathbf{b} , and \mathbf{C}) and the latent state, $\hat{\mathbf{x}}$, in the controller. Initially formulated for ideal, noise-free, linear dynamics in offline settings with extensive datasets (21), DD-DC has recently been expanded to accommodate noisy observations, non-linear dynamics, online applications, and limited datasets (23, 38–41) making it a potent model of computation in biological neurons. In the following sections, we explore the implications of this hypothesis and demonstrate its alignment with existing experimental evidence and novel analysis.

Basic DD-DC Accounts for Spike-Timing-Dependent Plasticity (STDP)

In this section, we demonstrate how even the most basic DD-DC neuron model can account for the principal characteristics of STDP (Fig. 1D). We start with the assumption that all dynamical variables are scalar ($n = 1$) and that the system is fully observed, thus $\hat{x} = y = x$ (Fig. 1B). This allows us to express equation Eq. (1) in a scalar form:

$$x_{t+1} = ax_t + bu_t. \quad [5]$$

We posit that the neuron aims to stabilize the environment's state at $x = 0$, even when its dynamics are unstable ($a > 1$). To achieve this, we employ a one-step time-horizon Linear Quadratic Regulator (LQR), where the optimal control signal u_t^* minimizes the sum of squared state error and control energy:

$$u_t^* = \arg \min_{u_t} q \|x_{t+1}\|^2 + r \|u_t\|^2. \quad [6]$$

This LQR objective is fulfilled by a linear control law:

$$u_t^* = w^* x_t, \quad [7]$$

with w^* , a scalar, representing the synaptic weight in place of \mathbf{K} from Eq. (3), as shown in Fig. 1B.

For simplicity, we initially address the limiting case of LQR with zero control cost ($r = 0$), and later present the solution for nonzero r . In the $r = 0$ scenario, Eq. (6) is minimized by $x_{t+1} = 0$. By substituting Eq. (7) into Eq. (5) and ensuring $x_{t+1} = 0$ for any given x_t we deduce a closed-form LQR solution, $w^* = -a/b$.

Of course, neurons must implement this control law without prior knowledge of a and b , a challenge adeptly addressed by the DD-DC model. Incorporating $x_{t+1} = 0$ into Eq. (4), we obtain:

$$\begin{bmatrix} 0 \\ x_t \\ u_t^* \end{bmatrix} = \begin{bmatrix} x_{1+1} & \dots & x_{k+1} \\ x_1 & \dots & x_k \\ u_1 & \dots & u_k \end{bmatrix} \mathbf{g} = \begin{bmatrix} \mathbf{X}_+ \\ \mathbf{X} \\ \mathbf{U} \end{bmatrix} \mathbf{g}, \quad [8]$$

where we introduced row-vector notation, $\mathbf{X} = [x_1 \dots x_k]$, $\mathbf{X}_+ = [x_{1+1} \dots x_{k+1}]$, and $\mathbf{U} = [u_1 \dots u_k]$.

To determine the optimal control signal u_t^* , we first solve the top two rows of Eq. (8) for \mathbf{g} . Given the underdetermined nature of the problem (k typically exceeds the combined dimensions of x and u), we express \mathbf{g} via a pseudoinverse:

$$\begin{aligned} \mathbf{g} &= [\mathbf{X}_+^\top \ \mathbf{X}^\top] \left(\begin{bmatrix} \mathbf{X}_+ \\ \mathbf{X} \end{bmatrix} [\mathbf{X}_+^\top \ \mathbf{X}^\top] \right)^{-1} \begin{bmatrix} 0 \\ x_t \end{bmatrix} = \\ &= \frac{[\mathbf{X}_+^\top \ \mathbf{X}^\top]}{\mathbf{X}_+ \mathbf{X}_+^\top \mathbf{X} \mathbf{X}^\top - (\mathbf{X}_+ \mathbf{X}^\top)^2} \begin{bmatrix} \mathbf{X} \mathbf{X}^\top & -\mathbf{X}_+ \mathbf{X}^\top \\ -\mathbf{X} \mathbf{X}_+^\top & \mathbf{X}_+ \mathbf{X}_+^\top \end{bmatrix} \begin{bmatrix} 0 \\ x_t \end{bmatrix}, \end{aligned} \quad [9]$$

Subsequently, we substitute this \mathbf{g} into the bottom row of Eq. (8) to obtain

$$u_t^* = \mathbf{U} \mathbf{g} = \frac{\mathbf{U} \mathbf{X}^\top \mathbf{X}_+ \mathbf{X}_+^\top - \mathbf{U} \mathbf{X}_+^\top \mathbf{X}_+ \mathbf{X}^\top}{\mathbf{X}_+ \mathbf{X}_+^\top \mathbf{X} \mathbf{X}^\top - (\mathbf{X}_+ \mathbf{X}^\top)^2} x_t, \quad [10]$$

This formulation can be interpreted as a control law Eq. (7), with

$$w^* = \frac{\mathbf{U} \mathbf{X}^\top - \mathbf{U} \mathbf{X}_+^\top \mathbf{X}_+ \mathbf{X}^\top (\mathbf{X}_+ \mathbf{X}_+^\top)^{-1}}{\mathbf{X} \mathbf{X}^\top - (\mathbf{X} \mathbf{X}_+^\top)^2 (\mathbf{X}_+ \mathbf{X}_+^\top)^{-1}}. \quad [11]$$

Notably, this control law obviates the need for a neuron to calculate \mathbf{g} at each timestep or retain all past values of \mathbf{U} , \mathbf{X} , and

\mathbf{X}_+ . Instead, it requires only the storage and update of their covariances, a biologically plausible process previously utilized in similarity matching networks (42). Rewriting these covariances as sums over recent history and omitting the denominator (a positive scalar independent of the control signal) yields:

$$w^* \sim \sum_{\tau=1}^k u_\tau x_\tau - \cos(\widehat{\mathbf{X} \mathbf{X}_+}) \sum_{\tau=1}^k u_\tau x_{\tau+1}, \quad [12]$$

where $\cos(\widehat{\mathbf{X} \mathbf{X}_+}) = (\mathbf{X}_+ \mathbf{X}_+^\top)^{-1} \mathbf{X}_+ \mathbf{X}^\top$.

In a neurophysiological context, x and u in Eq. (12) symbolize pre- and post-synaptic neuronal activities, respectively, with non-zero values during spikes. Consequently, the sums in Eq. (12) accrue contributions solely when pre- and post-synaptic spikes are temporally proximate and depending on their temporal order. Note that although the indices of u and x are the same in the first sum, this is an artifact of the discrete-time setting, and u must be delayed by at least a fraction of the time step to be computed from x . This model naturally accounts for the transition from potentiation to depression observed in STDP, as well as depression being weaker than potentiation (as $\cos(\widehat{\mathbf{X} \mathbf{X}_+}) < 1$ generically), aligning with empirical findings (24–26) (Fig. 1D).

The potentiation window of STDP can be viewed as the extension of the Hebbian rule (43) to the temporal interplay between spikes. However, the rationale behind the relatively narrow depression window in STDP has remained elusive, largely due to its seemingly anti-causal nature. Current explanations of this phenomenon, e.g. (44, 45), rely on ad hoc assumptions. In contrast, our work shows that the apparent anti-causal aspect of STDP is a natural outcome of conceptualizing neurons as feedback controllers. Specifically, a pre-synaptic spike following a post-synaptic spike conveys information to the neuron about the effectiveness or ineffectiveness of its control of the environment. Thus, what is initially perceived as an anti-causal feature in STDP transforms into a causal mechanism when viewed through the lens of the feedback loop inherent in the controller model.

Next, we consider a DD-DC LQR controller with $r > 0$, which in the scalar, one-step time-horizon case is given by:

$$w^* = \frac{\mathbf{U} \mathbf{X}^\top \|\mathbf{X}_+\|^2 - \mathbf{U} \mathbf{X}_+^\top \mathbf{X} \mathbf{X}_+^\top}{\|\mathbf{X}\|^2 \|\mathbf{X}_+\|^2 - (\mathbf{X} \mathbf{X}_+^\top)^2 + (\|\mathbf{U}\|^2 \|\mathbf{X}\|^2 - (\mathbf{U} \mathbf{X}^\top)^2) r/q}. \quad [13]$$

The full derivation of this solution is given in the Supplement but it aligns with the optimal LQR gain:

$$w^* = \frac{-ab}{b^2 + r/q}, \quad [14]$$

as confirmed by substituting $\mathbf{X}_+ = a\mathbf{X} + b\mathbf{U}$ into Eq. (13) and dividing both the numerator and the denominator of Eq. (13) by the common expression $(\|\mathbf{U}\|^2 \|\mathbf{X}\|^2 - (\mathbf{U} \mathbf{X}^\top)^2)$ thus reducing it to Eq. (14).

This solution provides a potential framework to interpret the variance in STDP profiles documented in various studies (46) in terms of variance of r .

However, as our Eq. (12) involves covariances with only two time lags, it does not fully describe the time-course shown in Fig. 1D. This limitation, inherent to the scalar dynamics model Eq. (5), motivates the exploration of higher order dynamics, as discussed in an upcoming section on reconstructing temporal filters.

Closed-loop DD-DC: malfunction under constant control law and restoration of function by adding noise to control

In this Section, we investigated the functioning of the DD-DC LQR controller through numerical simulations aimed at stabilizing a potentially unstable scalar dynamical system, Eq. (5). Initially, we operated the controller in an open-loop mode for four time steps, implementing white-noise control, u , and tracking the resultant state variable, x . Subsequently, we computed the controller gain, w , by integrating the recorded values of u and x into the general LQR solution, Eq. (13). Following this initialization, we transitioned to a closed-loop operation of the DD-DC LQR controller, recalculating w at each time step (see Supplementary Material for details). Our findings reveal that the DD-DC LQR controller successfully identifies and maintains the optimal value of w up to time = 25, Fig. 2 Left.

For the DD-DC controller to effectively replicate the adaptive behavior of a biological neuron, it must adjust to the evolving dynamics within a real-time, closed-loop framework. Accordingly, the update algorithm for w incorporates a discount factor, progressively diminishing the influence of older data on covariance calculations (see Supplementary Material for details). Initially, the controller learns and applies the optimal value of w . However, when the parameters a and b undergo a switch (at time = 25), the controller fails to adapt, Fig. 2 Left. To uncover the cause behind the DD-DC controller's failure to adjust following the static control law phase, we re-examined the data matrix entering Eq. (8):

$$\begin{bmatrix} x_{1+1} & \dots & x_{k+1} \\ x_1 & \dots & x_k \\ u_1 & \dots & u_k \end{bmatrix} = \begin{bmatrix} x_{1+1} & \dots & x_{k+1} \\ x_1 & \dots & x_k \\ wx_1 & \dots & wx_k \end{bmatrix}. \quad [15]$$

Note that the sub-matrix composed of the x and u rows is rank deficient, thereby contravening the persistence of excitation condition. This rank deficiency signifies a critical limitation in the DD-DC's learning capability, as it impairs the system's ability to extract meaningful information from the data. Operationally, this issue manifests in the denominator of Eq. (13) approaching zero.

Considering that sensory input may sometimes be constant (47) and control efforts typically aim for optimality, the question arises: how can the vulnerability of the online DD-DC controller be mitigated? Building upon the suggestions of control theorists (41), we propose that the brain deliberately generates variability and/or noise, denoted as η , to sustain the persistence of excitation condition even under static input or control regimes. Formally,

$$u_t^* = w^* x_t + \eta_t, \quad [16]$$

Implementing the control law Eq. (16) in the same dynamical system is illustrated in Figure 2 Right for $\text{VAR}(\eta_t) = 10^{-6}$. The addition of low-variance noise to the control signal re-establishes the functionality of the DD-DC in an online setting. This noise reinstates the persistence of excitation, thereby restoring the DD-DC's operational efficacy. However, introducing too high-variance noise could cause significant deviations of the control signal from the LQR optimal Eq. (14), implying that there might exist an ideal noise variance level for optimal control performance (see Supplementary Material for details).

Is there empirical evidence supporting the presence of such noise within the brain? Locally, such noise might originate from the inherent unreliability of synaptic transmission (28). On a broader scale, the brain encompasses specific circuits and neurons dedicated to introducing noise and/or variability, as evidenced in studies

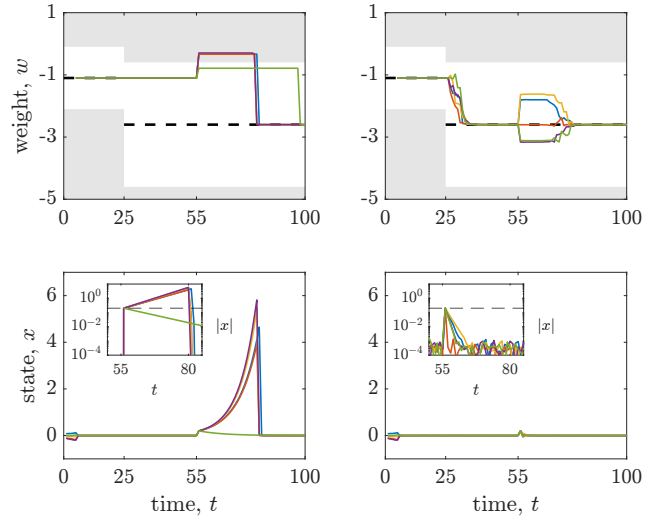


Fig. 2. Dynamics of the weight, w , and state variable, x , in the online DD-DC LQR controller over time with the dynamical systems switch at time = 25. Closed loop control begins at $t = 5$, at which point the weight, w , is constant and optimal (dashed line) and the state variable x rapidly converges to zero. Post-switch, in scenarios without control noise (Left), the DD-DC controller struggles to adapt due to the loss of persistence of excitation, leading to a sub-optimal w ; however, x remains at zero, indicating no direct loss impact. To demonstrate the controller's maladaptation, a jolt $\Delta x = .2$ (dashed line in inset) is applied to the state variable x at time $t = 55$. Introducing noise into the control law (Right), Eq. (16), facilitates exploration thus re-establishing persistence of excitation and overall controller performance. This is evidenced by x swiftly returning to zero and w reverting to its optimal state. The white bands within the gray shaded areas in the w plots represent regions of stability and instability respectively.

on songbirds (30) and *C. elegans* (29). Additionally, the brain can introduce variability in sensory inputs through the generation of corresponding motor outputs, exemplified by microsaccades (48). Given these mechanisms, the operational variability observed in neural representations (31, 32), appears less paradoxical and more a natural consequence of the brain's function as a DD-DC controller.

Reconstruction of feedforward and feedback temporal filters from data

We now explore the DD-DC of a dynamical system of order $n > 1$, equipped with a scalar control signal, u , and a scalar observation, y (Fig. 3A). In such systems, observations are partial and insufficient for direct control, necessitating that the controller estimate the latent state, x . For linear systems, this latent state can be inferred from recent sequences of observations and controls using time-delay embedding techniques (22),

$$\hat{x}(t) = [y_{t-n} \dots y_{t-1} \quad u_{t-n} \dots u_{t-1}]^T, \quad [17]$$

enabling us to reformulate the control law, Eq. (3), as:

$$u_t = [\mathbf{K}_{ff} \quad \mathbf{K}_{fb}] [y_{t-n} \dots y_{t-1} \quad u_{t-n} \dots u_{t-1}]^T. \quad [18]$$

Here, the feedforward, \mathbf{K}_{ff} , and feedback, \mathbf{K}_{fb} , temporal filters collectively form an Auto-Regressive Moving Average (ARMA) model of a neuron (Fig. 3A).

In this Section, rather than optimizing feedforward and feedback temporal filters, we estimate them from experimental data. In our experiments, neurons are isolated from the loop and stimulated with sensory input or injected current, y , and the neuronal

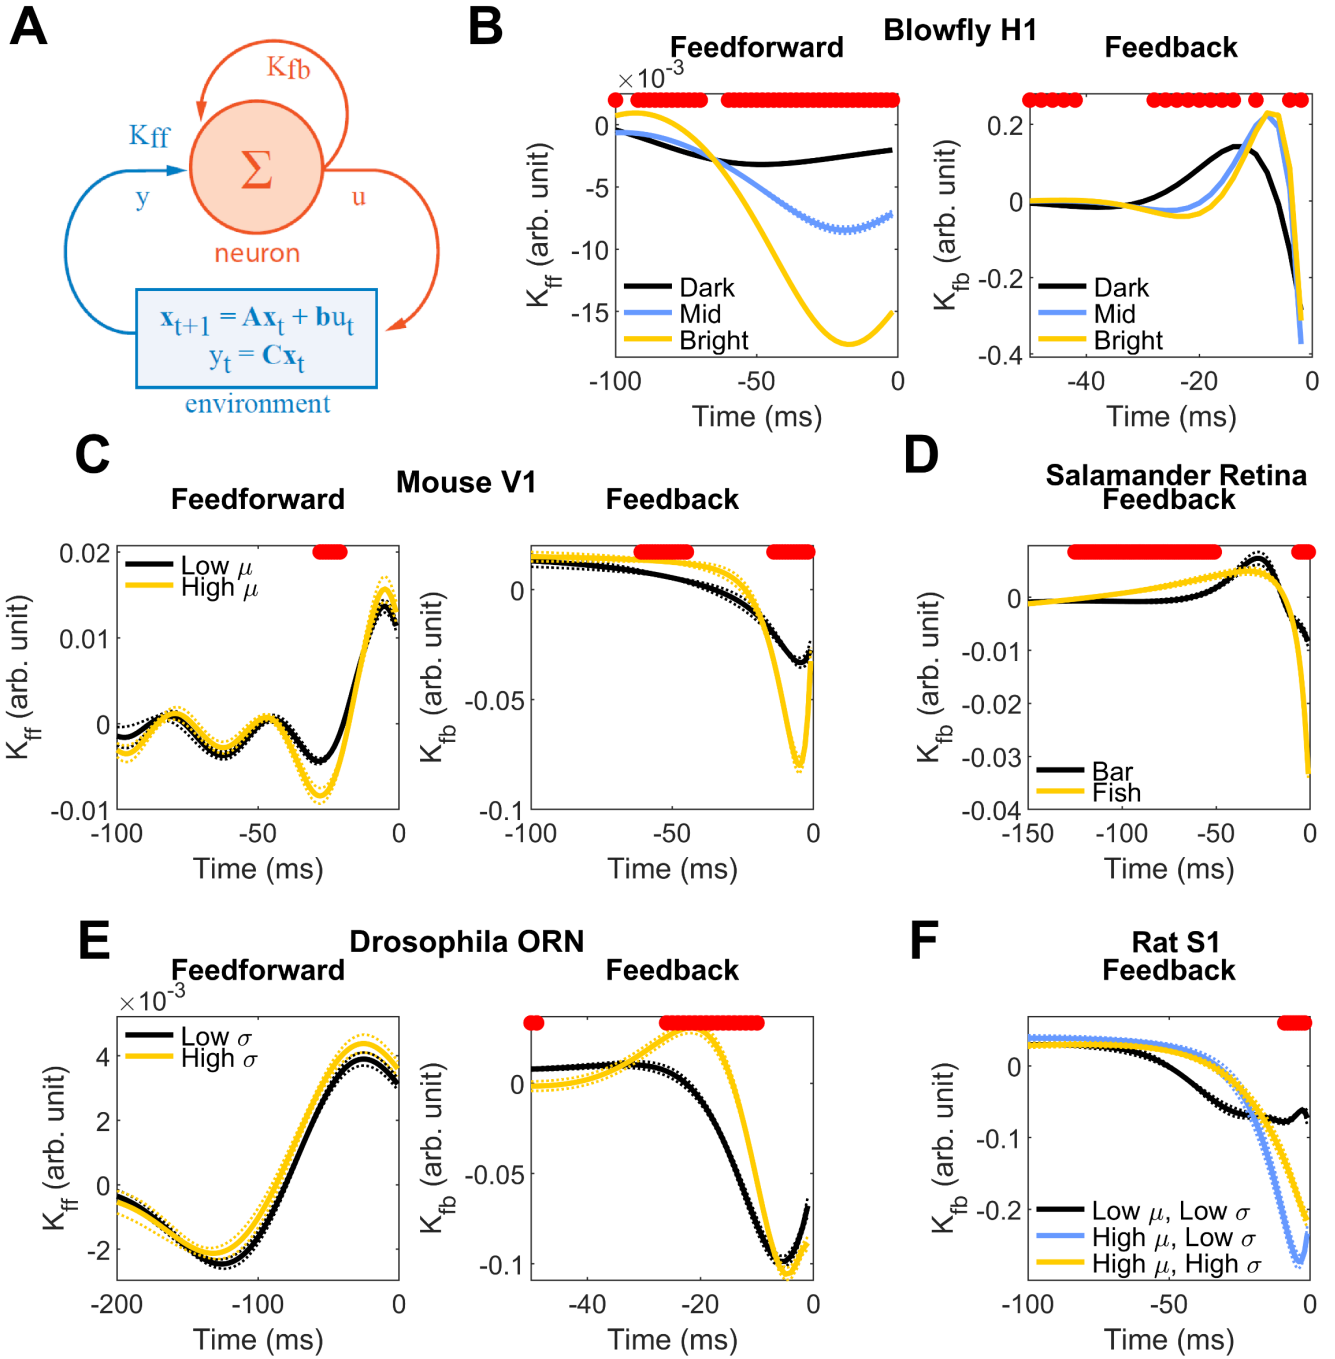


Fig. 3. A. Illustration of a neuron modeled as an Auto-Regressive Moving Average (ARMA) controller, characterized by feedforward, \mathbf{K}_{ff} , and feedback, \mathbf{K}_{fb} , temporal filters. **B-D.** Adaptation of experimentally measured temporal filters (depicted in black, yellow, and blue) to input signal statistics. Solid lines represent mean values, while thin dotted lines denote standard errors of the mean. Regions where differences are statistically significant (Wilcoxon rank-sum test with Bonferroni correction for multiple comparisons) are highlighted in red. **B.** Variation in feedforward (akin to decorrelated Spike-Triggered Average, STA) and feedback (analogous to spike-history-dependence) filters of the blowfly H1 neuron (49), responding to visual motion against different background luminance levels. **C.** Feedforward and feedback filters in pyramidal cells from mouse primary visual cortex (50) responding to current injections with varying mean levels. **D.** Feedback filters in a salamander retinal ganglion cell (8) for stimuli comprising a drifting bar and a fish movie (feedforward filter data unavailable). **E.** Adaptation of feedforward and feedback filters in a *Drosophila* Olfactory Receptor Neuron (ORN) (51) to odorant concentrations with varying variances. **F.** Feedback filters in rat somatosensory cortex pyramidal neurons (52), responding to current injections modulated by an Ornstein-Uhlenbeck process atop a DC component. Feedforward filters are provided in the Supplement.

response, u , is recorded. These data are compiled into matrices $\mathbf{U} = [u_1 \dots u_t]$ and $\hat{\mathbf{X}} = [\hat{x}_1 \dots \hat{x}_t]$, which are linearly related:

$$\mathbf{U} = [\mathbf{K}_{ff} \ \mathbf{K}_{fb}] \hat{\mathbf{X}}. \quad [19]$$

We solve for the filters by linear regression using a pseudoinverse,

$$[\mathbf{K}_{ff} \ \mathbf{K}_{fb}] = \mathbf{U} \hat{\mathbf{X}}^\top (\hat{\mathbf{X}} \hat{\mathbf{X}}^\top)^{-1}. \quad [20]$$

Utilizing Eq. (20), we reconstruct the feed forward, \mathbf{K}_{ff} , and feed-

back, \mathbf{K}_{fb} , temporal filters from experimental data across various model systems, Fig. 3 (see Supplement for details).

While these temporal filters have been previously measured (53–55), our controller-based perspective offers a novel interpretation. In a partially observed system, instantaneous observations alone are inadequate for control. Thus, feedforward and feedback filters with finite temporal extents are essential for control actions that are coherent with the latent state Eq. (17). The temporal extension of these filters aligns with the hypothesis that neuronal output influences the environment’s latent state and acknowledges the partial observation of the system.

In each system we studied, neurons were recorded under various conditions, with stimulus statistics changing between these conditions. These included the blowfly H1 neuron with varying background luminance (49), mouse V1 pyramidal neurons responding to different mean injected current waveforms (50), salamander retinal ganglion cells exposed to distinct visual stimuli (8), *Drosophila* olfactory receptor neurons reacting to varying odorant concentrations (51), and pyramidal neurons in the rat somatosensory cortex stimulated with current injections of diverse means and variances (52). As shown in Fig. 3, the filter shapes adapt to both the mean (μ) and variance (σ) of the input statistics.

The adaptation of feedforward filters to input changes is well-documented (2, 4, 14) and suggests a functional role beyond mere biological necessity, explainable by efficient coding and predictive information theories (3, 5, 12). Our analysis extends this understanding to feedback filters, which also adapt to changing stimulus statistics. This adaptation, not predicted by existing theories, calls for a new framework that treats feedforward and feedback filters equally, such as the controller neuron model.

Spike generation mechanism loses precision under constant input

Neuronal spike generation typically showcases remarkable precision: repeated injections of the same current waveform into a neuron yield highly reproducible spike trains, precise down to milliseconds (13, 14) (Fig. 4, Left). This level of precision in spike timing must incur metabolic cost and is therefore suggestive of a functional significance. Intriguingly, this precision deteriorates when the neuron is subject to a constant current input (14) (Fig. 4, Right), exposing a notable limitation in the spike-generation mechanism. The DD-DC model of neuronal function offers an

insightful explanation for this observed decline in spike-timing precision with constant input. The DD-DC model posits that a neuron reconstructs any state as a weighted sum of past states, which is effective only when these past states are sufficiently varied (21). This is rooted in the persistency of excitation condition, requiring the matrix of past states in Eq. (4) to have full row-rank. Under constant input, however, this condition fails as the lag vectors in Eq. (17) become uniform. Consequently, when a neuron processes recent history (approximately 100ms in Fig. 4, Right), the DD-DC model predicts erratic outputs in response to a constant current, mirroring the vulnerability of the spike-generation mechanism to such inputs.

Traditionally, the variability in spike timing under constant input was ascribed to intrinsic ion channel noise, not controlled for in experimental setups (56). This noise was thought to be inconsequential for spike timing in the presence of highly variable inputs, as the effects would be overshadowed by the abundance of open ion channels. However, even slight variations in injected current (STD \approx 50pA) are observed to restore spike-timing precision (52), posing questions about the underlying mechanisms of such sensitivity. Our model offers a different perspective, suggesting that this sensitivity to low-variance noise stems from the high condition number (ratio of the largest to the smallest singular values) of time-delay covariance matrices. These matrices must be inverted to compute a neuron’s response (akin to Eq. (20)). Our hypothesis posits that the singularity at constant current can be empirically validated by measuring spike time variability against noise variance below the threshold reported in (14), and correlating it with the condition number of the time-delay covariance.

Discussion

The power of the proposed DD-DC model of a neuron is in that, starting from a single postulate, it offers explanations for multiple seemingly unrelated neurophysiological phenomena, including the switch between potentiation and depression in STDP and its asymmetry, the extended nature and input-dependent adaptation of feedforward and feedback temporal filters, the imprecision of the spike-generation mechanism under constant input, and the prevalence of operational variability and noise in the brain. Although each of these explanations provides only a circumstantial evidence, their multitude and variety provide strong support for the DD-DC model. This perspective has the potential to deepen and refine our

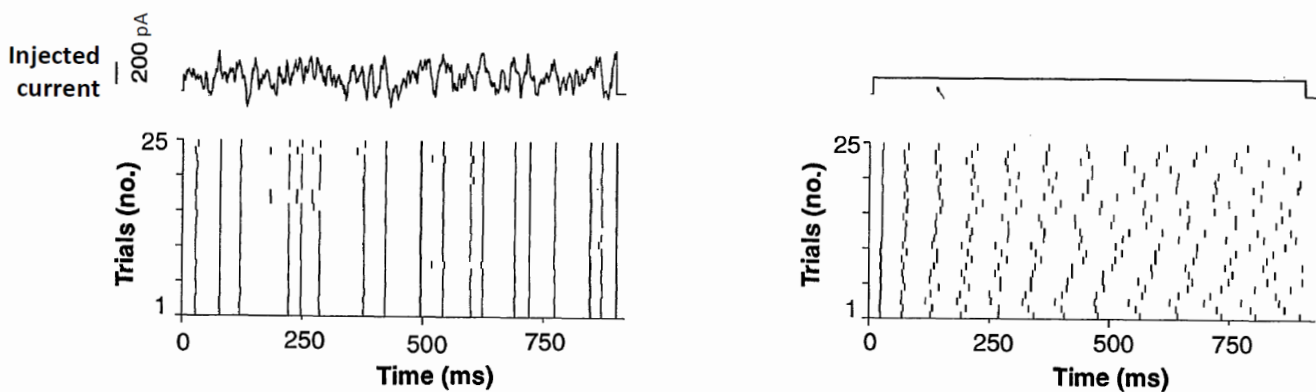


Fig. 4. Left: High-variance current injections into a neuron yield remarkably consistent spike trains over multiple trials, showcasing the precision of the spike-generation mechanism. Right: In contrast, a constant current input leads to notably variable spike trains, revealing a significant reduction in spike-timing precision. This dichotomy highlights the neuron’s differential response to varying and constant stimuli (14).

understanding of the brain, and may also aid in the development of biologically-inspired artificial neural networks.

Nonlinear dynamics and control. In this study, we focused on a DD-DC model that assumes discrete-time and linear dynamics of the environment. In reality, the dynamics are continuous-time and nonlinear. As is often the case in control theory, we expect that our framework can be naturally extended to continuous time. As nonlinear dynamics of the loop can be approximated locally as linear, we speculate that they can be modeled by a switching linear system controlled by a set of switching DD-DCs (see Fig. 5, *Left*). This would explain why layers of processing in the brain contain many neurons in parallel performing analogous functions.

How to derive a nonlinear controller model of a neuron from the normative perspective? Even for linear plant dynamics, apart from special cases like LQR, the optimal controller may not be linear. Perhaps, neuronal action potentials, in addition to having higher information transmission capacity relative to graded potentials in noisy environments (57–59), have other operational advantages similar to widely used bang-bang control (60). Deriving such a controller may help model neurons with active conductances and spikes (61) on the algorithmic level.

Stability and performance objectives. In control theory, the stability of the closed loop is of primary concern. While a stability criterion for DD-DC can be formulated (22), it allows for numerous solutions. How to select a specific solution out of the stable set is not clear. One approach could be to look for the most stable solution, which would remain stable even in the presence of noise and uncertainty about the parameters of the dynamical system. Another approach could be based on the observation that the brain operates at the edge of chaos (62, 63), favoring borderline stable solutions. A similar borderline stable solution is suggested by viewing a neuron as an integrator, which would require the top eigenvalue to have a unit norm (64). Such flexibility in the choice of the objective may allow one to use different solutions to model different neuronal classes.

A network of DD-DC neurons. The DD-DC model of a neuron presented here lumps the rest of the neurons into a single dynamical system, yet each fellow neuron can also be modeled as a DD-DC. This raises the question of how multiple DD-DC neurons interact with each other in a network. We leave this question to future work and comment only on several experimental observations that support this view.

First, measurements of synaptic plasticity may shed light on the temporal lag caused by the feedback loop. In the case of

STDP (Fig. 1, *Right*), pre- and post-synaptic spikes must be almost synchronous for plasticity to occur, indicating that the loop traverses an order of one synapse. This finding corresponds to the known abundance of short local feedback loops in the cortex (16, 65). At the other extreme, the plasticity of some synapses in the cerebellum and the hippocampus peaks when the spikes lag by tens of milliseconds (66, 67). This suggests longer loops involving different brain regions or even the external environment (68). The abundance of loops is not limited to mammalian brains (15) and has been reported in invertebrates as well (17, 18).

Second, the ability of individual neurons to control long (multi-synaptic and trans-environment) loops may seem unrealistic. However, theoretical analysis (69, 70) and experimental observations seem to support long-range propagation of signals from individual neurons (71). Specifically, rodents can be trained to behaviorally report single-neuron electrical stimulation in the barrel cortex (72) suggesting that the spikes of a single neuron make impact sufficiently far downstream to elicit behavior. Also, stimulation of single neurons in the motor cortex evoke whisker movements (73) suggesting that single neurons can produce observable effects on the environment. Taken together, these experiments support the propagation of individual neurons' spikes around long loops.

Third, modeling neurons as DD-DC controllers offers an explanation for the representation of movement outside of the motor cortex (34, 35).

As DD-DC neurons in every layer (Fig. 5, *Right*) combine systems identification and control, they acquire characteristics of both sensory and motor representations. Therefore, it is natural to expect a mixed representation at every layer with a gradual transformation from sensory to motor. The controller perspective also accounts for the context dependence of neuronal representations (33). As neuronal activity should not just reflect the sensory stimulus but rather optimal control, which is context dependent, it is natural to expect such representations.

Relationship to Other Work

The concept of modeling neurons as controllers intersects with several established research avenues. A notable idea in neuroscience relates neuronal output to a prediction of future inputs (2, 6, 11, 12). To optimize control the DD-DC neuron implicitly infers environmental dynamics, which could also be used for prediction. However, the controller neuron does not just predict the future input but aims to influence it through its output.

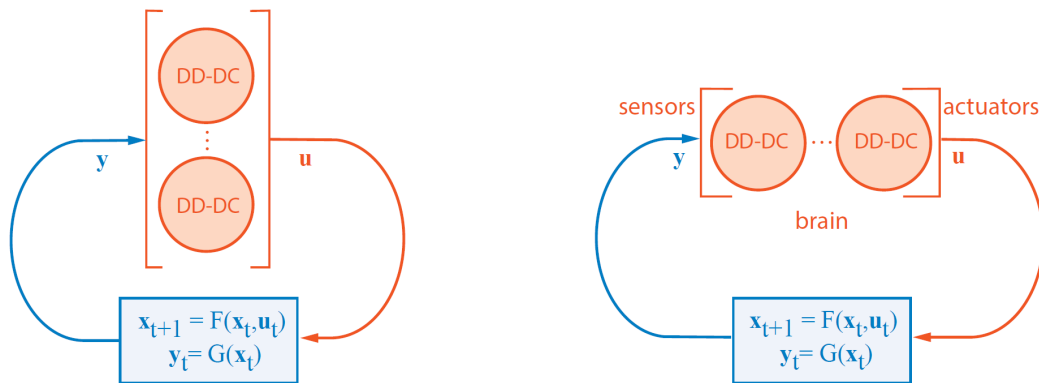


Fig. 5. Left: Illustration of controlling a nonlinear dynamical system using multiple switching DD-DCs. Right: Depiction of a deep network model where each neuron exerts control over its immediate environment, contributing to the broader control exerted by the entire brain over the external environment.

Normalization models, which have been both experimentally observed and theoretically justified (74–76), resemble the feedback filter in our model. However, these models focus on interactions among parallel channels under static stimuli, overlooking temporal correlations and stimulus dynamics. In contrast, the DD-DC model proactively controls inputs through its influence on the underlying dynamical system.

Our approach also aligns with the idea that spiking neuron networks encode temporally variable inputs (77). Although these networks are predominantly feedforward and don't allow neuron outputs to modify network inputs, both concepts emphasize learning generative dynamics.

Data-driven control in network contexts, including brain networks, has been investigated (78). However, these studies generally involve controlling networks through external perturbations and lack a focus on biological plausibility. Unlike our neuron-centric DD-DC model, they require access to multiple network nodes and are constrained by the resolution of technologies like fMRI.

The DD-DC ARMA model for neurons (Fig. 3A) shares similarities with Generalized Linear Models (GLMs) (55), notably in possessing feedforward and feedback filters. But, while GLMs are stochastic and nonlinear, the DD-DC model is deterministic, linear, and provides a rationale for the duration and adaptation of these filters. The concept of temporal integration in our model also echoes the principles of integrate-and-fire models (79), laying groundwork for future connections between these theories.

Differing from the conventional McCulloch-Pitts-Rosenblatt unit in artificial neural networks, the DD-DC neuron integrates inputs over time and features an auto-regressive loop, unlike the instantaneous response of standard units. Also, in contrast to network-wide optimization in artificial neural networks, the DD-DC model optimizes objectives at the neuronal level.

Neurons are sometimes conceptualized as agents in the reinforcement learning (RL) paradigm, (80, 81). While control theory and RL share commonalities, key distinctions include control theory's implicit dynamical systems model of the environment and its focus on optimizing specific objectives based on controls and observations, as opposed to the reward-maximization approach in RL.

Previous studies have conceptualized the whole brain as a controller acting on the external world (82–84), and used LQR with delays/noise to model internal feedback (85). Our DD-DC approach extends this concept to individual neurons. If both the entire brain and single neurons can be modeled as controllers, intermediate levels of brain structure might also fit this model (86, 87). Earlier models separated sensory system identification and motor control (88), but our unified neuron-as-controller model eliminates the need to match corresponding sensory and motor units. Future research could explore how coherent controller actions emerge in self-organized networks of such neuron-controllers.

- 1 Horace Barlow. Redundancy reduction revisited. *Network: computation in neural systems*, 12(3):241, 2001.
- 2 Mandyam Veeramabudi 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.
- 3 Joseph J Atick. Could information theory provide an ecological theory of sensory processing? *Network: Computation in neural systems*, 3(2):213–251, 1992.
- 4 Johannes H van Hateren. Theoretical predictions of spatiotemporal receptive fields of fly Imcs, and experimental validation. *Journal of Comparative Physiology A*, 171:157–170, 1992.
- 5 Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001.
- 6 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.

- 7 Na Young Jun, Greg Field, and John Pearson. Efficient coding, channel capacity, and the emergence of retinal mosaics. *Advances in neural information processing systems*, 35:32311–32324, 2022.
- 8 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.
- 9 Nicole C Rust and Stephanie E Palmer. Remembering the past to see the future. *Annual review of vision science*, 7:349–365, 2021.
- 10 Siwei Wang, Idan Segev, Alexander Borst, and Stephanie Palmer. Maximally efficient prediction in the early fly visual system may support evasive flight maneuvers. *PLoS computational biology*, 17(5):e1008965, 2021.
- 11 Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- 12 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.
- 13 Hugh L Bryant and José P Segundo. Spike initiation by transmembrane current: a white-noise analysis. *The Journal of physiology*, 260(2):279–314, 1976.
- 14 Zachary F Mainen and Terrence J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- 15 Anthony J Bell. Levels and loops: the future of artificial intelligence and neuroscience. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 354(1392):2013–2020, 1999.
- 16 Sen Song, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, 3(3):e68, 2005.
- 17 Lav R Varshney, Beth L Chen, Eric Paniagua, David H Hall, and Dmitri B Chklovskii. Structural properties of the caenorhabditis elegans neuronal network. *PLoS computational biology*, 7(2):e1001066, 2011.
- 18 Michael Winding, Benjamin D Pedigo, Christopher L Barnes, Heather G Patsolic, Youngser Park, Tom Kazimiers, Akira Fushiki, Ingrid V Andrade, Avinash Khandelwal, Javier Valdes-Aleman, et al. The connectome of an insect brain. *Science*, 379(6636):eadd9330, 2023.
- 19 Tohru Katayama et al. *Subspace methods for system identification*, volume 1. Springer, 2005.
- 20 Karl Johan Åström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2021.
- 21 Jan C Willems, Paolo Rapisarda, Ivan Markovsky, and Bart LM De Moor. A note on persistency of excitation. *Systems & Control Letters*, 54(4):325–329, 2005.
- 22 Claudio De Persis and Pietro Tesi. Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control*, 65(3):909–924, 2019.
- 23 Ivan Markovsky, Linbin Huang, and Florian Dörfler. Data-driven control based on the behavioral approach: From theory to applications in power systems. *IEEE Control Syst.*, 2022.
- 24 Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science*, 275(5297):213–215, 1997.
- 25 Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998.
- 26 Li I Zhang, Huizhong W Tao, Christine E Holt, William A Harris, and Mu-ming Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395(6697):37–44, 1998.
- 27 Michael E Rule, Timothy O'Leary, and Christopher D Harvey. Causes and consequences of representational drift. *Current opinion in neurobiology*, 58:141–147, 2019.
- 28 Lynn E Dobrunz and Charles F Stevens. Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron*, 18(6):995–1008, 1997.
- 29 Andrew Gordus, Navin Pokala, Sagi Levy, Steven W Flavell, and Cornelia I Bargmann. Feedback from network states generates variability in a probabilistic olfactory circuit. *Cell*, 161(2):215–227, 2015.
- 30 Bence P Ölveczky, Aaron S Andalman, and Michale S Fee. Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS biology*, 3(5):e153, 2005.
- 31 Daniel Deitch, Alon Rubin, and Yaniv Ziv. Representational drift in the mouse visual cortex. *Current biology*, 31(19):4327–4339, 2021.
- 32 Carl E Schoonover, Sarah N Ohashi, Richard Axel, and Andrew JP Fink. Representational drift in primary olfactory cortex. *Nature*, 594(7864):541–546, 2021.
- 33 Christopher M Niell and Michael P Stryker. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4):472–479, 2010.
- 34 Simon Musall, Matthew T Kaufman, Ashley L Juavinett, Steven Gluf, and Anne K Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience*, 22(10):1677–1686, 2019.
- 35 Edward Zagha, Jeffrey C Erlich, Soohyun Lee, Gyorgy Lur, Daniel H O'Connor, Nicholas A Steinmetz, Carsen Stringer, and Hongdian Yang. The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. *Journal of Neuroscience*, 42(8):1375–1382, 2022.
- 36 Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.
- 37 Frank Rosenblatt. Principles of neurodynamics, perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, 1961.
- 38 Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows of the deepc. In *2019 18th European Control Conference (ECC)*, pages 307–312. IEEE, 2019.
- 39 Julian Berberich, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. Linear tracking mpc for nonlinear systems—part ii: The data-driven case. *IEEE Transactions on Automatic Control*, 67(9):4406–4421, 2022.
- 40 Henk J van Waarde. Beyond persistent excitation: Online experiment design for data-driven modeling and control. *IEEE Control Systems Letters*, 6:319–324, 2021.
- 41 Monica Rotulo, Claudio De Persis, and Pietro Tesi. Online learning of data-driven controllers for unknown switched linear systems. *Automatica*, 145:110519, 2022.

- 42 Cengiz Pehlevan and Dmitri B Chklovskii. Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. *IEEE Signal Processing Magazine*, 36(6):88–96, 2019.
- 43 Donald O. Hebb. *The organization of behavior: A neuropsychological theory*. Wiley, New York, 1949.
- 44 Matthieu Gilson and Tomoki Fukai. Stability versus neuronal specialization for stdp: long-tail weight distributions solve the dilemma. *PLoS one*, 6(10):e25339, 2011.
- 45 Pau Vilimelis Aceituno, Matilde Tristany Farinha, Reinhard Loidl, and Benjamin F Grewe. Learning cortical hierarchies with temporal hebbian updates. *Frontiers in Computational Neuroscience*, 17:1136010, 2023.
- 46 Natalia Caporale and Yang Dan. Spike timing-dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.*, 31:25–46, 2008.
- 47 Roy M Pritchard. Stabilized images on the retina. *Scientific American*, 204(6):72–79, 1961.
- 48 Hee-kyoung Ko, Martina Poletti, and Michele Rucci. Microsaccades precisely relocate gaze in a high visual acuity task. *Nature neuroscience*, 13(12):1549–1553, 2010.
- 49 GD Lewen, W Bialek, and RR De Ruyter Van Steveninck. Neural coding of naturalistic motion stimuli. *Network: Computation in Neural Systems*, 12(3):317, 2001.
- 50 Corinne Teeter, Ramakrishnan Iyer, Vilas Menon, Nathan Gouwens, David Feng, Jim Berg, Aaron Szafer, Nicholas Cain, Hongkui Zeng, Michael Hawrylycz, et al. Generalized leaky integrate-and-fire models classify multiple neuron types. *Nature communications*, 9(1):709, 2018.
- 51 Srinivas Gorur-Shandilya, Mahmut Demir, Junjia Long, Damon A Clark, and Thierry Emonet. Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli. *Elife*, 6:e27670, 2017.
- 52 Alexander Rauch, Giancarlo La Camera, Hans-Rudolf Luscher, Walter Senn, and Stefano Fusi. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *J. Neurophysiol.*, 90(3):1598–1612, 2003.
- 53 Fred Rieke, David Warland, Rob de Ruyter van Steveninck, and William Bialek. Spikes: exploring the neural code, 1999.
- 54 Justin Keat, Pamela Reinagel, R Clay Reid, and Markus Meister. Predicting every spike: a model for the responses of visual neurons. *Neuron*, 30(3):803–817, 2001.
- 55 Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, EJ Chichilnisky, and Eero P Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999, 2008.
- 56 Elad Schneidman, Barry Freedman, and Idan Segev. Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural computation*, 10(7):1679–1703, 1998.
- 57 James Julian Bennett Jack, Denis Noble, and Richard W Tsien. Electric current flow in excitable cells. 1975.
- 58 Rahul Sarpeshkar. Analog versus digital: extrapolating from electronics to neurobiology. *Neural computation*, 10(7):1601–1638, 1998.
- 59 Simon B Laughlin, Rob R de Ruyter van Steveninck, and John C Anderson. The metabolic cost of neural information. *Nature neuroscience*, 1(1):36–41, 1998.
- 60 Morton I Kamien and Nancy Lou Schwartz. *Dynamic optimization: the calculus of variations and optimal control in economics and management*. courier corporation, 2012.
- 61 Christof Koch. *Biophysics of computation: information processing in single neurons*. Oxford university press, 2004.
- 62 Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.
- 63 John M Beggs and Dietmar Plenz. Neuronal avalanches in neocortical circuits. *Journal of neuroscience*, 23(35):11167–11177, 2003.
- 64 HS Seung. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. USA*, 93:13339–13344, 1996.
- 65 Y Zilberter. Dendritic release of glutamate suppresses synaptic inhibition of pyramidal neurons in rat neocortex. *The Journal of Physiology*, 528(Pt 3):489, 2000.
- 66 Joshua T Dudman, David Tsay, and Steven A Siegelbaum. A role for synaptic inputs at distal dendrites: instructive signals for hippocampal long-term plasticity. *Neuron*, 56(5):866–879, 2007.
- 67 Aparna Suvrathan, Hannah L Payne, and Jennifer L Raymond. Timing rules for synaptic plasticity matched to behavioral function. *Neuron*, 92(5):959–967, 2016.
- 68 Sriram Jayabal, Brandon J Bhasin, Aparna Suvrathan, Jennifer DiSanto, Mark S. Goldman, and Jennifer L. Raymond. Experience adaptively tunes the timing rules for associative plasticity. *bioRxiv*, 2022. . URL <https://www.biorxiv.org/content/early/2022/11/29/2022.11.28.518128>.
- 69 Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E Latham. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302):123–127, 2010.
- 70 Michael Monteforte and Fred Wolf. Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Physical Review X*, 2(4):041007, 2012.
- 71 Jason Wolfe, Arthur R Houweling, and Michael Brecht. Sparse and powerful cortical spikes. *Current opinion in neurobiology*, 20(3):306–312, 2010.
- 72 Arthur R Houweling and Michael Brecht. Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*, 451(7174):65–68, 2008.
- 73 Michael Brecht, Miriam Schneider, Bert Sakmann, and Troy W Margrie. Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature*, 427(6976):704–710, 2004.
- 74 Martin J Wainwright, Odelia Schwartz, and Eero P Simoncelli. 10 natural image statistics and divisive normalization. *Probabilistic models of the brain*, page 203, 2002.
- 75 Matteo Carandini and David J Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, 2012.
- 76 Naama Brenner, William Bialek, and Rob de Ruyter Van Steveninck. Adaptive rescaling maximizes information transmission. *Neuron*, 26(3):695–702, 2000.
- 77 Martin Boerlin, Christian K Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS computational biology*, 9(11):e1003258, 2013.
- 78 Giacomo Baggio, Danielle S Bassett, and Fabio Pasqualetti. Data-driven control of complex networks. *Nature communications*, 12(1):1429, 2021.
- 79 Louis Lapique. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *Journal de physiologie et de pathologie générale*, 9:620–635, 1907.
- 80 Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- 81 Peter Dayan. Matters temporal. *Trends in Cognitive Sciences*, 6(3):105–106, 2002.
- 82 Paul Cisek. Beyond the computer metaphor: Behaviour as interaction. *Journal of Consciousness Studies*, 6(11-12):125–142, 1999.
- 83 Paul Cisek and John F Kalaska. Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, 33:269–298, 2010.
- 84 Ehud Ahissar and Eldad Assa. Perception as a closed-loop convergence process. *elife*, 5:e12830, 2016.
- 85 Jing Shuang Li, Anish A Sarma, Terrence J Sejnowski, and John C Doyle. Why internal feedback is necessary in the perception-action loop. *arXiv preprint arXiv:2211.05922*, 2022.
- 86 DA Robinson. The use of control systems analysis in the neurophysiology of eye movements. *Annual review of neuroscience*, 4(1):463–503, 1981.
- 87 FA Miles and SG Lisberger. Plasticity in the vestibulo-ocular reflex: a new hypothesis. *Annual review of neuroscience*, 4(1):273–299, 1981.
- 88 Masahiko Haruno, Daniel M Wolpert, and Mitsuo Kawato. Mosaic model for sensorimotor learning and control. *Neural computation*, 13(10):2201–2220, 2001.
- 89 Rhys D Williams, Travis L Massey, and Michel M Maharbiz. Blowfly yaw control via electrical stimulation of the h1 lobula plate tangential cell. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1685–1688. IEEE, 2018.

Supplementary Information

A. Derivation of the DD-DC LQR controller

Given an initial state x_0 , controller finds a control signal u^* that brings the state to the value x^* minimizing LQR loss

$$\|x^*\|^2 + r/q\|u^*\|^2. \quad [21]$$

Corresponding DD-DC optimization problem takes the form:

$$\min_{\mathbf{g}} (\mathbf{g}\mathbf{X}_+^\top)^2 + r/q(\mathbf{g}\mathbf{U}^\top)^2 \quad [22]$$

$$\text{s.t. } (\mathbf{g}\mathbf{X}^\top) = x_0 \quad [23]$$

The chart in Fig. 6 shows sample vectors \mathbf{X} , \mathbf{U} , \mathbf{X}_+ which are co-planar due to the dynamics. Without loss of generality, We consider the solution vector \mathbf{g} to be co-planar with them. The blue line AC orthogonal to \mathbf{X} is the locus of solutions satisfying the constraint, so that $|OB| = x_0/\|\mathbf{X}\|$. We see that $|OC| = |OB|/\cos\phi$, $|OA| = |OB|/\cos\psi$, $|BC| = |OB|\tan\phi$, $|BA| = |OB|\tan\psi$. Denote $d := |DC|$, we'll be solving the problem in

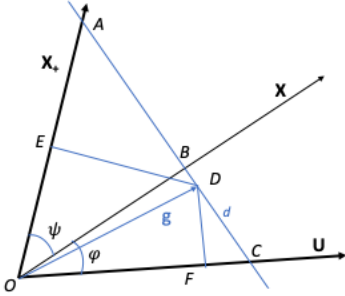


Fig. 6. Illustration for the derivation of DD-DC LQR controller.

this variable. Let DF, DE be orthogonal to \mathbf{U} , \mathbf{X}_+ , then observe that $\angle CDF = \phi$, $\angle ADE = \psi$ as angles with orthogonal sides. This gives expressions:

$$|OF| = |OB|/\cos\phi - d\sin\phi,$$

$$|OE| = |OB|/\cos\psi - |OB|(\tan\phi + \tan\psi)\sin\psi - d\sin\psi.$$

Note also that $\mathbf{g}\mathbf{X}_+^\top = |OE|\|\mathbf{X}_+\|$, $\mathbf{g}\mathbf{U}^\top = |OF|\|\mathbf{U}\|$. Substituting all that, we can rewrite the objective as quadratic expression in d . Solution takes the form

$$d = \frac{(-\sin\psi\cos\psi + \sin^2\psi\tan\phi)\|\mathbf{X}_+\|^2 + \tan\phi\|\mathbf{U}\|^2 r/q}{\cos\phi\cos\psi(\sin^2\psi\|\mathbf{X}_+\|^2 + \sin^2\phi\|\mathbf{U}\|^2 r/q)} \frac{x_0}{\|\mathbf{X}\|}$$

The optimal control signal we seek is $u^* = \mathbf{g}\mathbf{U}^\top = |OF|\|\mathbf{U}\|$, which gives:

$$u^* = \frac{\cos\phi\sin\phi + \cos\psi\sin\psi - \cos\psi\sin\psi\sin^2\phi - \cos\phi\sin\phi\sin^2\psi}{\cos\phi\cos\psi(\sin^2\psi\|\mathbf{X}_+\|^2 + \sin^2\phi\|\mathbf{U}\|^2 r/q)} \times \frac{x_0\|\mathbf{U}\|\|\mathbf{X}_+\|^2\sin\psi}{\|\mathbf{X}\|}$$

For online calculation it is convenient to have an expression in variances and covariances of the variables, so we reformulate trigonometric expressions through cosines only. For that purpose we group terms first and fourth, second and third from the numerator of the first fraction:

$$\begin{aligned} & \cos\phi\sin\phi(1 - \sin^2\psi) + \cos\psi\sin\psi(1 - \sin^2\phi) \\ &= \cos\phi\sin\phi\cos^2\psi + \cos\psi\sin\psi\cos^2\phi \\ &= \cos\phi\cos\psi\sin(\phi + \psi) \end{aligned}$$

Bringing in $\sin\psi$ from the second fraction numerator we use cosine of differences rule to observe:

$$\cos\phi\cos\psi[\sin(\phi + \psi)\sin\psi] = \cos\phi\cos\psi[\cos\phi - \cos\psi\cos(\phi + \psi)]$$

Now we can rewrite the expression for u^* :

$$u^* = \frac{\cos\phi - \cos\psi\cos(\phi + \psi)}{(1 - \cos^2\psi)\|\mathbf{X}_+\|^2 + (1 - \cos^2\phi)\|\mathbf{U}\|^2 r/q} \frac{\|\mathbf{U}\|\|\mathbf{X}_+\|^2}{\|\mathbf{X}\|} x_0$$

Replacing $\cos\phi = \frac{\|\mathbf{X}\mathbf{U}^\top\|}{\|\mathbf{X}\|\|\mathbf{U}\|}$, $\cos\psi = \frac{\|\mathbf{X}\mathbf{X}_+^\top\|}{\|\mathbf{X}\|\|\mathbf{X}_+\|}$, and $\cos(\phi + \psi) = \frac{\|\mathbf{X}_+\mathbf{U}^\top\|}{\|\mathbf{X}_+\|\|\mathbf{U}\|}$, we obtain $u^* = w^*x_0$ where w^* is defined in Eq. (13).

B. Simulation of the DD-DC LQR controller

The simulation presented in Fig. 2 begins in the open loop mode with random initialization of x and random Gaussian control signal with standard deviation 0.01. Open loop runs for 4 steps, after that the control weight, w , is estimated using Eq. (13) and then updated at each time step in the closed loop setting using the variances and covariances of the variables x, u, x_+ . In the online setting, these variances and covariances are updated at each time step according to

$$\text{cov}_{\alpha,\beta}(t+1) = \gamma\text{cov}_{\alpha,\beta}(t) + (1-\gamma)\alpha(t)\beta(t),$$

where α, β are any of x, u, x_+ . For this simulation we set $\gamma = 0.5$. For the noise process, Fig. 2 Right, Gaussian noise with standard deviation $\sigma = 0.001$ was added to the control signal u .

To check the ability of the solution to adapt to changes of the control system parameters a, b were abruptly switched at steps 25 as follows: $a = 1.1, 1.3$; $b = 1, 0.5$. In addition, a jolt was added to the system state $x: 0.2$ at step 55. Throughout the simulation, the LQR penalties were kept constant at $r/q = 0$.

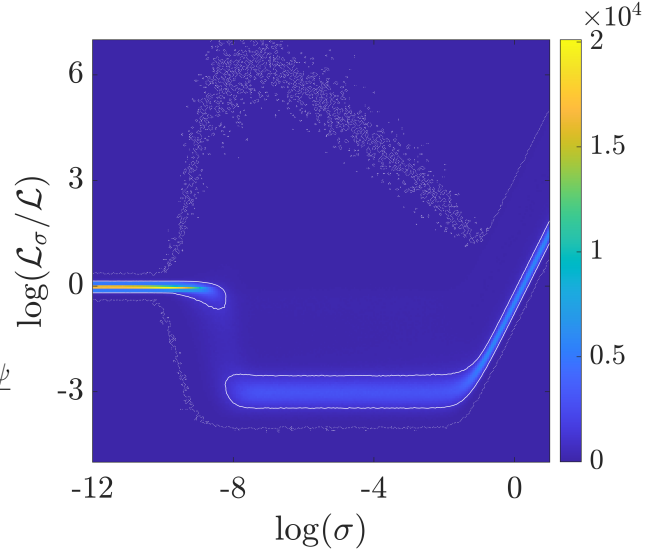


Fig. 7. Heat map of the ratio of losses $\mathcal{L}_\sigma = \sum_t x_t^2(\sigma)$ for different levels of Gaussian noise σ over the no-noise loss $\mathcal{L} = \sum_t x_t^2(0)$. For each noise level, $\sim 5 \times 10^4$ simulations were run ($\sim 10^7$ trials total). The dotted line encloses all trials, and the solid line encloses 95% of all observed trials. For the majority of trials the ratio is less than one indicating that the addition of noise is advantageous. The average loss ratio is observed to be lowest for σ in the range $10^{-8} - 10^{-2}$.

C. Datasets

C.1 Blowfly neuron H1 (Fig. 3B in the main text). Flies maintain their flight heading by utilizing visual feedback to modulate wingbeat-amplitude asymmetry. The bilaterally symmetric pair of H1 neurons (one for each direction of motion), processes information for horizontal direction (yaw) control (53). The firing of these neurons not only signals the horizontal velocity of optic flow (53) but also plays a role in controlling wingbeat asymmetry (89). In the experiment a computer controlled stepper motor was used to rotate the fly along a vertical axis while spike trains from H1 were recorded extracellularly. The fly sat immobilized in wax in a plexiglass cylinder to allow temperature to be controlled at 22 degrees Celsius.

The setup was situated outdoors in a wooded area on a bright day (49). The experiment started 42 minutes before sunset and ended 48 minutes after sunset. During this 90 minute interval, light intensities decreased by a factor of approximately 6×10^5 . The fly was subjected to a repeated yaw motion (period 10 seconds), presented 540 times. The yaw velocity was synthesized by a Markov model derived from a video recording of flies in pursuit flight, and is therefore representative of natural flight. Regression was performed separately for 4 parts of the 540 sessions: two during bright light time, one - middle, one - dark. The two bright parts were very similar, so filters are presented for the 1st, 3rd, and 4th part.

C.2 Mouse V1 pyramidal neuron (Fig. 3C in the main text). Response profiles of pyramidal neurons in primary visual cortex were measured using patch clamp recordings in brain slices from adult mice between the ages of P45 to P70 (50). The specific neuron analyzed here was from a fluorescently-labeled neuron from a mouse expressing tdTomato in a Cre-dependent manner in a Cux2-Cre transgenic mouse line, which labels excitatory neurons in cortical layers 2/3 and 4.

Voltage was recorded while injecting a stimulus composed of pink noise riding on top of square wave pulses. Pulses were 3 seconds long at rheobase (“Low μ ”) or 1.5 times rheobase (“High μ ”). Pink noise with a coefficient of variation equal to 0.2 was generated using 2 different seeds. Each stimulus was applied to the neuron 4 times for a total of 8 trials of Low and High μ . Recordings were initially performed at a sampling rate of 200 kHz. Action potentials were detected using standard approaches and the binary spiking activity was downsampled to 1 kHz for analysis.

C.3 Salamander retinal ganglion cell (Fig. 3D in the main text). In response to the visual stimulation, retinal ganglion cells generate reproducible spike trains which can be predicted using a combination of feedforward and feedback temporal filters (54, 55). Retinal ganglion cells from larval tiger salamander retina were extracellularly recorded from freshly dissected retina pressed onto a multi-electrode array.

Movies were presented to the retina at 60 frames/second and voltages were recorded at 10 kHz. The “Bar” stimulus consisted of a black bar moving against a gray background according to the Brownian motion of a particle bound by a spring to the center of the display. This was repeated 62 times, with each trial lasting 8 seconds. The “Fish” stimulus consisted of a 18-second clip of a fish swimming in a tank with swaying plants in the background. This stimulus was repeated 102 times. Action potentials were detected using standard approaches and the binary spiking activity was downsampled to 1 kHz for analysis.

C.4 Drosophila olfactory receptor neuron (Fig. 3E in the main text). Responses of olfactory receptor neurons (ORN) in Drosophila was studied by stimulating them with odors of varying mean concentration (bias) and variance of the stochastically varying component (51). Neurons were recorded extracellularly from the ab3 sensillum of Drosophila. The neuron analyzed here was an ab3A ORN.

A stimulus of fluctuating ethyl acetate was delivered to the fly antenna by blowing air over pure monomolecular odorants in liquid phase. The gas concentration was determined by the flow rate of air. For the recordings analyzed here, the variance of the signal changed every 5 seconds around a constant mean, with the variance switching a total of 60 times for 30 trials of “High σ ” and 30 trials of “Low σ ” stimulation. These stimuli were Gaussian distributed, generated by optimizing control signals to Mass Flow Controllers which regulated airflows. The mean stimulus corresponded to an odorant flux of approximately $2.5 \mu \text{ mol/s}$. The “High σ ” condition corresponded to a standard deviation of approximately $0.75 \mu \text{ mol/s}$, while the “Low σ ” condition corresponded to a standard deviation of approximately $0.3 \mu \text{ mol/s}$. Identified action potentials were binarized and downsampled to 1 kHz for analysis. Odorant concentrations were also recorded at 1 kHz.

C.5 Rat somatosensory pyramidal neuron (Fig. 3F in the main text). Response profiles of pyramidal neurons in rat somatosensory cortex were measured using patch clamp recordings in brain slices (52). The neuron analyzed here was recorded in current-clamp whole cell configuration from the soma of a layer 5 regular spiking pyramidal cell. Data were recorded at 5 kHz and low-pass filtered at 2.5 kHz.

Current was modeled by an Ornstein-Uhlenbeck process with different resulting means μ and standard deviations σ . Low μ corresponded to a mean of approximately 700 pA. Low σ corresponded to a standard deviation of approximately 200 pA. High μ corresponded to approximately 1800 pA. High σ corresponded to 300-400 pA.

D. Filter Fitting Procedure

D.1 Data Preparation. First, spiking vector $\mathbf{U} = [u(1) \dots u(T)]$ and stimulus vector $\mathbf{Y} = [y(1) \dots y(T)]$ were z-scored across time to have 0 mean and unit variance. Lag matrices were then constructed by stacking n -dimensional lag vectors together for the whole length of the dataset, T :

$$\mathbf{U}_{\text{lag}} = \begin{bmatrix} u(1) & u(2) & \dots & u(T-n+1) \\ u(2) & u(3) & \dots & u(T-n+2) \\ \vdots & \vdots & \ddots & \vdots \\ u(n) & u(n+1) & \dots & u(T) \end{bmatrix}$$

where \mathbf{U}_{lag} has dimensions $n \times (T-n+1)$. Similarly:

$$\mathbf{Y}_{\text{lag}} = \begin{bmatrix} y(1) & y(2) & \dots & y(T-n+1) \\ y(2) & y(3) & \dots & y(T-n+2) \\ \vdots & \vdots & \ddots & \vdots \\ y(n) & y(n+1) & \dots & y(T) \end{bmatrix}$$

where \mathbf{Y}_{lag} has dimensions $n \times (T-n+1)$.

D.2 Dimensionality Reduction. To avoid overfitting and promote smoothness of the filters, we projected the matrices \mathbf{U}_{lag} and \mathbf{Y}_{lag} onto a low-dimensional subspace. The stimulus matrix \mathbf{Y}_{lag}

was projected onto the top N principal components, where N is the number of principal components needed to capture 75% of the variance of the data. If the singular value decomposition of $\mathbf{Y}_{\text{lag}} = \mathbf{W}\Sigma\mathbf{V}^\top$, then define $\mathbf{S} = \Sigma\mathbf{V}^\top$, and define \mathbf{Y}_{PCA} as the first N rows of \mathbf{S} , i.e., \mathbf{Y}_{PCA} has dimensions $N \times (T - n + 1)$.

The spiking matrix \mathbf{U}_{lag} was dimensionally reduced using Laguerre polynomials, a family of orthogonal polynomials under the exponential kernel. The first few Laguerre polynomials, L_l , are:

l	$L_l(z)$
0	1
1	$-z + 1$
2	$\frac{1}{2}(z^2 - 4z + 2)$
3	$\frac{1}{6}(-z^3 + 9z^2 - 18z + 6)$
4	$\frac{1}{24}(z^4 - 16z^3 + 72z^2 - 96z + 24)$
5	$\frac{1}{120}(-z^5 + 25z^4 - 200z^3 + 600z^2 - 600z + 120)$

To define a natural time course for the filters and to ensure orthogonality, the Laguerre polynomials were each multiplied by a decaying exponential. Also a timescale τ was introduced, producing functions $\Lambda_{l,\tau}(x) = L_l(x/\tau) \exp(-x/2\tau)$. These functions were calculated on values of $x = 0, 1, \dots, n - 1$. For any given timescale τ and number of polynomials p we assembled a $n \times p$ matrix $\mathbf{\Lambda}$, where $\Lambda_{i,j} = \Lambda_{i,\tau}(j - 1)$. We can now define $\mathbf{U}_{\text{LGR}} = \mathbf{\Lambda}^\dagger \mathbf{U}_{\text{lag}}$, where \mathbf{U}_{LGR} has dimensions $p \times (T - n + 1)$ and \dagger signifies Moore-Penrose pseudo-inverse.

D.3 Computing Filters. Regressor weights were then calculated as $[\tilde{\mathbf{K}}_{ff} \tilde{\mathbf{K}}_{fb}] = \mathbf{U}\tilde{\mathbf{X}}^\top(\tilde{\mathbf{X}}\tilde{\mathbf{X}}^\top)^{-1}$, where $\tilde{\mathbf{X}} = [\mathbf{Y}_{\text{PCA}}^\top \mathbf{U}_{\text{LGR}}^\top]^\top$, and $\tilde{\mathbf{X}}$ has dimensions $(N+p) \times (T-n+1)$, $\tilde{\mathbf{K}}_{ff}$ has dimensions $1 \times N$, and $\tilde{\mathbf{K}}_{fb}$ has dimensions $1 \times p$. Then the time-dependent kernels were calculated as:

$$\mathbf{K}_{ff} = \tilde{\mathbf{K}}_{ff} \mathbf{W}(1 : N, :)$$

$$\mathbf{K}_{fb} = \tilde{\mathbf{K}}_{fb} \mathbf{\Lambda}^\dagger$$

D.4 Parameter Search. The number of Laguerre polynomials, p , and the timescale, τ , were chosen to minimize the reconstruction error of the fit using a grid search, with the number of polynomials restricted between 2 and 7 and τ chosen from $\{1, 2, 4, 8, 16, 32, 64\}$ ms.

For each parameter set, a 2-fold cross-validation was performed by splitting the data into equal-sized halves. The kernels were computed from data in the first half, and the reconstruction error was calculated for the second half. The procedure was then repeated by estimating from the second half and reconstructing the first. The reconstructed spiking vector was computed as $[\tilde{\mathbf{K}}_{ff} \tilde{\mathbf{K}}_{fb}]\tilde{\mathbf{X}}$.

Reconstruction error was defined as:

$$\text{Error} = \frac{\|\mathbf{U} - [\tilde{\mathbf{K}}_{ff} \tilde{\mathbf{K}}_{fb}]\tilde{\mathbf{X}}\|_2^2}{\|\mathbf{U}\|_2^2}$$

This estimation procedure was done individually for each trial. The filters plotted in Fig. 3 in the main text represent the mean across all trials, with dotted lines indicating mean plus or minus the standard error of the mean. The number and length of trials for each dataset in Fig. 3 in the main text were as follows:

E. Control Analyses

E.1 Mean Firing Rate. In the fly H1 dataset spike rates decrease considerably in darker conditions, so we performed supplementary analyses to rule out the possibility that the change in filter shape is simply a consequence of a different firing rate. We repeated the

	Trials	Trial Length	Time Step	Maximum Shift*
2B (Dark)	135	5000	2 ms	100 data points
2B (Mid)	135	5000	2 ms	100 data points
2B (Light)	135	5000	2 ms	100 data points
2C (Low)	8	3000	1 ms	100 data points
2C (High)	8	3000	1 ms	100 data points
2D (Bar)	62	8270	1 ms	200 data points
2D (Fish)	102	17705	1 ms	200 data points
2E (Low)	30	5000	1 ms	200 data points
2E (High)	30	5000	1 ms	200 data points
2F (Each Condition)	8	2375	1 ms	150 data points

Table 1. * Maximum Shift refers to the maximum number of time steps that the stimulus is shifted when constructing the lagged stimulus matrix.

analysis used in Fig. 3 in the main text for bright segment of the dataset while keeping only those data where the number of spikes in the lag vector were below a threshold (100 or 200 spikes/s). Similarly, for the dark segment, the analysis was repeated while keeping data points where the number of spikes in the lag vector was above some threshold (100 or 200 spikes/s). As shown in Fig. 8, the filter shape is determined primarily by belonging to the dark/bright segment rather than by the firing rate. Therefore, this control supports our assertion that the filter shapes adapt to the stimulus statistics.

E.2 Sunrise and Sunset. A possible explanation for the change in feedforward and feedback filters in the fly H1 dataset is that the effect was due to the elapsed recording time. To rule this out, we repeated the analysis on a separate H1 recording performed during the time around sunrise rather than sunset. The response profiles were remarkably similar for periods of the experiment with comparable light conditions, Fig. 9, indicating that the change in filter shape is primarily due to an effect of environmental brightness rather than recording time.

E.3 Effect of Z-Scoring Spike Trains. Spike trains with different mean firing rates will be affected differently by z-scoring (the first step in Data Preparation). Time bins with spikes will take on larger values at low rates than at high rates. This disparity in the magnitude of the non-zero components of the spike train may affect the shape or scaling of the estimated filters. To control for this, we re-ran the fitting procedure for all datasets while not z-scoring the spiking vector \mathbf{U} , Fig. 10. The filter shapes and differences across conditions are extremely similar to those obtained with z-scoring, aside from a scaling factor for the feedforward filters.

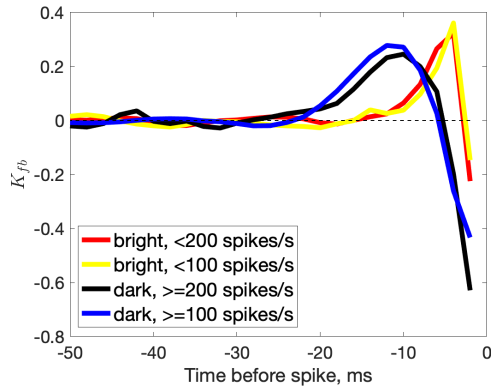


Fig. 8. Blowfly H1 feedback filter shape is primarily determined by the statistics of sensory stimuli (background luminance) rather than the H1 firing rate. Even when we analyze data subsets from the bright background experiment with firing rate lower than that in data subsets from the dark background experiment, the difference between the filters is qualitatively similar to Fig. 3B in the main text.

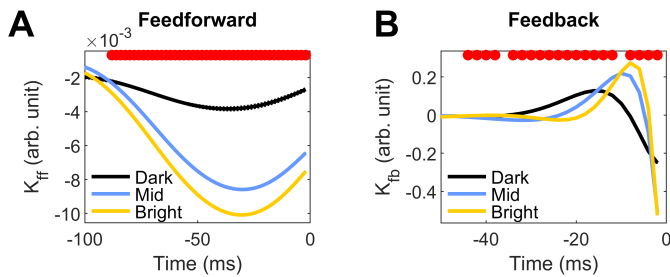


Fig. 9. Changes in Blowfly H1 feedforward and feedback filters are due to changes in environmental brightness, not overall recording time. The above filters were obtained from a dataset recorded during sunrise, in which the Dark condition came first, followed by Mid and then Bright. This is in contrast to the analyses in Fig. 3B which was recorded during sunset, in which Bright came first, followed by Mid and then Dark. The changes in the filters are qualitatively the same between these conditions, indicating that this is driven by the environmental brightness and not an artifact of overall recording time.

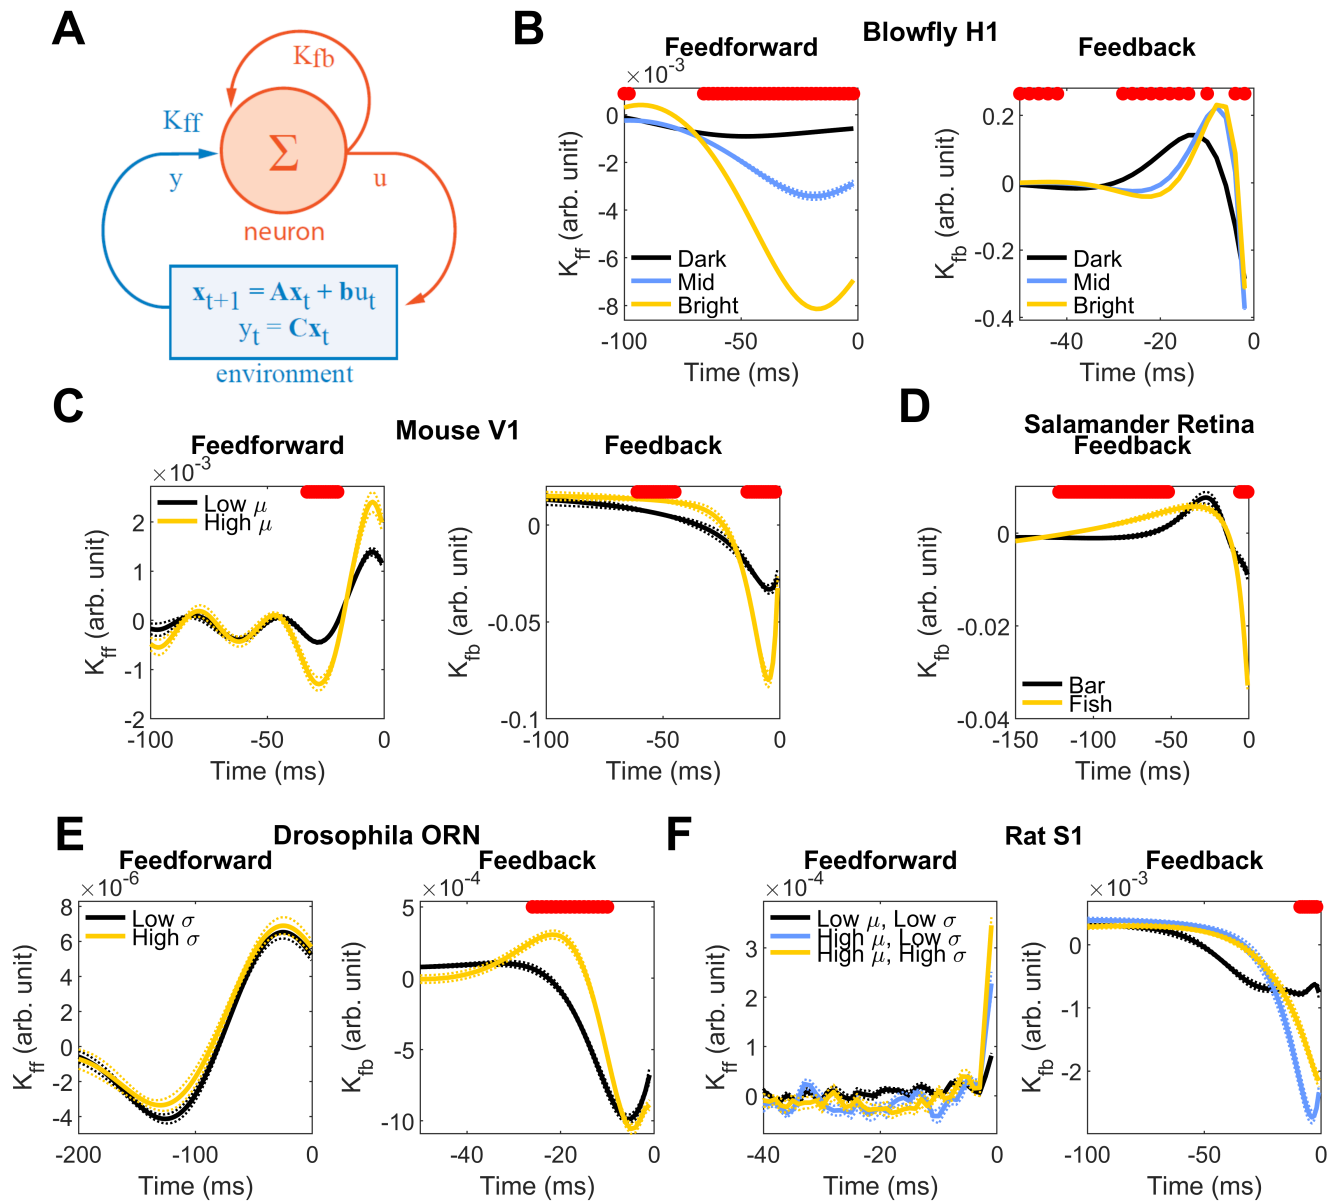


Fig. 10. Estimated filters without z-scoring the spiking vector $u(t)$ are qualitatively similar to the ones with z-scoring, Fig. 3, in the main text. Naming and coloring conventions are as in Fig. 3 in the main text. Feedforward filters in *F* were not substantially different from those obtained using the Z-scored data.