



The neuron as a direct data-driven controller

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

Affiliations are included on p. 8.

Edited by Yuhai Tu, International Business Machines Corp, Yorktown Heights, NY; received August 22, 2023; accepted April 12, 2024, by Editorial Board Member Herbert Levine

In the quest to model neuronal function amid 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 current normative models, which primarily optimize prediction, by conceptualizing neurons as optimal feedback controllers. We posit that neurons, especially those beyond early sensory areas, steer their environment toward 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. To model neurons as biologically feasible controllers which implicitly identify loop dynamics, infer latent states, and optimize control we utilize the contemporary direct data-driven control (DD-DC) framework. Our DD-DC neuron model explains various neurophysiological phenomena: the shift from potentiation to depression in spike-timing-dependent plasticity 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 modern, 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 incomplete 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 de-inactivation 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 is 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 premotor 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.

Significance

Building upon the efficient coding and predictive information theories, we present a perspective that neurons not only predict but may also actively influence their future inputs through their outputs. We model neurons as feedback controllers of their environments, a role traditionally considered computationally demanding, particularly when the dynamical system characterizing the environment is unknown. By harnessing an advanced data-driven control framework, we illustrate the feasibility of biological neurons functioning as effective feedback controllers. This innovative approach enables us to coherently explain various experimental findings that previously seemed unrelated. Our research has multiple potential implications, from the modeling of neuronal circuits to enabling biologically inspired artificial intelligence systems.

Author contributions: R.R.d.R.v.S. and D.B.C. designed research; J.J.M., A.G., M.T., R.R.d.R.v.S., and D.B.C. performed research; J.J.M., A.G., M.T., J.L.P.-S., and R.R.d.R.v.S. contributed new reagents/analytic tools; J.J.M. and A.G. analyzed data; J.J.M., A.G., and D.B.C. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. Y.T. is a guest editor invited by the Editorial Board.

Copyright © 2024 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: mitya@flatironinstitute.org.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2311893121/-/DCSupplemental>.

Published June 24, 2024.

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. 1*A*. 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 appropriate for the whole brain (19–21) may seem daunting for a single 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. 1*A*) from historical observations and control signals (22). These parameters form the basis for deriving a control law that optimizes specific objectives, like optimal or robust control (23). When dealing with low-dimensional or noisy observations, the control law must be based not only on immediate observations, **y**, but on an estimated state, $\hat{\mathbf{x}}$, (23) derived from past data (see output feedback control and Kalman filtering) (Fig. 1*A*). Although, for linear dynamics, the above tasks have known solutions (22, 23), 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 contemporary direct data-driven control (DD-DC) framework (24–26). 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 autoregressive moving average (ARMA) process.

Conceptualizing neurons as controllers in general and modeling them as DD-DCs in particular provides insights into multiple seemingly unrelated experimental observations. First, it can explain the potentiation/depression transition in spike-timing-dependent plasticity (STDP) and its asymmetry (Fig. 1*D*) (27–29). Second, it can account for the temporal extent (noninstantaneous nature) of feedforward (STA) and feedback (spike-history dependent) filters and their adaptation to input statistics (Fig. 3 *B–F*). Third, it explains the loss of temporal precision in the neuronal spike-generation mechanism under constant input (Fig. 4, *Right*) (14). Fourth, the operation of DD-DC in the online setting requires variability and/or noise (Fig. 2, *Right*), which is consistent with many neurophysiological observations (30–36). Finally, viewing neurons as controllers is consistent with the observations of movement-related activity throughout much of the brain including traditionally sensory areas (37, 38).

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 (39, 40). This neuronal unit is overly simplistic in that it lacks internal feedback and temporal dynamics (for more details, see Relationship to other work). Therefore, our proposed DD-DC model of the neuron could serve as an alternative foundational building block for constructing biologically inspired artificial neural networks.

A preliminary publication of this work has appeared in the abstract form (41, 42).

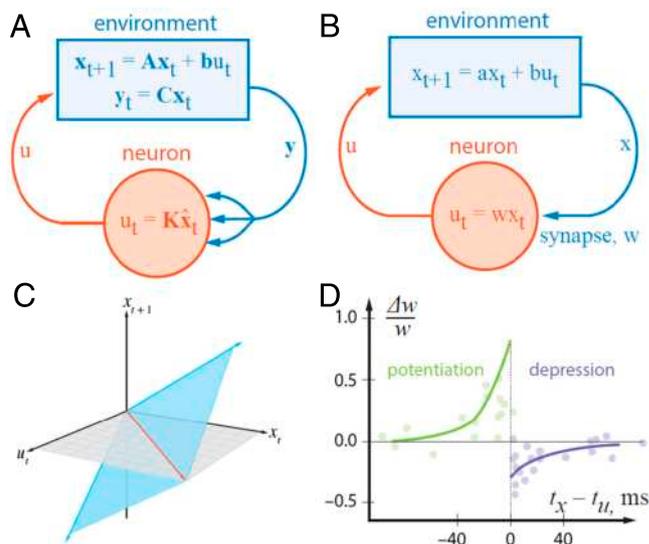


Fig. 1. (A) A schematic representation of the neuron as a feedback controller in a closed loop. (B) A scalar fully observed dynamical system controlled by tuning the weight of a synapse, w , in the control law. (C) The subspace of valid pairings of observations and controls (blue plane) is spanned by 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) STDP: the relative change in the synaptic weight, $\Delta w/w$, vs. the time interval between the pre- and postsynaptic spikes, $t_x - t_y$, showing the potentiation (causal) and depression (anticausal) windows (27–29). Adapted from ref. 28. Copyright 1998 Society for Neuroscience.

The DD-DC Framework

In this section, we provide an overview of the DD-DC framework (24–26). 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. 1*A*):

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

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

where $\mathbf{x}_t \in \mathbb{R}^n$ represents the latent state of the environment at time t , u_t is the control signal generated by the neuron, and \mathbf{y}_t is the observation received by the neuron. In the realm of model-based control, the dynamics parameters **A**, **b**, and **C** are typically predefined, an unrealistic presumption in a biological context. We consider the system Eqs. 1 and 2 to be fully controllable and observable. In some cases, the optimal control signal is linearly related to the estimated latent state variable, $\hat{\mathbf{x}}$,

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

When the parameters **A**, **b**, and **C** are unknown to the controller, the optimal control signal can be constructed using DD-DC. In this case, control signals are generated directly from the observations, bypassing an explicit representation of the system dynamics and the latent state. The DD-DC controller learns the mapping from observations to control signals from the history of such pairings. 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. 1*C*). This intuition is formalized by Willems' fundamental lemma (24), which posits that each observation-control pairing (in the

fully observed scenario) can be expressed as a linear combination of k historical pairings ($1, 2, \dots, k < t$),

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

assuming the matrix composed of \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 (24), DD-DC has recently been expanded to accommodate noisy observations, nonlinear dynamics, online applications, and limited datasets (26, 43–46) 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 original analysis.

Basic DD-DC Accounts for 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 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, in place of \mathbf{K} from Eq. 3 because it is naturally implemented by the synaptic weight, 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, $w^* = -a/b$.

As argued above, 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 compute \mathbf{g} with a minimum l_2 -norm found 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 (47). 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 the neurophysiological context, x and u in Eq. 12 symbolize pre- and postsynaptic neuronal activities, respectively, with nonzero values during spikes. Consequently, the sums in Eq. 12 accrue contributions solely when pre- and postsynaptic spikes are temporally proximate, 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 (27–29) (Fig. 1D).

Whereas the potentiation window of STDP can be viewed as the extension of the Hebbian rule (48) to the temporal interplay between spikes, the rationale behind the relatively narrow depression window in STDP has remained elusive, largely due to its seemingly anticausal nature. Current explanations of this phenomenon, e.g., refs. 49 and 50, rely on ad hoc assumptions. In contrast, our work shows that the apparent anticausal aspect of STDP is a natural outcome of conceptualizing neurons as feedback controllers. Specifically, the absence or presence of a presynaptic spike following a postsynaptic spike conveys information to the neuron about the effectiveness or ineffectiveness of its control of the environment. Thus, what is initially perceived as an anticausal 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 a nonzero control cost, $r > 0$, which in the scalar, one-step time-horizon case is given by:

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

The full derivation of this solution is given in *SI Appendix, section A*. One can see that it aligns with the optimal LQR gain:

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

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 $(\|\mathbf{U}\|^2 \|\mathbf{X}\|^2 - (\mathbf{UX}^\top)^2)$, thus reducing it to Eq. 14. This solution provides a potential framework to interpret the variance in STDP profiles documented in various studies (51) 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 investigate the functioning of the DD-DC LQR controller through numerical simulations aimed at stabilizing a potentially unstable scalar dynamical system, Eq. 5. Initially, we operate 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 compute the controller gain, w , by integrating the recorded values of u and x into the general LQR solution, Eq. 13. Following this initialization, we transition to a closed-loop operation of the DD-DC LQR controller, recalculating w at each time step (see *SI Appendix, section B* 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 *SI Appendix, section B* for details). Initially, the controller learns and applies the optimal value of w . When the parameters a and b undergo a switch (at time = 25), x remains at zero initially, indicating no direct loss impact, Fig. 2, *Left*. To demonstrate that the controller fails to adapt, we apply a jolt $\Delta x = 0.2$ (dashed line in *Inset*) to the state variable x at time $t = 55$. This results in significant and often unstable deviations from equilibrium, Fig. 2, *Left*.

To uncover the cause behind the DD-DC controller's failure to adjust following the static control law phase, we reexamined 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 submatrix composed of the x and u rows is rank deficient, thereby contravening the persistence of excitation

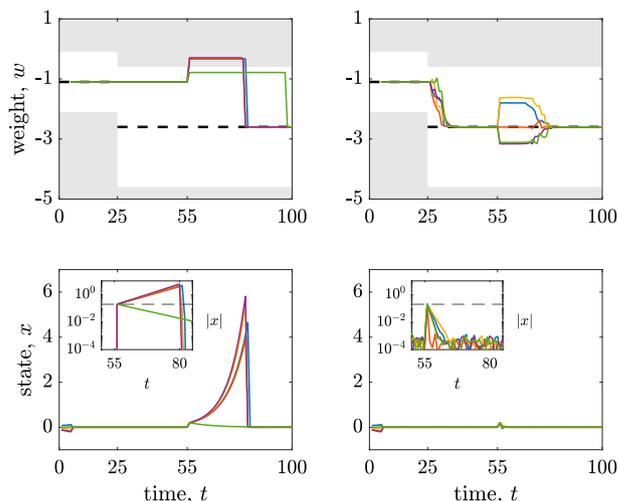


Fig. 2. The gain, w , of the online DD-DC LQR controller (*Top*) and the state variable, x , of the scalar dynamical plant (*Bottom*) as a function of time for switching system dynamics parameters. *Left:* In the absence of noise in control, the DD-DC controller fails to adapt to the changing conditions because of the loss of persistence of excitation. Each of the five colored lines represents a different simulation trial. *Right:* Adding noise to the control law Eq. 16 enables exploration that restores the persistence of excitation and performance of the controller (also see *SI Appendix, section B*). The dashed line shows optimal LQR values of weights for every time step. The white bands within the gray shaded areas in the w plots represent regions of stability and instability respectively. *Insets* plot $|x|$ on a log scale—representative of the control loss.

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 (52) 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 (46), 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 Fig. 2, *Right*, for $\text{VAR}(\eta_t) = 10^{-6}$. The addition of low-variance noise to the control signal facilitates exploration thus reestablishing persistence of excitation and overall controller performance. This is evidenced by x swiftly returning to zero and w reverting to its optimal state. 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 *SI Appendix, section B* for details).

There are multiple empirical observations of noise that could play such a role within the brain. Locally, such noise might originate from the inherent unreliability of synaptic transmission (31). On a broader scale, the brain encompasses specific circuits and neurons dedicated to introducing noise and/or variability, as evidenced in studies on songbirds (33) and *Caenorhabditis elegans* (32). Additionally, the brain can introduce variability in sensory inputs through the generation of corresponding motor outputs, exemplified by microsaccades (53). Given these mechanisms, the

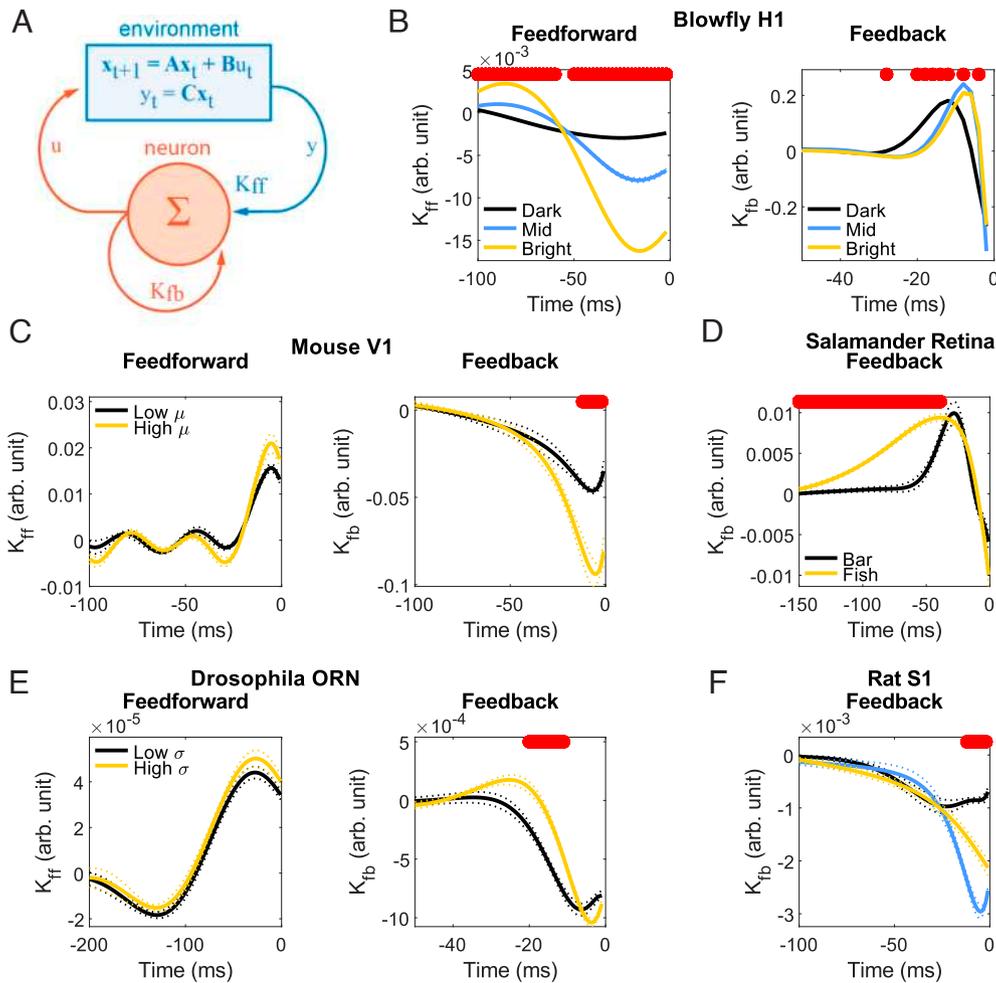


Fig. 3. (A) Illustration of the neuron modeled as an ARMA controller, characterized by feedforward, \mathbf{K}_{ff} , and feedback, \mathbf{K}_{fb} , temporal filters. (B–D) Adaptation of experimentally measured temporal filters (depicted in black, yellow, and blue) to input signal statistics. Solid lines represent mean values, while thin dotted lines denote standard errors of the mean. Regions where differences are statistically significant ($P < 0.05$, 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 (54), responding to visual motion against different background luminance levels. (C) Feedforward and feedback filters in pyramidal cells from the mouse primary visual cortex (55) 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 (meta data for the feedforward filter is unavailable, see *SI Appendix*). (E) Adaptation of feedforward and feedback filters in a *Drosophila* olfactory receptor neuron (ORN) (56) to odorant concentrations with varying variances. (F) Feedback filters in rat somatosensory cortex pyramidal neurons (57), responding to current injections modulated by an Ornstein–Uhlenbeck process atop a DC component. Feedforward filters are provided in *SI Appendix*.

operational variability observed in neural representations (34, 35) 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, \mathbf{x} . For linear systems, this latent state can be inferred from recent sequences of observations and controls using time-delay embedding techniques (25),

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

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 ARMA model of the neuron (Fig. 3A).

In this section, rather than optimizing feedforward and feedback temporal filters, we estimate them from experimental data. In these experiments, neurons are isolated from the loop and stimulated with sensory input or injected current, y , and the neuronal response, u , is recorded. These data are compiled into matrices $\mathbf{U} = [u_1 \quad \dots \quad u_t]$ and $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1 \quad \dots \quad \hat{\mathbf{x}}_t]$, which are linearly related:

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

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

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

Utilizing Eq. 20, we reconstruct the feedforward, \mathbf{K}_{ff} , and feedback, \mathbf{K}_{fb} , temporal filters from experimental data across various model systems, Fig. 3 (see *SI Appendix, section C* for details).

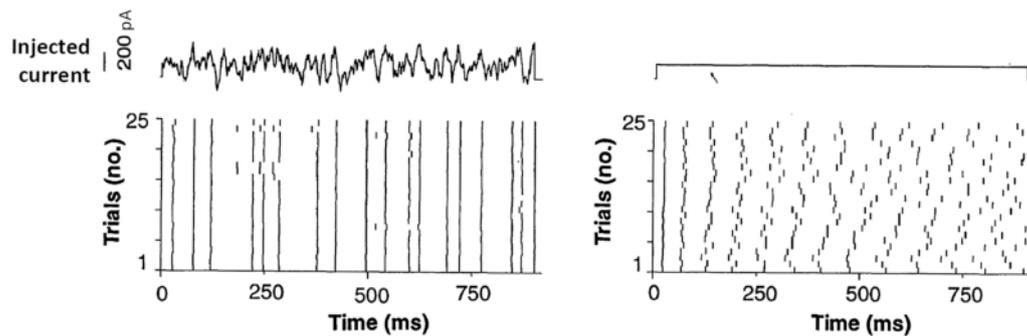


Fig. 4. High-variance current injections into a neuron yield remarkably consistent spike trains over multiple trials, showcasing the precision of the spike-generation mechanism. *Right:* In contrast, a constant current input leads to notably variable spike trains, revealing a significant reduction in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

While these temporal responses have been previously measured (58–60) and interpreted as optimal feedforward filters, our controller-based perspective offers a different 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 (54), mouse V1 pyramidal neurons responding to different mean injected current waveforms (55), salamander retinal ganglion cells exposed to distinct visual stimuli (8), *Drosophila* olfactory receptor neurons reacting to varying odorant concentrations (56), and pyramidal neurons in the rat somatosensory cortex stimulated with current injections of diverse means and variances (57). 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). The adaptation of feedback filters to changing stimulus statistics demonstrated by our analysis is not predicted by existing theories. This calls for a 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 (24). 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 100 ms 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 (61). This noise was thought to be inconsequential for spike timing in the presence of highly variable inputs, as its effects would be overshadowed by the abundance of open ion channels. However, under constant input, even slight variations in injected current ($STD \approx 50$ pA) are observed to restore spike-timing precision (57), 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 ref. 14, and correlating it with the condition number of the time-delay covariance.

Discussion

The power of the proposed DD-DC model of the neuron is in that, starting from a single postulate, it offers explanations for multiple previously 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 circumstantial evidence, their multitude and variety provide strong support for the DD-DC model. This perspective has the potential to deepen and refine our 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

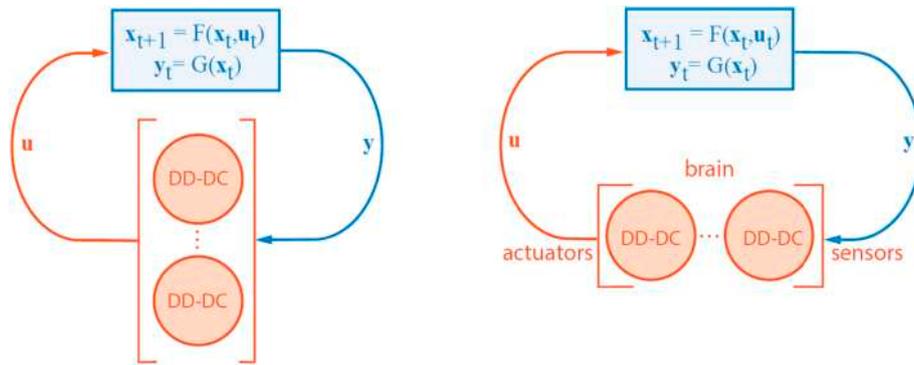


Fig. 5. *Left:* Illustration of controlling a nonlinear dynamical system using multiple switching DD-DCs. *Right:* Depiction of a deep network model where each neuron exerts control over its immediate environment, contributing to the broader control exerted by the entire brain over the external environment.

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 (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 (62–64), have other operational advantages similar to widely used bang-bang control (65). Deriving such a controller may help model neurons with active conductances and spikes (66) 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 (25), 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 (67, 68), 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 (69). 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 the 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*D*), pre- and postsynaptic 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, 70). At the other extreme, the plasticity of some synapses in the cerebellum and the hippocampus peaks when the spikes lag by tens of milliseconds (71, 72). This suggests longer loops involving different brain regions or even the external environment (73). 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 (multisynaptic and transenvironment) loops may seem unrealistic. However, theoretical analysis (74, 75) and experimental observations seem to support long-range propagation of signals from individual neurons (76). Specifically, rodents can be trained to behaviorally report single-neuron electrical stimulation in the barrel cortex (77) suggesting that the spikes of a single neuron can make an impact sufficiently far downstream to elicit behavior. Also, stimulation of single neurons in the motor cortex evoke whisker movements (78), 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 (37, 38). 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 (36). 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.

Normalization of neuronal responses, which have been both experimentally observed and theoretically modeled (79–81), resembles the feedback filter in our model. However, normalization 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 (82). Although these networks are predominantly feedforward and do not allow neuron outputs to modify network inputs, both concepts emphasize learning generative dynamics. Moreover, our approach allows for unstable, partially observed dynamical systems.

Data-driven control in network contexts, including brain networks, has previously been investigated (83). 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) (60), 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 (84), laying groundwork for future connections between these theories. The intraneuronal feedback loop similar to ours has been previously used to combine encoder and decoder functions in the same neuron (85). In that model, an integrate-and-fire neuron spikes when the error between its input and reconstructed signal crosses an adapting spike threshold. This is similar in spirit to our approach in that a neuron spikes when it is attempting to align its input to a target value. Our model adds the interpretation that the neuron is not just reporting the error but acting in an effort to modify the incoming input.

Differing from the conventional McCulloch–Pitts–Rosenblatt (MCPR) unit in artificial neural networks, the DD-DC neuron integrates inputs over time and features an autoregressive loop, unlike the instantaneous response of the MCPR units. Also, in contrast to network-wide optimization in artificial neural networks, the DD-DC model optimizes objectives at the neuronal level. The DD-DC neuron uses its inputs as a teaching signal as opposed to biologically implausible error backpropagation in artificial neural networks.

Neurons are sometimes conceptualized as agents in the reinforcement learning (RL) paradigm, (86, 87). 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 (19–21), and used LQR with delays/noise to model internal feedback (88). 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 (89, 90). Earlier models separated sensory system identification and motor control (91), but our unified neuron-as-controller model eliminates the need to match corresponding sensory and motor units.

Our results may apply to cell types other than neurons and allow modeling their function as feedback controllers of their environments. This view is natural in the light of evolution as single-cell organisms must act as feedback controllers to survive. Future research could explore how coherent controller actions emerge in self-organized networks of neuron (or cell) controllers.

Data, Materials, and Software Availability. Previously published data were used for this work (8, 54–57). Analysis code available at <https://zenodo.org/records/11399185> (92).

ACKNOWLEDGMENTS. We would like to express our sincere gratitude to Anirvan Sengupta, Luca Mazzucato, Christian Pehle, Zachary Mainen, and Subham Dey for their invaluable discussions and insights. We thank Stephanie Palmer, Benjamin Hoshal, Thierry Emonet, and Giancarlo La Camera for providing datasets and technical expertise. Our special thanks to Lucy Reading-Ikkanda for her exceptional assistance with the illustrations. Additionally, we acknowledge the use of ChatGPT in the writing process to improve wording and grammar.

Author affiliations: ^aNeuroscience Institute, New York University Grossman School of Medicine, New York City, NY 10016; ^bCenter for Computational Neuroscience, Flatiron Institute, New York City, NY 10010; and ^cPhysics Department, Indiana University, Bloomington, IN 47405

- H. Barlow, Redundancy reduction revisited. *Network* **12**, 241 (2001).
- M. V. Srinivasan, S. B. Laughlin, A. Dubs, Predictive coding: A fresh view of inhibition in the retina. *Proc. R. Soc. London B* **216**, 427–459 (1982).
- J. J. Atick, Could information theory provide an ecological theory of sensory processing? *Network* **3**, 213–251 (1992).
- J. H. van Hateren, Theoretical predictions of spatiotemporal receptive fields of fly LMCS, and experimental validation. *J. Comp. Physiol. A* **171**, 157–170 (1992).
- E. P. Simoncelli, B. A. Olshausen, Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
- R. P. N. Rao, D. H. Ballard, Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
- N. Y. Jun, G. Field, J. Pearson, Efficient coding, channel capacity, and the emergence of retinal mosaics. *Adv. Neural Inf. Process. Syst.* **35**, 32311–32324 (2022).
- S. E. Palmer, O. Marre, M. J. Berry, W. Bialek, Predictive information in a sensory population. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 6908–6913 (2015).
- N. C. Rust, S. E. Palmer, Remembering the past to see the future. *Annu. Rev. Vis. Sci.* **7**, 349–365 (2021).
- S. Wang, I. Segev, A. Borst, S. Palmer, Maximally efficient prediction in the early fly visual system may support evasive flight maneuvers. *PLoS Comput. Biol.* **17**, e1008965 (2021).
- N. Tishby, F. C. Pereira, W. Bialek, The information bottleneck method. arXiv [Preprint] (2000). <https://doi.org/10.48550/arXiv.physics/0004057> (Accessed 29 May 2024).
- M. Chalk, O. Marre, G. Tkačik, Toward a unified theory of efficient, predictive, and sparse coding. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 186–191 (2018).
- H. L. Bryant, J. P. Segundo, Spike initiation by transmembrane current: A white-noise analysis. *J. Physiol.* **260**, 279–314 (1976).
- Z. F. Mainen, T. J. Sejnowski, Reliability of spike timing in neocortical neurons. *Science* **268**, 1503–1506 (1995).
- A. J. Bell, Levels and loops: The future of artificial intelligence and neuroscience. *Philos. Trans. R. Soc. London B, Biol. Sci.* **354**, 2013–2020 (1999).
- S. Song, P. J. Sjöström, M. Reigl, S. Nelson, D. B. Chklovskii, Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* **3**, e68 (2005).
- L. R. Varshney, B. L. Chen, E. Paniagua, D. H. Hall, D. B. Chklovskii, Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS Comput. Biol.* **7**, e1001066 (2011).
- M. Winding *et al.*, The connectome of an insect brain. *Science* **379**, eadd9330 (2023).
- P. Cisek, Beyond the computer metaphor: Behaviour as interaction. *J. Conscious. Stud.* **6**, 125–142 (1999).
- P. Cisek, J. F. Kalaska, Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
- E. Ahissar, E. Assa, Perception as a closed-loop convergence process. *eLife* **5**, e12830 (2016).
- T. Katayama *et al.*, *Subspace Methods for System Identification* (Springer, 2005), vol. 1.
- K. J. Åström, R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers* (Princeton University Press, 2021).
- J. C. Willems, P. Rapisarda, I. Markovsky, B. L. M. De Moor, A note on persistency of excitation. *Syst. Control Lett.* **54**, 325–329 (2005).
- C. De Persis, P. Tesi, Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Trans. Autom. Control* **65**, 909–924 (2019).
- I. Markovsky, L. Huang, F. Dörfler, Data-driven control based on the behavioral approach: From theory to applications in power systems. *IEEE Control Syst. Mag.* **43**, 28–68 (2022).
- H. Markram, J. Lübke, M. Frotscher, B. Sakmann, Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).
- G. Bi, M. Poo, Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* **18**, 10464–10472 (1998).
- L. I. Zhang, H. W. Tao, C. E. Holt, W. A. Harris, M. Poo, A critical window for cooperation and competition among developing retinotectal synapses. *Nature* **395**, 37–44 (1998).
- M. E. Rule, T. O'Leary, C. D. Harvey, Causes and consequences of representational drift. *Curr. Opin. Neurobiol.* **58**, 141–147 (2019).
- L. E. Dobrunz, C. F. Stevens, Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron* **18**, 995–1008 (1997).

32. A. Gordus, N. Pokala, S. Levy, S. W. Flavell, C. I. Bargmann, Feedback from network states generates variability in a probabilistic olfactory circuit. *Cell* **161**, 215–227 (2015).
33. B. P. Ölveczky, A. S. Andalman, M. S. Fee, Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS Biol.* **3**, e153 (2005).
34. D. Deitch, A. Rubin, Y. Ziv, Representational drift in the mouse visual cortex. *Curr. Biol.* **31**, 4327–4339 (2021).
35. C. E. Schoonover, S. N. Ohashi, R. Axel, A. J. P. Fink, Representational drift in primary olfactory cortex. *Nature* **594**, 541–546 (2021).
36. C. M. Niell, M. P. Stryker, Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
37. S. Musall, M. T. Kaufman, A. L. Juavinett, S. Gluf, A. K. Churchland, Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
38. E. Zagha *et al.*, The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. *J. Neurosci.* **42**, 1375–1382 (2022).
39. W. S. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133 (1943).
40. F. Rosenblatt, "Principles of neurodynamics. Perceptrons and the theory of brain mechanisms" (Tech. Rep., Cornell Aeronautical Lab Inc., Buffalo, NY, 1961).
41. J. J. Moore, R. de Rob, R. van Steveninck, D. B. Chklovskii, A. Genkin, "A neuron as a direct data-driven controller (DD-DC) model of a neuron" (Computational Neuroscience Annual Meeting, 2023).
42. T. Tesileanu, A. Genkin, D. Chklovskii, M. Tournoy, J. Moore, "Spike-timing dependent plasticity (STDP) emerges from a direct data-driven controller (DD-DC) model of a neuron" (Society for Neuroscience Annual Meeting, 2023).
43. J. Coulson, J. Lygeros, F. Dörfler, "Data-enabled predictive control: In the shallows of the deep" in 2019 18th European Control Conference (ECC) (IEEE, 2019), pp. 307–312.
44. J. Berberich, J. Köhler, M. A. Müller, F. Allgöwer, Linear tracking MPC for nonlinear systems—Part II: The data-driven case. *IEEE Trans. Autom. Control* **67**, 4406–4421 (2022).
45. H. J. van Waarde, Beyond persistent excitation: Online experiment design for data-driven modeling and control. *IEEE Control Syst. Lett.* **6**, 319–324 (2021).
46. M. Rotulo, C. De Persis, P. Tesi, Online learning of data-driven controllers for unknown switched linear systems. *Automatica* **145**, 110519 (2022).
47. C. Pehlevan, D. B. Chklovskii, Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. *IEEE Signal Process. Mag.* **36**, 88–96 (2019).
48. D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (Wiley, New York, NY, 1949).
49. M. Gilson, T. Fukai, Stability versus neuronal specialization for STDP: Long-tail weight distributions solve the dilemma. *PLoS ONE* **6**, e25339 (2011).
50. P. V. Aceituno, M. T. Farinha, R. Loidl, B. F. Grewe, Learning cortical hierarchies with temporal Hebbian updates. *Front. Comput. Neurosci.* **17**, 1136010 (2023).
51. N. Caporale, Y. Dan, Spike timing-dependent plasticity: A Hebbian learning rule. *Annu. Rev. Neurosci.* **31**, 25–46 (2008).
52. R. M. Pritchard, Stabilized images on the retina. *Sci. Am.* **204**, 72–79 (1961).
53. H. Ko, M. Poletti, M. Rucci, Microsaccades precisely relocate gaze in a high visual acuity task. *Nat. Neurosci.* **13**, 1549–1553 (2010).
54. G. D. Lewen, W. Bialek, R. R. De Ruyter Van, Steveninck., Neural coding of naturalistic motion stimuli. *Network* **12**, 317 (2001).
55. C. Teeter *et al.*, Generalized leaky integrate-and-fire models classify multiple neuron types. *Nat. Commun.* **9**, 709 (2018).
56. S. Gorur-Shandilya, M. Demir, J. Long, D. A. Clark, T. Emonet, Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli. *eLife* **6**, e27670 (2017).
57. A. Rauch, G. La Camera, H.-R. Luscher, W. Senn, S. Fusi, Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *J. Neurophysiol.* **90**, 1598–1612 (2003).
58. F. Rieke, D. Warland, R. de Ruyter van Steveninck, W. Bialek, Spikes: Exploring the neural code (MIT Press, Cambridge, MA, 1999).
59. J. Keat, R. Pamela Reinagel, C. Reid, M. Meister, Predicting every spike: A model for the responses of visual neurons. *Neuron* **30**, 803–817 (2001).
60. J. W. Pillow *et al.*, Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
61. E. Schneidman, B. Freedman, I. Segev, Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural Comput.* **10**, 1679–1703 (1998).
62. J. J. B. Jack, D. Noble, R. W. Tsien, Electric current flow in excitable cells (Clarendon/Oxford University Press, New York, 1975).
63. R. Sarpeshkar, Analog versus digital: Extrapolating from electronics to neurobiology. *Neural Comput.* **10**, 1601–1638 (1998).
64. S. B. Laughlin, R. de Rob, R. van Steveninck, J. C. Anderson, The metabolic cost of neural information. *Nat. Neurosci.* **1**, 36–41 (1998).
65. M. I. Kamien, N. L. Schwartz, Dynamic optimization: The calculus of variations and optimal control in economics and management (Courier Corporation, North Chelmsford, 2012).
66. C. Koch, *Biophysics of Computation: Information Processing in Single Neurons* (Oxford University Press, 2004).
67. W. Maass, T. Natschläger, H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
68. J. M. Beggs, D. Plenz, Neuronal avalanches in neocortical circuits. *J. Neurosci.* **23**, 11167–11177 (2003).
69. H. S. Seung, How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. U.S.A.* **93**, 13339–13344 (1996).
70. Y. Zilberter, Dendritic release of glutamate suppresses synaptic inhibition of pyramidal neurons in rat neocortex. *J. Physiol.* **528**, 489 (2000).
71. J. T. Dudman, D. Tsay, S. A. Siegelbaum, A role for synaptic inputs at distal dendrites: Instructive signals for hippocampal long-term plasticity. *Neuron* **56**, 866–879 (2007).
72. A. Suvrathan, H. L. Payne, J. L. Raymond, Timing rules for synaptic plasticity matched to behavioral function. *Neuron* **92**, 959–967 (2016).
73. S. Jayabal, *et al.*, Experience adaptively tunes the timing rules for associative plasticity. bioRxiv [Preprint] (2022). <https://doi.org/10.1101/2022.11.28.518128> (Accessed 29 May 2024).
74. M. London, A. Roth, L. Beeren, M. Häusser, P. E. Latham, Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* **466**, 123–127 (2010).
75. M. Monteforte, F. Wolf, Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Phys. Rev. X* **2**, 041007 (2012).
76. J. Wolfe, A. R. Houweling, M. Brecht, Sparse and powerful cortical spikes. *Curr. Opin. Neurobiol.* **20**, 306–312 (2010).
77. A. R. Houweling, M. Brecht, Behavioural report of single neuron stimulation in somatosensory cortex. *Nature* **451**, 65–68 (2008).
78. M. Brecht, M. Schneider, B. Sakmann, T. W. Margrie, Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* **427**, 704–710 (2004).
79. M. J. Wainwright, O. Schwartz, E. P. Simoncelli, Chapter 10 Natural image statistics and divisive normalization in the book. Probabilistic models of the brain (2002), p. 203.
80. M. Carandini, D. J. Heeger, Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012).
81. N. Brenner, W. Bialek, R. de Ruyter, V. Steveninck, Adaptive rescaling maximizes information transmission. *Neuron* **26**, 695–702 (2000).
82. M. Boerlin, C. K. Machens, S. Denève, Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput. Biol.* **9**, e1003258 (2013).
83. G. Baggio, D. S. Bassett, F. Pasqualetti, Data-driven control of complex networks. *Nat. Commun.* **12**, 1429 (2021).
84. L. Lapique, Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen.* **9**, 620–635 (1907).
85. D. L. Jones, E. C. Johnson, R. Ratnam, A stimulus-dependent spike threshold is an optimal neural coder. *Front. Comput. Neurosci.* **9**, 61 (2015).
86. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, 2018).
87. P. Dayan, Matters temporal. *Trends Cogn. Sci.* **6**, 105–106 (2002).
88. J. S. Li, A. A. Sarma, T. J. Sejnowski, J. C. Doyle, Internal feedback in the cortical perception-action loop enables fast and accurate behavior. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2300445120 (2022).
89. D. A. Robinson, The use of control systems analysis in the neurophysiology of eye movements. *Annu. Rev. Neurosci.* **4**, 463–503 (1981).
90. F. A. Miles, S. G. Lisberger, Plasticity in the vestibulo-ocular reflex: A new hypothesis. *Annu. Rev. Neurosci.* **4**, 273–299 (1981).
91. M. Haruno, D. M. Wolpert, M. Kawato, Mosaic model for sensorimotor learning and control. *Neural Comput.* **13**, 2201–2220 (2001).
92. J. J. Moore, Moore_et_al_PNAS_2024_Code. Zenodo. <https://zenodo.org/records/11399185>. Deposited 31 May 2024.