



Internal feedback in the cortical perception–action loop enables fast and accurate behavior

Jing Shuang Li^{a,1}, Anish A. Sarma^{a,b,1}, Terrence J. Sejnowski^{c,d,2}, and John C. Doyle^{a,2}

Contributed by Terrence Sejnowski; received January 9, 2023; accepted July 18, 2023; reviewed by Mark Schnitzer and Domitilla Del Vecchio

Animals move smoothly and reliably in unpredictable environments. Models of sensorimotor control, drawing on control theory, have assumed that sensory information from the environment leads to actions, which then act back on the environment, creating a single, unidirectional perception–action loop. However, the sensorimotor loop contains internal delays in sensory and motor pathways, which can lead to unstable control. We show here that these delays can be compensated by internal feedback signals that flow backward, from motor toward sensory areas. This internal feedback is ubiquitous in neural sensorimotor systems, and we show how internal feedback compensates internal delays. This is accomplished by filtering out self-generated and other predictable changes so that unpredicted, actionable information can be rapidly transmitted toward action by the fastest components, effectively compressing the sensory input to more efficiently use feedforward pathways: Tracts of fast, giant neurons necessarily convey less accurate signals than tracts with many smaller neurons, but they are crucial for fast and accurate behavior. We use a mathematically tractable control model to show that internal feedback has an indispensable role in achieving state estimation, localization of function (how different parts of the cortex control different parts of the body), and attention, all of which are crucial for effective sensorimotor control. This control model can explain anatomical, physiological, and behavioral observations, including motor signals in the visual cortex, heterogeneous kinetics of sensory receptors, and the presence of giant cells in the cortex of humans as well as internal feedback patterns and unexplained heterogeneity in neural systems.

internal feedback | speed–accuracy trade-off | optimal control | sensorimotor control

Feedback control is an essential strategy for both engineered and biological systems to achieve reliable movements in unpredictable environments (1). Optimal and robust control theory, which provides a general mathematical foundation to study feedback systems, has been used successfully to explain behavioral observations by modeling the sensorimotor system as a single control loop, also called the perception–action cycle or perception–action loop (2–4). In these models, the sensorimotor system senses the environment, communicates signals from sensors to the brain, computes actions, and then acts on the environment, feeding back to the sensors and forming a single unidirectional loop as shown in Fig. 1.

Consider the canonical model of localized function in the primate visuomotor cortical pathway, depicted in Fig. 2: A visual signal is encoded on the retina, then travels to the lateral geniculate nucleus (LGN) of the thalamus, and on to the primary visual cortex (V1), progressing through successive transformations until it reaches the primary motor cortex (M1), the spinal cord, and ultimately the muscles. Although intuitive, this feedforward model neglects a well-known and ubiquitous feature of sensorimotor processing: internal feedback, which is the main focus of this paper (5).

The perception–action control model does not have a direct role for internal feedback connections. Internal feedback includes all signals that do not flow from sensing toward action. We can divide internal feedback into two broad categories: counterdirectional to feedforward projections and lateral interactions within or between areas at the same processing level. Counterdirectional internal feedback is in the opposite direction of the single-loop model (for instance, from V2 to V1); these signals flow from action toward sensing. Lateral internal feedback consists of recurrent connections that are used for exchanging information within a cortical area (for instance, within V2), or between areas at the same level in parallel streams (such as between areas MT and IT in the dynamical and object recognition streams, respectively). This distinction emphasizes the importance of where signals are spatially located in cortical hierarchies (Fig. 2).

The single-loop model offers a set of tools from control theory and a conceptual framework that allows subsystems to be treated as successive transformations that can be

Significance

Internal feedback projections—signals flowing from motor areas or late sensory processing regions back to early sensory processing regions such as primary visual and auditory areas—are ubiquitous in the sensorimotor nervous system and are as or more numerous than feedforward projections. However, the function of internal feedback is poorly understood, particularly in the context of task performance. We leverage control theory and simple models to demonstrate that internal feedback facilitates good task performance when there are communication limitations such as internal time delays and speed–accuracy trade-offs, which motivate compensatory feedback signals to counter self-generated and predictable movements. Control theory explains why motor-related signals are found throughout the sensory cortex and why the motor cortex is dominated by internal dynamics.

Author contributions: J.S.L., A.A.S., T.J.S., and J.C.D. designed research; performed research; and wrote the paper.

Reviewers: M.S., Stanford University; and D.D.V., Massachusetts Institute of Technology.

The authors declare no competing interest.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹J.S.L. and A.A.S. contributed equally to this work.

²To whom correspondence may be addressed. Email: doyle@caltech.edu or terry@salk.edu.

Published September 22, 2023.

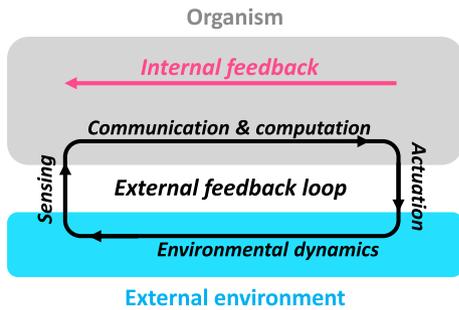


Fig. 1. Single-loop model of sensorimotor control. The organism receives information from the external environment via sensors, communicates this information through the body, computes actions, and then acts on the environment; this forms the external feedback loop, or single loop model (black). Internal signals that flow opposite to the direction of the external feedback loop are classified as internal feedback (pink). Thus, the internal feedback is counterdirectional. Internal feedback also includes lateral interactions within an area or between areas at the same processing stage (not shown).

studied in isolation. However, these subsystems are not isolated. With internal feedback, each subsystem has access to bottom up, top down, and lateral information. The eye is itself a site of computation and control: as the eye moves and senses different parts of the visual scene, lateral interactions within the retina control spatial and temporal filter properties that can adapt and identify important features under a wide range of illumination and scene dynamics (6, 7). Retinal ganglion cells project to relay neurons in the LGN, which then project to the primary visual cortex, V1, but a much greater number of feedback neurons project from V1 to LGN (8–10) (Fig. 2).

Projections from motor and later sensory areas in the cortex to early visual areas have a wide range of morphology, myelination, and synaptic kinetics (8, 11, 12). Given the position of M1 in the final common pathway, one might expect activity in M1 to be driven by current visual stimuli or current movements, but instead, autonomous internal dynamics dominate the data (13). Counterintuitively, signals related to movements of the whole body are found in areas typically associated with particular parts of the body, such as the hand area, as well as sensory areas such as the primary visual cortex (14–17). Indeed, recent analysis of the correlation structure between neurons during a visual discrimination task revealed a task-related global mode in the correlations between cortical neurons associated with the task response rather than the sensory stimulus, strongly supporting the idea that *Top-Down* feedback is an important element of sensory processing (18). These motor-related signals in sensory pathways, which span subsystems and tasks, are generated by internal feedback and are the focus of this study.

Internal feedback has been studied in the context of sensory predictive coding (19, 20) and has been invoked in other modeling studies and theory frameworks (21–25). However, these models focus on sensory or motor systems separately and do not account for key constraints on neuronal communication in both space and time to achieve sensorimotor tasks. Achieving fast and accurate computation and communication across brain areas is difficult, or even impossible, because communication may be slow, limited in bandwidth and constrained to spatially localized populations.

Here, we build on the foundations of recent work in distributed control theory (26–30) and show that internal feedback is a solution to achieving rapid and accurate control given the spatial and temporal constraints on brain components and communication systems. We analyze an idealized class of control models

and prove mathematically that internal feedback is necessary for achieving optimal performance in these idealized models. Internal feedback serves at least three functions: state estimation, localization of function, and focused attention, all of which are crucial for effective sensorimotor control and survival. This theory explains why there are differences in population responses between M1 and V1, why different projections predominantly activate AMPA or NMDA glutamate receptors, the functions of giant pyramidal cells in visuomotor control, and both the uses and limitations of localization of function in the cortex. There is a general principle behind all of these physiological properties.

Task Model and Performance. We analyze expected values and theoretical bounds on task performance for highly simplified control loop models motivated by a well-studied and ethologically relevant tracking task—reaching for a moving object. The goal of the task is continuous pursuit, such as catching a fly ball in baseball, rather than one-time contact between limb and object.

The complete tracking task requires identification of the object in a cluttered visual scene, prediction of the object’s movement, and generation and execution of bimanual limb movement. We make many simplifying assumptions that allow us to study internal feedback in an accessible way using familiar linear dynamical systems. Models of greater detail and complexity are discussed briefly.

Single-loop feedback control. Consider the task of tracking a moving object with the endpoint of a limb on a plane. The variable to be controlled is the tracking error—the distance between the hand and the object. We start by assuming that the system controlling the limb can perfectly sense the position of the limb and object at every instant, which will be relaxed in

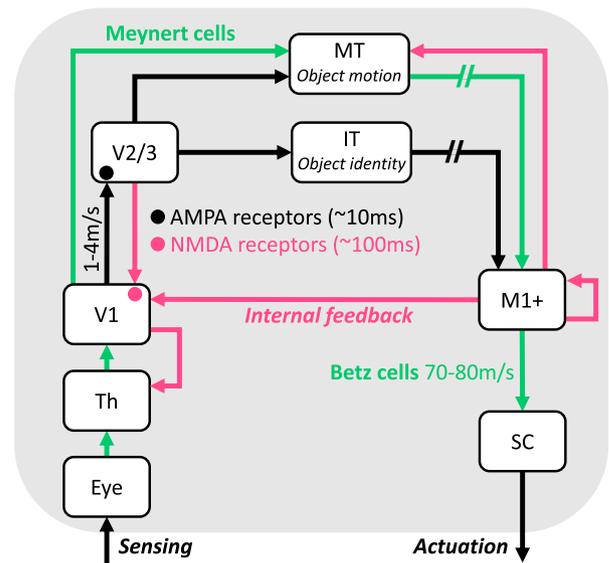


Fig. 2. A partial, simplified schematic of sensorimotor control. We focus on key cortical and subcortical areas and communications between them. Black and green arrows indicate communications that traverse from sensing toward actuation; green arrows are particularly fast pathways, which enable the tracking of moving objects in our model. Pink arrows indicate internal feedback signals, which traverse from actuation toward sensing. Solid lines indicate direct neuronal projections, while broken lines include both direct and indirect connections. SC, spinal cord; Th, thalamus; V1, primary visual cortex; M1+, primary motor cortex and additional motor areas; V2/3, secondary and tertiary visual cortex; IT, inferotemporal cortex; MT, mediotemporal cortex (V5). Only a subset of the internal feedback pathways are shown (e.g., not included are internal feedback signals from M1+ to V2 and IT).

later models. The cost is defined as the squared Euclidean norm of the tracking error over time, normalized by the total amount of time, with a smaller cost indicating better tracking.

Let x , u , and w represent the tracking error, the control action on the limb, and the action of the object, respectively. We will refer to x as the state of the system. Let A be a matrix that represents the intrinsic dynamics of x , including features such as the movement of the object or mechanical coupling between two dimensions of limb movement. The time-evolving dynamics of the tracking error follows from a linear equation of motion:

$$x(t+1) = Ax(t) + u(t) + w(t). \quad [1]$$

In general, the difficulty of a task will depend on properties of A such as its eigenvalues and the strength of coupling between states. For example, if the spectral radius of A is less than 1, this corresponds to a task in which tracking error x decreases with no limb action, an easy task.

The actions u provide feedback control on the tracking error, computed by an arbitrary function K that has access to all past and present tracking errors $x(1:t)$, as follows:

$$u(t) = K(x(1:t)). \quad [2]$$

The optimal solution to this problem is the linear quadratic regulator (LQR) and the optimal controller is $K(x(1:t)) = -Ax(t)$ if the action of the object w behaves as white noise (1). This controller is compatible with the single-loop model of sensorimotor control, as there is no internal feedback, and the addition of internal feedback does not provide any additional performance advantage.

Controllers without internal feedback are optimal for a large but special class of problems, including standard state feedback and full control problems from control theory. Though mathematically elegant, these controllers make assumptions that are impractical when applied to biological systems. In subsequent sections, we relax some of the assumptions implicit in this single-loop model and show that small deviations from assumptions relevant to biological systems introduce the need for internal feedback.

Any of the controllers in subsequent sections can be implemented in a variety of ways, although whether or not a particular controller needs internal feedback is generic across all possible implementations. We choose particular nonunique controller implementations with internal feedback for which the optimal solution is relatively transparent and easy to interpret.

State Estimation Requires Counterdirectional Internal Feedback.

Internal feedback facilitates implicit estimation in the presence of sensor delays. Simple modifications to the control problem described above lead to an optimal controller K whose implementation requires internal feedback. One such modification is the introduction of sensor delays, which are ubiquitous in biological systems (for example, the neuronal conduction time from the eye to the motor cortex is on the order of tens to hundreds of milliseconds). Sensor delays can be modeled by introducing a virtual internal state x_s , which represents the adjusted tracking error from the previous time step (29). This formulation allows us to pose the delayed-sensor tracking problem as a standard control problem which can be optimally solved by a linear quadratic

regulator (LQR):

$$\begin{aligned} \tilde{A} &= \begin{bmatrix} A & 0 \\ I & 0 \end{bmatrix}, C = [0 \quad I] \\ \begin{bmatrix} x(t+1) \\ x_s(t+1) \end{bmatrix} &= \tilde{A} \begin{bmatrix} x(t) \\ x_s(t) \end{bmatrix} + u(t) + \begin{bmatrix} w(t) \\ 0 \end{bmatrix} \\ u(t) &= KC \begin{bmatrix} x(t) \\ x_s(t) \end{bmatrix}, \end{aligned} \quad [3]$$

where the virtual internal state x_s contains delayed information about the tracking error. Controller K can be partitioned into two block-matrices, ($K = [K_1^T \quad K_2^T]^T$). The resulting system is shown in Fig. 3. Here, the controller does not directly “perceive” the tracking error x and only has access to the virtual internal state x_s . However, the controller can freely take actions that affect both the tracking error and the virtual state. The action on the virtual state, as shown in Fig. 3, is an example of counterdirectional internal feedback with gain K_2 .

For the delayed sensing problem, the optimal controller has a simple analytical form: $K_2 = -A$ is the internal feedback and $K_1 = -A^2$. If no internal feedback is allowed (i.e., we enforce $K_2 = 0$), then the optimal controller is $K_1 = -A^2/4$. We compare the performance of these two controllers in Fig. 4 and see that the controller with internal feedback far outperforms the controller without internal feedback. We also note that as the task becomes more difficult (spectral radius of $A > 2$), the controller without internal feedback is unable to stabilize the closed-loop system and tracking breaks down.

For a controller with sensory delays, internal feedback is required for optimal performance. This also applies to controllers with actuator delays (29). In both cases, internal feedback adjusts delayed signals to compensate for actions taken and information received during the delay; in other words, internal feedback implicitly compensates for the delays.

The linear quadratic problem that we consider here could be straightforwardly tested in a laboratory setting. A simple real-world task described by linear dynamics is the action of a stable limb, resting on a surface or manipulandum, tracking an object over a line or plane. This can be modeled by neutrally stable double-integrator dynamics ($A = \begin{bmatrix} I & I \\ 0 & I \end{bmatrix}$, $\alpha = 1$) with single time-step delays corresponding to internal loop delays on the order of 100 ms. More complex models of tasks corresponding to more realistic scenarios, such as movement against gravity or adversarial disturbances, would tend to increase the penalty of control without internal feedback.

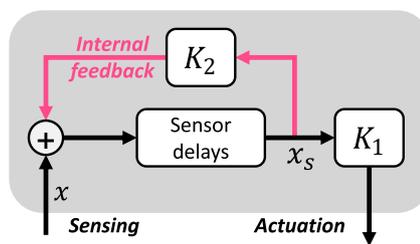


Fig. 3. Optimal control model for a system with sensor delays. Tracking error x is sensed, and then communicated by the sensor with some delay to the K_1 block, which computes the appropriate actuation. Counterdirectional internal feedback (pink) conveys information from actuation back toward sensing. Internal computation K_2 adjusts the sensor signal to compensate for actions taken by the system; this results in improved performance.

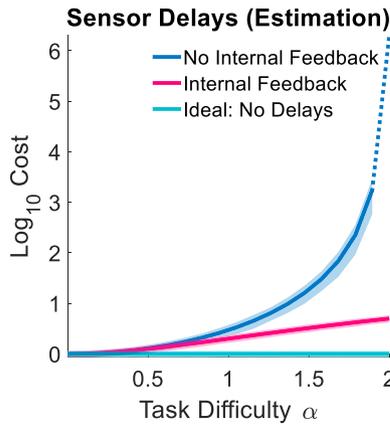


Fig. 4. Internal feedback improves performance when there are internal delays in sensing. The scalar problem of tracking a moving target over a line was simulated, varying the task difficulty (α = spectral radius of A , representing the dynamics). The “Ideal” controller contains no sensor delays. The “Internal Feedback” controller contains sensor delays, and uses internal feedback to compensate for the delays. The “No Internal Feedback” controller contains sensor delays but uses no internal feedback. As α approaches 2, the task becomes infeasible without internal feedback (broken line). Shaded areas indicate SDs.

We next describe controllers with internal feedback that include sensor delays, actuator delays, and imperfect sensing that will motivate a general model for sensorimotor control in brains. **Intrinsic internal feedback in the Kalman filter.** We now consider the case in which sensing is instantaneous, but imperfect and noisy, and actuation is imperfect. Consider the following system:

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + w(t) \\ y(t) &= Cx(t) + v(t), \end{aligned} \quad [4]$$

where y is the sensor input and v is the sensor noise, assumed to be white noise. Matrix B represents the effect of action u on tracking error x , and matrix C represents how sensor input y is related to tracking error x . This is a standard formulation in control theory, and the optimal controller makes use of controller gain K and estimator gain L as follows:

$$\begin{aligned} \hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + L(y(t) - C\hat{x}(t)) \\ u(t) &= K\hat{x}(t), \end{aligned} \quad [5]$$

where \hat{x} is an internal estimate of tracking error x . This optimal controller uses the Kalman filter, which inherently contains three counterdirectional internal feedback pathways irrespective of delays being present. These pathways are represented by the blue arrows in Fig. 5 and play a central role in state estimation. The pathway through A estimates state evolution in the absence of noise and actuation; the pathway through B accounts for controller action, and the pathway through C predicts incoming sensory signals based on the internal estimated state.

The implementation shown in Fig. 5 is not unique. We now briefly discuss a few equivalent implementations that use less internal feedback and explain why they are less advantageous in the sensorimotor context than the implementation in Fig. 5.

In one alternative implementation, we can remove the internal feedback through B and replace A with $A + BK$. However, this requires duplication of K ; in the sensorimotor context, which requires duplicating of motor structures within visual structures. In another alternative implementation, we can remove the internal feedback through C and replace A with $A - LC$. This

requires a duplication of L . Additionally, filtering out predictable sensory input via C earlier (as is done in Fig. 5) can be preferable to filtering it out later (as in our alternative implementation). This is because the filtered information is typically much smaller in bandwidth and requires less resources to communicate: The earlier we perform this filtering, the less resources we require to pass this information forward. If communications are subject to a speed-accuracy trade-off (described below), then earlier filtering allows us to pass sensory information forward with less delay.

Sources of internal feedback are preserved in a Kalman filter with delays. We now synthesize a model that combines features from previous sections: sensor delays, actuator delays, and imperfect sensing. The model can be constructed using virtual states as follows:

$$\begin{aligned} \begin{bmatrix} x(t+1) \\ x_a(t+1) \\ x_s(t+1) \end{bmatrix} &= \begin{bmatrix} A & B & 0 \\ 0 & 0 & 0 \\ C & 0 & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ x_a(t) \\ x_s(t) \end{bmatrix} \\ &+ \begin{bmatrix} 0 & 0 \\ I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} u(t) \\ u_s(t) \end{bmatrix} + \begin{bmatrix} w(t) \\ 0 \\ 0 \end{bmatrix} \quad [6] \\ y(t) &= x_s(t), \end{aligned}$$

where x_a and x_s are virtual internal states corresponding to delayed actuator commands and delayed sensor signals, respectively, and u_s represents compensation on virtual internal states. We can use standard control theory to obtain the optimal controller gain K and optimal estimator gain L . Due to the block-matrix structure of the system matrices, the optimal gains have the following structure: $K = [K_1 \ K_2 \ 0]$, and $L = [L_1^T \ 0 \ L_2^T]^T$ (29). The controller can be implemented as follows:

$$\begin{aligned} \delta(t+1) &= Cx(t) - C\hat{x}(t) - L_2\delta(t) \\ \hat{x}(t+1) &= A\hat{x}(t) + Bx_a(t) + L_1\delta(t) \\ u(t) &= K_1\hat{x}(t) + K_2x_a(t), \end{aligned} \quad [7]$$

where δ is the delayed difference between the estimated sensor input and true sensor input, adjusted by the $L_2\delta(t)$ term. The resulting controller, shown in Fig. 6, contains two internal feedback pathways related to delay: one pathway compensates for sensor delays, and the other compensates for actuator delays. The remaining internal feedback is inherent to the Kalman filter, as described in the previous section and shown in Fig. 5. Overall, the

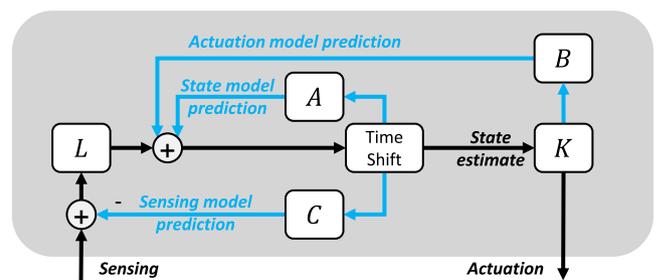


Fig. 5. Internal feedback in a controller with instantaneous but imperfect sensing and actuation. A , B , and C represent the state, actuation, and sensing matrices of the physical plant; K represents the optimal controller, and L represents the optimal observer. The Time Shift block shifts $\hat{x}(t+1)$ to $\hat{x}(t)$ in Eq. 5. The internal feedback pathways (blue) are inherent to the Kalman filter; these use state, actuation, and sensing models to create an internal estimate of the tracking error or state. All internal feedback depicted in this diagram is counterdirectional and assumed to have no delay and infinite bandwidth.

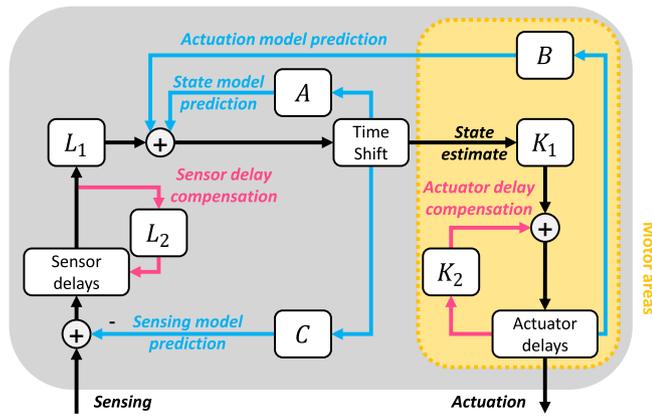


Fig. 6. Internal feedback in a controller with sensor and actuator delays. A , B , and C represent the state, actuation, and sensing matrices of the physical plant; K_1 , K_2 , L_1 , and L_2 are submatrices of the optimal controller and observer gains. The internal feedback pathways (pink) through L_2 and K_2 compensate for sensor and actuator delays, respectively. Other internal feedback pathways (blue) are inherent to the Kalman filter. All internal feedback depicted in this diagram is counterdirectional. The yellow box contains parts of the controller that roughly correspond to motor areas in the cortex.

inclusion of sensor delays, actuator delays, and imperfect sensing results in an optimal controller with several internal feedback pathways, each of which serves a specific, interpretable purpose.

Localization of Cortical Function Requires Lateral Internal Feedback. Almost all muscles in the body are engaged in even the simplest actions, such as reaching. Controlling a system with many degrees of freedom is a difficult problem for motor control even without delays. Localization of function is well established in the motor cortex, with different body parts controlled by different cortical areas; however, communication and computation between localized cortical areas typically have spatial and temporal constraints, compared with signals within areas.

Consider two motor areas and partition tracking errors into two sets x_1 and x_2 , representing two distinct but possibly coupled subsystems (e.g., two distinct limbs that are mechanically coupled) using the problem formulation described by Eq. 1. The overall tracking error is $x = [x_1^T \ x_2^T]^T$. Correspondingly, we partition actuators into two sets u_1 and u_2 that act on their respective subsystems via local controllers: $u = [u_1^T \ u_2^T]^T$.

Each local controller senses and controls one subsystem; i.e., local controller 1 senses x_1 and computes u_1 , and local controller 2 senses x_2 and computes u_2 . Local controllers may communicate with another; however, due to localization constraints, the cross-communication is delayed. Thus, local controller 1 cannot directly access x_2 without some delay and similarly for local controller 2.

We observe that without the constraint of localized communication, the optimal controller for Eq. 1 is $u = -Ax$. If A is block-diagonal (i.e., x_1 and x_2 are uncoupled), then this controller obeys localized communication constraints—in fact, no cross-communication (internal feedback) is required between the two local controllers. However, if the two subsystems are coupled, then this controller requires rapid, global communication, which violates localized communication constraints.

To enforce localized communication, we reformulate the problem by introducing virtual states x'_1 and x'_2 , which represent delayed cross-communication between the two local controllers. x'_1 is information sent from local controller 1 to local controller 2,

with delay; and similarly for x'_2 . We also define u'_1 and u'_2 , which model interconnections between virtual states and real tracking errors. For simplicity, we assume unit delay. The reformulated problem then becomes

$$\tilde{x} = \begin{bmatrix} x_1 \\ x'_2 \\ x'_1 \\ x_2 \end{bmatrix}, \tilde{u} = \begin{bmatrix} u_1 \\ u'_2 \\ u'_1 \\ u_2 \end{bmatrix}, \tilde{w} = \begin{bmatrix} w_1 \\ 0 \\ 0 \\ w_2 \end{bmatrix},$$

$$\tilde{A} = \begin{bmatrix} A_{11} & 0 & 0 & A_{12} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ A_{21} & 0 & 0 & A_{22} \end{bmatrix}, K = \begin{bmatrix} * & \Delta & * & 0 \\ \Delta & * & \Delta & * \\ * & \Delta & * & \Delta \\ 0 & * & \Delta & * \end{bmatrix} \quad [8]$$

$$\tilde{x}(t+1) = \tilde{A}\tilde{x}(t) + \tilde{u}(t) + \tilde{w}(t)$$

$$\tilde{u} = K\tilde{x}.$$

The zeros in the *Top Right* and *Bottom Left* corners of the K matrix preserve localized communication; they enforce that the two local controllers cannot communicate instantaneously to one another. Asterisks and triangles indicate free values: Triangles represent sites of potential cross-communication, or lateral internal feedback. When these free values are optimized to achieve optimal performance with localized communication, the resulting K matrix is

$$K = \begin{bmatrix} -A_{11} & 0 & -A_{12} & 0 \\ 0 & 0 & -A_{12} & A_{12} \\ A_{21} & -A_{21} & 0 & 0 \\ 0 & -A_{21} & 0 & -A_{22} \end{bmatrix}. \quad [9]$$

The resulting local controllers are shown in Fig. 7. Note that the $-A_{12}$ term in the second row and the $-A_{21}$ term in the fourth row of K correspond to lateral internal feedback. Here, these internal feedback signals carry predicted values of the unsensed tracking errors for each controller, after taking control action into account; for instance, internal feedback from local controller 2 to local controller 1 conveys the predicted value of x_2 , after taking control action from controller 2 into account.

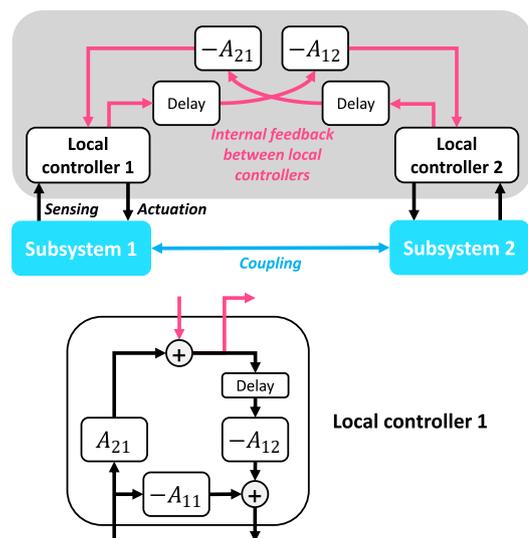


Fig. 7. Optimal localized control of two coupled subsystems. (Top) Overall schematic. Each subsystem has its own corresponding local controller, which senses and actuates only its assigned subsystem. Local controllers communicate to each other via lateral internal feedback (pink), with some delay. (Bottom) Circuitry of local controller 1. Local controller 2 has identical circuitry, with different matrices; A_{12} instead of A_{21} , A_{22} instead of A_{11} , etc.

We can develop intuition for this implementation by following an impulse w through time:

$$\tilde{x}(1) = \begin{bmatrix} w_1 \\ 0 \\ 0 \\ w_2 \end{bmatrix} \rightarrow \tilde{x}(2) = \begin{bmatrix} A_{12}w_2 \\ A_{12}w_2 \\ A_{21}w_1 \\ A_{21}w_1 \end{bmatrix} \rightarrow \tilde{x}(3) = 0. \quad [10]$$

The performance of this controller can be compared to the controller without internal feedback in Fig. 8. The best possible linear controller for the controller without internal feedback results in severe performance degradation. As task difficulty increases, this controller is unable to stabilize the closed-loop system and tracking becomes infeasible. With internal feedback, task performance stays near the centralized optimal case where local controllers can communicate freely without delay.

This analysis shows that when motor function is localized to specialized parts of the motor cortex that control particular parts of the body, cross-communication via internal feedback between local controllers is essential. Local circuits in the two hemispheres must also be coordinated—indeed, they are connected by a massive corpus callosum that crosses the midline.

The cross-talk between local controllers is supported by the presence of global signals from movements of the whole body to the local controller in the motor cortex specialized to particular parts of the body. In reality, all body movements are mechanically coupled, something which the motor system can conceal through effective localization and coordination using internal feedback.

Speed-Accuracy Trade-Offs Drive the Use of Layering and Internal Feedback for Attention. We have shown that state estimation and localization of function require internal feedback to correct for self-generated or predictable movements. We now show how to efficiently track a moving object with the limitations imposed by neural components and internal time delays using an attention mechanism.

Up to this point, we have assumed that the controller can directly sense the position of the object (perhaps with some delay). In the real world, a scene can comprise many objects,

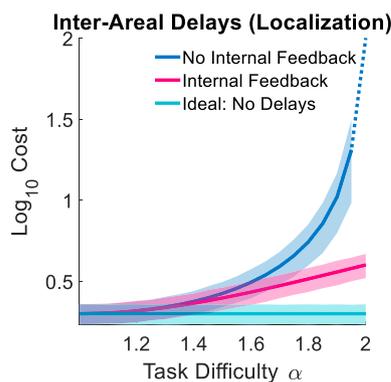


Fig. 8. Localization of function within the motor-related cortex: Although different parts of the cortex control different parts of the body, these parts of the body are inherently mechanically coupled. As a result, internal feedback is useful and in some cases necessary to maintain localization of function. In simulations, we consider the problem of tracking a moving target over a two-dimensional space, varying the task difficulty. The “Ideal” controller is centralized (i.e., no delays between local controllers) and obtains the best performance. The localized controller with internal feedback achieves similar performance. The localized controller without internal feedback suffers from substantially worse performance (higher cost). As task difficulty increases, the task becomes infeasible without internal feedback (broken line). Shaded areas indicate SDs.

which makes it more difficult for a sensorimotor system to localize an object in the scene. However, a moving object, once identified, can be more easily discriminated from a static visual scene. This illustrates the distinction between scene-related tasks (such as object identification) and error-related tasks (such as object tracking), which in the visual cortex is accomplished by the ventral and dorsal streams, respectively.

This distinction also mirrors the separation between bumps and trails in the mountain-biking task studied in ref. 27, allowing us to build on the control architecture in that task. The main difference is that instead of separating into two control loops, we use layering and internal feedback to supplement the control actions of the main control loop.

We consider a one-dimensional problem (tracking on a line) and use as the metric $\|x\|_\infty$ (worst-case tracking error for adversarial object action) rather than $\|x\|_2$ (average-case tracking error for random object action). Worst-case error is a more realistic model of many ethological tasks, and optimal solutions to worst-case control problems can have additional internal feedback pathways compared to average-case; however, the worst-case setting is less familiar in neuroscience models than the average-case setting we have considered to this point (31). We have some object whose position, r , is governed by the dynamics

$$r(t+1) = r(t) + w_r(t) + w_b(t), \quad [11]$$

where w_r represents object movement, and w_b represents changes in the background scene. Limb position p is governed by the dynamics

$$p(t+1) = p(t) - u(t), \quad [12]$$

where $u(t)$ is some limb action. The tracking error $x := r - p$ then obeys the dynamics

$$x(t+1) = x(t) + w_r(t) + w_b(t) + u(t), \quad [13]$$

where the task difficulty is implicitly equal to 1.

We assume that object movement and background changes are bounded: $|w_r(t)| \leq \epsilon_r$ and $|w_b(t)| \leq \epsilon_b$ for all t . Additionally, we assume that background changes are much slower than object movement: $\epsilon_b \ll \epsilon_r$, i.e.,

$$\epsilon_r + \epsilon_b \approx \epsilon_r. \quad [14]$$

Consider a movable sensor that senses some interval of size β on the continuous line. Information from the sensor must be communicated to the controller via axon bundles, which are subject to speed-accuracy trade-offs—that is, the higher bandwidth a signal, the slower it can be sent. Thus, roughly speaking, axonal communication can be low-bandwidth and fast, or high-bandwidth and slow. We can formalize this as follows for a volume of cortex axons with uniform radius, adapting from ref. 27:

We first observe that delay T is inversely proportional to axon radius ρ with proportionality constant α :

$$T = \frac{\alpha}{\rho}. \quad [15]$$

Firing rate per axon, ϕ , is proportional to axon radius with proportionality constant β :

$$\phi = \beta\rho. \quad [16]$$

Cross-sectional area s is related to axon radius ρ and the number of axons in the nerve n via:

$$s = n\pi\rho^2. \quad [17]$$

And finally, signaling rate R of the entire nerve, which is related to the resolution of information sent about the sensed interval, is represented by

$$R = n\phi. \quad [18]$$

These equations can be combined to obtain the speed–accuracy trade-off

$$R = \lambda T, \quad \lambda = \frac{s\beta}{\pi\alpha}. \quad [19]$$

The constant λ is proportional to s , the cross-sectional area; for projections of fixed length, this represents the spatial and metabolic cost to build and maintain the axons. In general, given some fixed cortical volume, we can either build few thick axons, which will have low delay but low information rate, or we can build many thin axons, which will have high information rate but high delay.

We implement this speed–accuracy trade-off using a static, memoryless quantizer Q with uniform partition, followed by a communication delay, as shown in Fig. 9, *Top*. This choice of quantizer does not add to the cost since it recovers the optimal cost over all possible quantizers (32).

The controller can move the sensor around; the interval sensed by the sensor remains constant, but the controller can choose where the interval lies. Assume the initial position of the object is known—we can select an initial sensor location and β appropriately such that $r(t)$ always falls within the sensed interval. In this case, the best possible tracking error for any delay T is

$$\epsilon_r T + \frac{\beta}{2\lambda T}. \quad [20]$$

The first term represents error from delay, object movement, and drift. In the time taken for information to reach the controller, the most adversarial action possible by the object and background would contribute a tracking error of $(\epsilon_r + \epsilon_b)T$; we apply simplification Eq. 14 to obtain $\epsilon_r T$. The second term represents

quantization error. For an interval of size β divided into N uniform subintervals, the worst-case error is $\frac{\beta}{N}$; we then use the fact that $N = 2^R = 2^{\lambda T}$.

This is achieved by the controller depicted on the Fig. 9, *Top*. The cost, as a function of T , is plotted in Fig. 10, *Left* with the label “No Internal Feedback” (where $T = T_s$). Here, the speed–accuracy trade-off is implicit. Very low values of T correspond to very low signaling rates—the controller does not receive enough information to act accurately, so performance is poor. The opposite problem occurs at very high values of T ; though the information is high-resolution, the time elapsed between information and action is too long, leading to poor performance. The best performance occurs at a sweet spot between these two extremes (27).

We can improve this performance by nearly an order of magnitude by adding an additional communication pathway and the requisite internal feedback. We now have two communication paths from the sensor, each with its own quantizer and delay block. The slower communication path uses quantizer Q_s with delay T_s , while the faster path uses Q_f with delay T_f . To further facilitate speed in the fast path, we allow it to send only a subset of information from the sensor (i.e., only send information about a small part of the sensed scene). Mathematically, let the fast path send information about an interval of size β_f , with $\beta_f < \beta$, and let this smaller subinterval be contained within the sensor interval. This subinterval is an implementation of attention. The fast path is the main actuation path, while the slower path provides compensatory signals via internal feedback; this is shown at Fig. 9, *Bottom*. In this case, the best possible cost is

$$\epsilon_r T_f + \frac{\beta_f + E_s}{2\lambda T_f} \quad [21]$$

$$E_s = \epsilon_b T_s + \frac{\beta}{2\lambda T_s}.$$

The first term represents error from delay and object movement, similar to Eq. 20. The second term represents a combination of quantization error from the fast communication pathway (β_f) and performance error of the slow pathway (E_s), which informs the fast pathway of where to place the subinterval. Notice that E_s takes the same form as Eq. 20.

The cost, as a function of T_s , is plotted in Fig. 10, *Left* with the label “Internal Feedback.” In this plot, we assume T_f to be its smallest possible value: one unit of delay. We see that using two quantizers in combination with internal feedback is superior to using one quantizer. We remark that this holds only when the two quantizers are different; if we simply use two quantizers with the same interval, bit rate, and delay, no such performance boost occurs. In general, holding T_s constant and decreasing T_f improves performance, as shown in Fig. 10, *Right*.

Functionally, the inclusion of a faster communication pathway allows action to be taken in a more timely manner than in the single-pathway case. Unlike in the single-pathway case, we are not encumbered by issues of low-resolution information; the slower communication pathway corrects the fast pathway through internal feedback. Here, as in previous examples, the internal feedback carries signals correcting for self-generated and slow, predictable changes. Overall, despite speed–accuracy trade-offs in communication, the system achieves fast and accurate behavior with the help of internal feedback, under reasonable assumptions about the dynamics of the scene and environment.

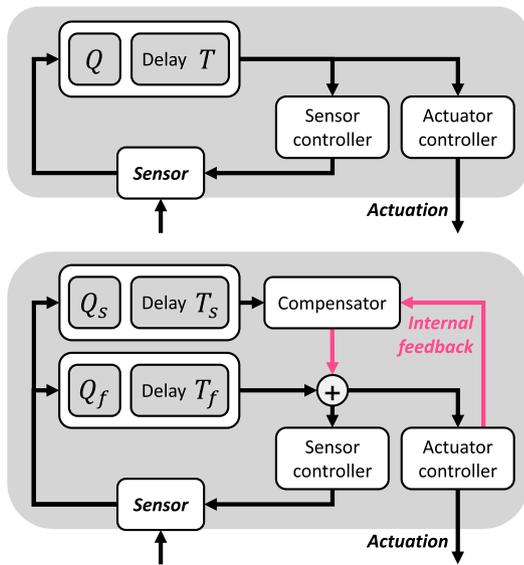


Fig. 9. Optimal control model of attention, with moveable sensor. (*Top*) Model with one communication path, in which information is quantized by quantizer Q and conveyed to the controller with delay T . (*Bottom*) Model with two communication paths, and two separate quantizers Q_s , Q_f , and respective delays. This model can be considered lateral (e.g., V1-V1) or counterdirectional (V2-V1) internal feedback (pink) between the two controller paths.

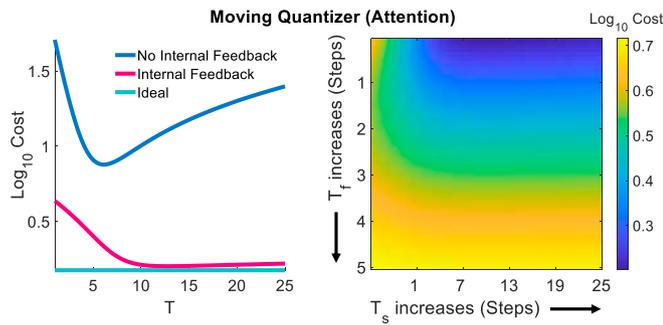


Fig. 10. (Left) Internal feedback and layering achieve superior performance when sensor-controller communications are subject to speed-accuracy trade-offs. The “No Internal Feedback” controller uses one layer, while the “Internal Feedback” controller uses two layers, with internal feedback between the layers. The two-layer case consists of a fast-forward pathway compensated by slow internal feedback, which takes slow background changes into account; this achieves better performance (Lower cost) than the case without internal feedback. The “Ideal” controller, where the sensor directly senses the moving object, is also shown. The layered system with internal feedback achieves performance close to the ideal. T represents delay. For the “No Internal Feedback” controller, it represents the delay of the single layer; for the “Internal Feedback” controller, it represents the delay of the slow layer, i.e., $T = T_s$. The delay of the fast layer is held constant. (Right) Performance (log cost) of the two-layer controller with internal feedback as delays of both layers are varied. Performance is best when T_f is low and T_s is sufficiently high.

Discussion

We analyzed a set of minimal control models to explore internal feedback in a perception-action loop with time delays and limited communication bandwidth. We showed how internal feedback, which is ubiquitous in brains and has been poorly understood, facilitates state estimation and localization of function and how attention facilitates sensorimotor performance. This is a step toward an end-to-end model of sensorimotor processing in neural systems.

The mathematical framework explored here can be applied across a wide range of experimental systems. The simple models are meant to provide an intuitive understanding of the control strategies for handling internal time delays and limited signaling bandwidth. The framework can be elaborated to make more specific predictions for more complex biological systems.

Task-Oriented Whole-System Frameworks Reveal Roles for Internal Feedback. There may be other functions for internal feedback in addition to compensating for time delays and limited cortical communication bandwidth. Additional functions that have been suggested are computation through dynamics (21), deploying recurrent networks (22), performing Bayesian inference (23, 24), generating predictive codes (19, 20), and many others.

These frameworks emphasize prediction but are largely confined to the context of either sensory processing or motor processing separately and do not explicitly model closed loop task performance. Our analysis considers sensorimotor control from an ethological perspective: The ultimate selection pressure on sensory processing is to support actions that ensure survival (33).

Our framework is consistent with recent neural recordings from the cortex showing that motor signals and the influence of past and current actions account for substantial cortical activity, previously considered spontaneous, background, or noise (14–17, 34). These internal feedback signals carry information about how actions propagate through the body and

its environment, ameliorating communication limitations that affect both plans and future actions.

The predictions of these models can be used to interpret how ablation, suppression, or delay of counterdirectional internal feedback would degrade performance in many visuomotor tasks. The performance gap should be more pronounced in tasks that involve quickly changing conditions. Similar interventions that disrupt lateral feedback interactions within the motor cortex should degrade performance in tasks where body parts are highly coupled mechanically—such as hands and shoulders—and less so for tasks where body parts are loosely coupled—such as speech and walking. This can be tested experimentally.

Our theory predicts that the fastest neurons from V1 to V2 or V1 to MT (including Meynert cells) should be the most highly activated when there are unpredictable changes in the visual scene. Similarly, we attribute autonomous dynamics in M1 to the communication of predicted actions laterally and counterdirectionally. This leads to the prediction that unexpected perturbations during a motor task should lead subpopulations of cells in M1 (including Betz cells) to transmit rapid change-related signals in contrast to the more slowly changing responses that accompany unperturbed control.

Standard Optimal Control Models Neglect Key Physiological Limitations. Optimal control theory is a general framework for sensorimotor modeling. Given a mathematical description of a system and some task specification, the optimal controller probably gives the best possible performance. However, these proofs assume that the components are fast and accurate, with instantaneous communication and control circuits implemented with fast and accurate electronics. Using these components, a single sense-compute-actuate loop is generally sufficient to achieve optimal behavior.

Applying control theory to model physiological circuitry requires a distinction between behavior and implementation. The same optimal performance may be achieved through a number of different implementations in the underlying circuitry. Although traditional control theory excels as a model of sensorimotor behavior, it does not incorporate the component-level constraints that are prevalent in biology; as a consequence, the ways that traditional control theory models are implemented may not be directly relevant to biological control.

Recent advances have extended traditional control theory to allow distributed control and incorporation of component-level constraints (26, 27, 29, 30). We build on this body of work to describe how constraints on components affect the implementation of an optimal distributed biological controller. In particular, we show how and why internal feedback arises in controllers whose components exhibit the speed-accuracy trade-offs found in brains.

Fast long-range association fibers in the cortex are metabolically and developmentally expensive, have low bandwidth (compared to slower fibers with higher bandwidth), require constant maintenance and repair, and are limited in number. Internal feedback from the motor cortex to earlier sensory areas can regulate communication along these pathways by suppressing self-generated and other predictable signals, freeing the fast pathways to selectively transmit the unpredicted changes needed by the motor system to make fast decisions. This virtualizes the behavior of the control system to produce actions that are both fast and accurate despite internal time delays and limited communication bandwidth.

Evidence for the Suppression of Self-Generated Sensory Signals by Corollary Discharge Signals. Efference copies of motor signals are ubiquitous throughout brains and serve several functions. Fast suppressive internal feedback signals originate before motor commands are executed and target sensory pathways before self-generated signals can subsequently reach higher levels of processing (35)

Perhaps the best understood neural system that suppresses predictable self-generated sensory signals is found in electric fish, which generate electric fields for navigation and communication (36). The electrosensitive lateral-line lobe (ELL), a cerebellum-like structure, receives both a corollary discharge of the generated electric field and sensory input from electroreceptors on the body of the fish. These two signals are subtracted in the ELL to detect externally generated electric fields. Suppression is learned using anti-Hebbian synaptic plasticity, in which the coincidences of incoming spikes and outgoing spikes lead to a decrease in synaptic strength. A similar arrangement is found in the dorsal cochlear nuclei of mammals, which receive corollary discharge signals from brainstem nuclei associated with vocalization and respiration as well as proprioceptive input from body movement (36).

Biophysical Speed–Accuracy Trade-Offs Drive Internal Feedback. Biological control systems do not have components that are both fast and accurate. Spiking neurons, though fast relative to other biological signaling mechanisms, are many orders of magnitude slower than electronics and face severe speed–accuracy trade-offs that constrain communication and control. For example, some neurons can rapidly convey a few bits of information, and others can slowly convey many bits of information, but neurons that rapidly convey many bits of information are expensive and correspondingly rare. Speed–accuracy trade-offs include the number of neurons (information rate) and their axonal diameter (conduction speed) in nerve bundles (27, 37). By cleverly combining components with different speed–accuracy trade-offs and using internal feedback as demonstrated above, brains are able to perform survival-critical sensorimotor tasks with speed and accuracy. Additional trade-offs include spike averaging versus spike timing. These trade-offs have consequences for the performance of sensorimotor systems that we can study in our control models.

The range of neural conduction speeds in humans spans several orders of magnitude (37). The fastest components are used in the feedforward loop, sending information from sensing areas toward motor areas. Internal feedback compensates for accuracy by filtering out slowly changing, predictable, or task-irrelevant stimuli, such that the fewest possible bits need to travel along the fastest possible neurons. From an evolutionary perspective, once a system can achieve fast responses, additional layers of control can be added to achieve more accurate and flexible behavior without sacrificing performance.

The reason internal feedback is limited in most engineered control systems is that internal time delays are negligible. But in biological systems, even the fastest neurons used in the feedforward loop give rise to significant time delays. This is why it is essential to include delays in control-based analyses of the forward loop in neural models of control (Fig. 4).

Fast Feedforward Conduction Is Key to Successful Sensorimotor Task Performance. In the cortex, the fastest, largest, and most striking neurons are the large pyramidal cells: Meynert cells in the primary visual cortex carry signals from rapidly moving objects; Betz cells in the motor cortex that project to the spinal cord are

responsible for rapid responses to perturbations from planned movements; and although their role is less clear, Von Economo cells in the prefrontal cortex (anterior cingulate and fronto-insular areas) project rapid signals to subcortical areas involved in the regulation of cognitive and emotional behaviors (38–41).

The visual hierarchy diverges into the dorsal and ventral streams, which are responsible for object motion and object identity, respectively. In natural scenes, object locations may change quickly, but object identities change relatively slowly; a mouse may move around rapidly in the visual field of a predator barn owl, but it remains a mouse, and its status as prey does not change. Thus, speed is crucial for the dorsal stream, but not the ventral stream.

This difference has physiological consequences in our minimal model of attention that could explain differences between cortical projections in the two streams: the giant Meynert cells project from V1 to MT (an object motion area in the dorsal stream; see Fig. 1), but there are no equivalently large cells projecting from V1 to upstream areas leading to the inferotemporal cortex (IT, an object identity area in the ventral stream). Reaching tasks could test the predictions of our control model for rapidly and unpredictably moving objects on a fixed background compared with predictably moving objects on nonstationary backgrounds.

Neurons in MT respond selectively to the direction of moving objects and provide signals that are used by the oculomotor system for the smooth pursuit of moving objects (42). There are two visual pathways from the retina to the area MT. In addition to the cortical pathway that projects from V1 to area MT (Fig. 1), the retina also projects to the pulvinar, another thalamic relay, to extrastriate areas of the visual cortex (43). These two pathways could implement the optical control model in Fig. 9, *Bottom*, where the fast, direct pathway is from the pulvinar and the delayed, indirect pathway from V1 corresponds to the slower pathway.

Internal Feedback Facilitates Fast Feedforward Signals in the Visual Cortex. In recent years, large-scale recordings from the visual cortex have uncovered nonvisual signals that challenge the traditional single-loop view of sensorimotor control. In the traditional view, visuomotor processing consists of a series of successive transformations from stimulus to response, with each cortical area along the way tuned to some aspect of stimulus space (44). However, although V1 does respond to visual stimuli, the activity of these neurons also carries motor-based internal feedback signals (also called corollary discharge) and task- or attention-related modulatory internal feedback (12, 16, 17, 33). Our attention model could be implemented as the enhancement of responses in neurons that selectively respond to an attended stimulus through internal feedback (45).

The number of projections from V1 to V2 is roughly the same as the number of neurons, of similar conduction speed, that project from V2 to V1 (8–10). However, these neurons are very different in morphological and molecular characteristics: the neurons that project feedforward from V1 to V2 primarily activate AMPA receptors, while the feedback neurons that project from V2 to V1 have a strong NMDA receptor component and terminate almost exclusively on excitatory pyramidal neurons (46, 47). Both of these receptors are activated by glutamate, but AMPA-mediated currents are fast, lasting only a few milliseconds, while NMDA-mediated currents can linger in the postsynaptic neurons for hundreds of milliseconds (48). This feedback could be relevant for top-down signaling to shape and control perception during actions. Because NMDA receptors trigger

synaptic plasticity, the feedback could also be important for learning how to suppress self-generated sensory signals as well as perceptual learning.

Pharmacologically blocking NMDA receptors in the visual cortex disrupts figure-ground discrimination; that is, a loss in capacity to contextually interpret the visual scene (49). In the context of our theory and minimal model of attention, internal feedback from V2 informs V1 of predictable elements of inputs arriving from sensory stimuli in the near future. Since the visual space cannot be sampled losslessly, these feedback signals could be helping V1 suppress predictable features, making the unpredictable features more salient (46).

Corticothalamic feedback releases bursts of spikes in thalamocortical neurons during awake states, which might have a role in the detectability of sensory stimuli (50). A study that activated the layer 6 corticothalamic feedback (51) concluded that L6 projection does not modulate ongoing sensory processing, but rather serves to briefly speed up thalamic inputs in specific behavioral contexts, consistent with the role of cortico-cortico feedback for speeding up sensory processing.

Internal Feedback Facilitates Localization of Function in the Motor Cortex. The primary motor cortex (M1) is dominated by its own past activity rather than static representations (13). In the context of the state estimation problem we considered in Fig. 6, these dynamics in the motor cortex are driven neither by motor representations nor by pattern generation, but by predictions of the consequences of self-action through local internal feedback, which need to be sent throughout the brain because the entire body is affected.

By the same principle, the localization of function within the motor cortex that we considered in Fig. 8 reconciles the conventional view of homuncular organization with, for example, the body-related signals found in putatively hand-related parts of the motor cortex, as well as signals related to the contralateral hand (14, 15, 52). As with motor signals in the visual cortex, these broad body movement signals in the motor cortex are crucial for identifying predictable consequences of motor signals from other body movements and separating them from unpredictable signals of critical importance for rapidly controlling localized body parts. This provides each body part with the context it needs to compensate for the movement of other body parts.

In our analysis, we assumed the existence of a distinct motor cortex that generates motor commands and a visual cortex that interprets visual scenes and is essentially an extension of the retina. With these assumptions, we consider where to place internal feedback.

Is it possible for the entire estimator to be implemented in the visual cortex? Since the estimator uses predictions of future actions, the estimator requires at least some input from the motor cortex.

Is it possible for the entire estimator to be implemented in the motor cortex? The dynamical structure of responses in the motor cortex is compatible with a predictive and delay-compensating function of exactly the kind our model suggests. However, our model also shows that the information transmitted from the sensor to the motor cortex depends on the estimator. Thus, the motor cortex would need to contain all of the visual cortex along with all of the internal feedback connections to perform the entire function. The evolution of the cortex favored a more distributed architecture.

In addition to feedback loops between the motor cortex and other cortical areas, there are also loops with the cerebellum

and the basal ganglia. These provide additional information about sensory predictions and sequences of future actions, respectively. Regions of the motor cortex that interdigitate between projections to body parts have recently been identified that are associated with stimulation-evoked complex actions and connectivity to internal organs such as the adrenal medulla that are associated with goals (53). These regions may be responsible for integrating skeletal body movements with visceral states and goals.

Learning on Internal Feedback Pathways Fine-Tunes Performance. Internal feedback pathways carry attentional signals that activate slow NMDA receptors, which in turn regulate the strengths of synapses (54). We have shown that internal feedback pathways are needed for ignoring self-generated and other predictable signals. Early in brain development, activation of NMDA receptors in the primary visual cortex before the first visuomotor experience is needed to suppress predictable feedback and the selection of unpredictable stimuli (49). Blocking these NMDA receptors during development impairs ongoing visuomotor skill learning later in life. Learning to reduce self-generated sensory prediction error can be implemented locally through the same internal feedback system that broadcasts motor predictions.

Reinforcement learning governed by circuits in the basal ganglia may also benefit from the internal feedback pathways in the cortex. Transient dopamine release, which carries reward prediction error, does not specify which sensory inputs were responsible for the reward, partly explaining why it is a much slower form of learning. Attentional internal feedback in the cortex automatically selects and represents the currently most salient information to guide motor actions. Attentional information projects to the striatum and makes it easier for the basal ganglia to associate the causally relevant sensory inputs with reward signals (33).

We have proposed that prediction is an essential aspect of performance in visuomotor tasks where fast and accurate responses are needed. Prediction is useful in the model because it enables compressed representations of the current state, which can then be transmitted more quickly across the nervous system because of the speed-accuracy trade-off discussed above. In the same way, high performance can be achieved by simpler representations of tasks, which in turn allows faster responses.

After a flexible but slow learning system has successfully mastered a task, it can then “load” a model onto a simpler, more rigid, and faster subcortical system. Internal feedback can facilitate this transfer. We have focused here on the fast pathways represented by large axons in the visuomotor cortex, but similar variation in timescales of conduction can be found throughout the sensorimotor system. We therefore propose that during the acquisition of fast and accurate motor skills, control would shift from slower learning systems in the cortex to less flexible subcortical parts of the motor system.

Attention has been studied primarily in the context of sensory processing. The importance of attentional signals for reducing time delays in making motor decisions adds a direction for future experimental studies. Attention is linked to conscious awareness and rides atop the global representation of the body throughout the cortex. This makes internal feedback a candidate feature of the nervous system that helps explain the sense of unity that we experience, which would otherwise be difficult to achieve within a balkanized control architecture built on body parts.

Materials and Methods

Our main methodological approach was mathematical analysis of simplified models, deriving general properties and proving theorems about classes of mathematical models (See refs. 28–30 for more technical analyses). We focused here on insights that simulating specific mathematical models might have for neuroscience (Figs. 4 and 8). The methods used to generate these models and simulations are described here.

Sensor Delays. We start with the mathematical methods used to generate the data plotted in Fig. 4.

The scalar dynamics are governed by:

$$\begin{aligned} x(t+1) &= \alpha x(t) + u(t) + w(t) \\ u(t) &= kx(t), \end{aligned} \quad [22]$$

where we assume $w(t)$ is independently and identically drawn from a standard normal distribution, and k is the controller. First, let us assume that the system has no sensing delay. Let us penalize the term $x(t) \cdot b \cdot b \cdot x(t)$, where $b = 1$. Then, the optimal controller is $k = -\alpha$. To find the expected value of the cost, we first solve a discretized Lyapunov equation:

$$AXA^T - X + Q = 0, \quad [23]$$

where A represents the closed-loop dynamics, and $Q = b \cdot b$. In the scalar case, the expected cost is equal to $b \cdot X \cdot b$. In our case, $A = \alpha - k = 0$, giving a cost of $X = 1$ for all values of α . This is plotted in Fig. 4 with the label "Ideal."

Now, we add sensory delays. Consider augmented system

$$\begin{aligned} \begin{bmatrix} x(t+1) \\ x_s(t+1) \end{bmatrix} &= \begin{bmatrix} \alpha & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ x_s(t) \end{bmatrix} + \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix} + \begin{bmatrix} w(t) \\ 0 \end{bmatrix} \\ \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix} &= \begin{bmatrix} 0 & k_1 \\ 0 & k_2 \end{bmatrix} \begin{bmatrix} x(t) \\ x_s(t) \end{bmatrix} \end{aligned} \quad [24]$$

Here, $x_s(t)$ is $x(t)$ after one time step of delay; we do not allow the controller to directly access $x(t)$, and we only allow it to access the delayed signal $x_s(t)$. $u_1(t)$ represents true actuation (equivalent to $u(t)$ from the undelayed case), and $u_2(t)$ and k_2 represent internal feedback; this is depicted in Fig. 3. Let us penalize the term $\begin{bmatrix} x(t) & x_s(t) \end{bmatrix} \cdot b \cdot b^T \begin{bmatrix} x(t) \\ x_s(t) \end{bmatrix}$, where $b = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. This corresponds to penalizing $x(t)$ but not $x_s(t)$. If we allow internal feedback, i.e., nonzero values of u_2 and k_2 , then the optimal controller is $k_1 = -\alpha^2$ and $k_2 = -\alpha$. We can once again solve the discrete Lyapunov equation Eq. 23, this time setting $A = \begin{bmatrix} \alpha & 0 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} 0 & k_1 \\ 0 & k_2 \end{bmatrix} = \begin{bmatrix} \alpha & -\alpha^2 \\ 1 & -\alpha \end{bmatrix}$ and $Q = b \cdot b^T$. The expected cost is equal to $b^T \cdot X \cdot b$. This is plotted in Fig. 4 with the label "Internal Feedback".

To plot the case of no internal feedback, we restrict $k_2 = 0$. In this case, with the same penalty as in the previous case, the optimal controller is $k_1 = -\alpha^2/4$. As with before, we solve the discrete Lyapunov equation Eq. 23, this time setting $A = \begin{bmatrix} \alpha & 0 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} 0 & k_1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \alpha & -\alpha^2/4 \\ 1 & 0 \end{bmatrix}$. The expected cost is, again, $b^T \cdot X \cdot b$, and is plotted in Fig. 4 with the label "No Internal Feedback."

SDs of cost were obtained experimentally by running 1,000 simulations for each value of α with the appropriate controller and initial conditions $x(t=0)$ and random disturbances w drawn from independent standard normal distributions. Each simulation was $T = 100$ time steps long, and the costs of each simulation are calculated as:

$$\frac{1}{T} \sum_{t=0}^T x(t) \cdot x(t), \quad [25]$$

for all cases. Note that in the case of delays, $x_s(t)$ is not penalized, and we are only concerned with $x(t)$. We then computed the mean and SD of the costs over these 1,000 simulations. In all cases, mean values of simulated costs coincided with the expected costs obtained analytically.

Localization. Here, we provide the detailed methods used to generate the data plotted in Fig. 8. The open-loop dynamics of the two-block system are given by:

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}. \quad [26]$$

To parameterize the task difficulty for simulations, we considered the two-state system with $A_{12} = A_{21} = \gamma$, where γ represents mechanical coupling between the states ($\gamma = 0$ means no coupling, while $\gamma = 1$ means that what happens on one state is felt by the coupled state one time step later, and $\gamma > 1$ means that what happens on one state is amplified when it affects the coupled state). We additionally set $A_{11} = A_{22} = 1$. In Fig. 8, $\alpha = \gamma + 1$ to give the spectral radius of A .

$$A = \begin{bmatrix} 1 & \gamma \\ \gamma & 1 \end{bmatrix}. \quad [27]$$

Three controllers were considered: one with instantaneous interareal communication (ideal), one with internal feedback allowing delayed interareal communication, and one with no internal feedback and no interareal communication. We compute the ideal cost with the discrete-time Lyapunov equation, as before, with $K = -A$.

To compute the controller and cost with internal feedback, we again sought to optimize the infinite-horizon quadratic state cost. Just as in the centralized case the optimal solution is to negate a disturbance in a single time step, in this case, the optimal solution is to negate the disturbance as quickly as possible, in two time steps. This solution is identical to the solution produced by computing the controller with System Level Synthesis (26).

In the case with no internal feedback and no interareal communication, we restricted our consideration to static and linear controllers. (In the other cases considered, this restriction is unnecessary, because the linear solution is globally optimal).

$$K = \begin{bmatrix} K_{11} & 0 \\ 0 & K_{22} \end{bmatrix}. \quad [28]$$

Under this restriction, the best available control policy is to neglect the off-diagonal terms. For $\gamma > 1$, the task is infeasible without internal feedback. We computed the ideal cost with the discrete-time Lyapunov equation.

Data, Materials, and Software Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. Stephen Lisberger helped clarify our discussion on the oculomotor system. J.C.D. and T.J.S. were supported by NSF NCS-FO 1735004. J.C.D. was supported by NIH NS121918. T.J.S. was supported by ONR N00014-16-1-282. J.S.L. was in part supported by NSERC PGSD3-557385-2021. This paper is based on the doctoral research of A.A.S. and J.S.L.

Author affiliations: ^aControl and Dynamical Systems, Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125; ^bSchool of Medicine, Vanderbilt University, Nashville, TN 37232; ^cDepartment of Neurobiology, Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037; and ^dDepartment of Neurobiology, Division of Biological Sciences, University of California San Diego, La Jolla, CA 92093

1. J. Doyle, B. Francis, A. Tannenbaum, *Feedback Control Theory* (Macmillan, 1992).
2. D. M. Wolpert, Z. Ghahramani, M. I. Jordan, An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).
3. E. Todorov, Optimality principles in sensorimotor control. *Nat. Neurosci.* **7**, 907–915 (2004).

4. D. W. Franklin, D. M. Wolpert, Computational mechanisms of sensorimotor control. *Neuron* **72**, 425–442 (2011).
5. E. Zagha, Shaping the cortical landscape: Functions and mechanisms of top-down cortical feedback pathways. *Front. Syst. Neurosci.* **14**, 33 (2020).
6. L. Zhaoqing, *Understanding Vision* (Oxford University Press, 2014).

7. T. Gollisch, M. Meister, Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron* **65**, 150–164 (2010).
8. E. M. Callaway, Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Netw.* **17**, 625–632 (2004).
9. A. Angelucci, J. Bullier, Reaching beyond the classical receptive field of V1 neurons: Horizontal or feedback axons? *J. Physiol. Paris* **97**, 141–154 (2003).
10. Y. El-Shamayleh, R. D. Kumbhani, N. T. Dhruv, J. A. Movshon, Visual response properties of v1 neurons projecting to v2 in macaque. *J. Neurosci.* **33**, 16594–16605 (2013).
11. D. J. Felleman, D. C. Van Essen, Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* **1**, 1–47 (1991).
12. L. Muckli, L. S. Petro, Network interactions: Non-geniculate input to V1. *Curr. Opin. Neurobiol.* **23**, 195–201 (2013).
13. M. M. Churchland *et al.*, Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
14. F. R. Willett *et al.*, Hand knob area of premotor cortex represents the whole body in a compositional way. *Cell* **181**, 396–409.e26 (2020).
15. S. D. Stavisky *et al.*, Neural ensemble dynamics in dorsal motor cortex during speech in people with paralysis. *eLife* **8** (2019).
16. C. Stringer *et al.*, Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364** (2019).
17. 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).
18. S. Ebrahimi *et al.*, Emergent reliability in sensory cortical coding and inter-area communication. *Nature* **605**, 713–721 (2022).
19. 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).
20. G. B. Keller, T. D. Mrsic-Flogel, Predictive processing: A canonical cortical computation. *Neuron* **100**, 424–435 (2018).
21. S. Vyas, M. D. Golub, D. Sussillo, K. V. Shenoy, Computation through neural population dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
22. K. Kar, J. Kubilius, K. Schmidt, E. B. Issa, J. J. DiCarlo, Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nat. Neurosci.* **22**, 974–983 (2019).
23. A. Bastos *et al.*, Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
24. E. Libby, T. J. Perkins, P. S. Swain, Noisy information processing through transcriptional regulation. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7151–7156 (2007).
25. L. Perrinet, R. Adams, K. Friston, Active inference, eye movements and oculomotor delays. *Biol. Cybern.* **108**, 777–801 (2014).
26. J. Anderson, J. C. Doyle, S. H. Low, N. Matni, System level synthesis. *Annu. Rev. Control* **47**, 364–393 (2019).
27. Y. Nakahira, Q. Liu, T. J. Sejnowski, J. C. Doyle, Diversity-enabled sweet spots in layered architectures and speed-accuracy trade-offs in sensorimotor control. *Proc. Natl. Acad. Sci. U.S.A.* **118**, 1–11 (2021).
28. A. A. Sarma *et al.*, "Internal feedback in biological control: Architectures and examples" in *Proceedings of the IEEE American Control Conference* (2022), pp. 456–461.
29. J. Stenberg, J. S. Li, A. A. Sarma, J. C. Doyle, "Internal feedback in biological control: Diversity, delays, and standard theory" in *Proceedings of the IEEE American Control Conference* (2022), pp. 462–467.
30. J. S. Li, "Internal feedback in biological control: Locality and system level synthesis" in *Proceedings of the IEEE American Control Conference* (2022), pp. 474–479.
31. J. Doyle, K. Glover, P. Khargonekar, B. Francis, State-space solutions to standard H/sub 2/ and H/sub infinity/control problems. *IEEE Trans. Autom. Control* **34**, 831–847 (1989).
32. A. A. Sarma, J. C. Doyle, "Flexibility and cost-sensitivity in a quantized control loop" in *Proceedings of the IEEE American Control Conference* (2019).
33. P. S. Churchland, V. S. Ramachandran, T. J. Sejnowski, "A critique of pure vision" in *Large-Scale Neuronal Theories of the Brain*, C. Koch, J. Davis, Eds. (MIT Press, 1994), pp. 23–60.
34. M. Leinweber, D. R. Ward, J. M. Sobczak, A. Attinger, G. B. Keller, A sensorimotor circuit in mouse cortex for visual flow predictions. *Neuron* **95**, 1420–1432.e5 (2017).
35. E. V. Holst, M. H., An internal model for sensorimotor integration. *Naturwissenschaften* **37**, 464–467 (1950).
36. C. C. Bell, V. Han, N. B. Sawtell, Cerebellum-like structures and their implications for cerebellar function. *Annu. Rev. Neurosci.* **31**, 1–24 (2008).
37. P. Sterling, S. B. Laughlin, *Principles of Neural Design* (MIT Press, 2015).
38. M. S. Livingstone, Mechanisms of direction selectivity in macaque v1. *Neuron* **20**, 509–526 (1998).
39. V. Chan-Palay, S. L. Palay, S. M. Billings-Gagliardi, Meynert cells in the primate visual cortex. *J. Neurocytol.* **3**, 631–658 (1974).
40. E. E. Fetz, "Functional organization of motor and sensory cortex: Symmetries and parallels" in *Dynamic Aspects Of Neocortical Function*, W. C. G. M. Edelman, W. E. Gall, Eds. (John Wiley, 1984), pp. 453–474.
41. J. Allman *et al.*, The von Economo neurons in frontoinsular and anterior cingulate cortex in great apes and humans. *Brain Struct. Function* **214**, 495–517 (2010).
42. S. G. Lisberger, Visual guidance of smooth-pursuit eye movements: Sensation, action, and what happens in between. *Neuron* **66**, 477–491 (2010).
43. C. E. Warner, Y. Goldshmit, J. A. Bourne, Retinal afferents synapse with relay cells targeting the middle temporal area in the pulvinar and lateral geniculate nuclei. *Front. Neuroanat.* **4**, 8 (2010).
44. D. H. Hubel, T. N. Wiesel, Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* **148**, 574–591 (1959).
45. J. Reynolds, T. Pasternak, R. Desimone, Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703–714 (2000).
46. M. W. Self, R. N. Koopmans, H. Supèr, V. A. Lamme, P. R. Roelfsema, Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 11031–11036 (2012).
47. J. Anderson, K. Martin, "Chapter 6 - interareal connections of the macaque cortex: How neocortex talks to itself" in *Axons and Brain Architecture*, K. S. Rockland, Ed. (Academic Press, San Diego, 2016), pp. 117–134.
48. D. Attwell, A. Gibb, Neuroenergetics and the kinetic design of excitatory synapses. *Nat. Rev. Neurosci.* **6**, 841–849 (2005).
49. F. C. Widmer, S. M. O'Toole, G. B. Keller, NMDA receptors in visual cortex are necessary for normal visuomotor integration and skill learning. *eLife* **11**, e71476 (2020).
50. S. M. Sherman, Tonic and burst firing: Dual modes of thalamocortical relay. *Trends Neurosci.* **24**, 122–126 (2001).
51. J. Ansorge, D. Humanes-Valera, F. P. Pazuin, M. K. Schwarz, P. Krieger, Cortical layer 6 control of sensory responses in higher-order thalamus. *J. Physiol.* **598**, 3973–4001 (2020).
52. K. C. Ames, M. M. Churchland, Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *eLife* **8** (2019).
53. E. M. Gordon, R. J. Chauvin, A. N. Van, A somato-cognitive action network alternates with effector regions in motor cortex. *Nature* **617**, 351–359 (2023).
54. F. Li, J. Z. Tsien, Memory and the NMDA receptors. *N. Engl. J. Med.* **16**, 302–303 (2009).