

Internal Models in Biological Control

Daniel McNamee^{1,2} and Daniel M. Wolpert^{1,3}

¹Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom; email: d.mcnamee@eng.cam.ac.uk

²Institute of Neurology, University College London, London WC1E 6BT, United Kingdom

³Zuckerman Mind Brain Behavior Institute, Department of Neuroscience, Columbia University, New York, NY 10027, USA; email: wolpert@columbia.edu

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Control Robot. Auton. Syst. 2019.
2:339–64

The *Annual Review of Control, Robotics, and
Autonomous Systems* is online at
control.annualreviews.org

<https://doi.org/10.1146/annurev-control-060117-105206>

Copyright © 2019 by Annual Reviews.
All rights reserved

Keywords

internal model, state estimation, predictive control, planning, optimal feedback control, Bayesian inference

Abstract

Rationality principles such as optimal feedback control and Bayesian inference underpin a probabilistic framework that has accounted for a range of empirical phenomena in biological sensorimotor control. To facilitate the optimization of flexible and robust behaviors consistent with these theories, the ability to construct internal models of the motor system and environmental dynamics can be crucial. In the context of this theoretic formalism, we review the computational roles played by such internal models and the neural and behavioral evidence for their implementation in the brain.

1. INTRODUCTION

Over the last half century, the hypothesis that the nervous system constructs predictive models of the physical world to guide behavior has become a major focus in neuroscience (1–3). In his 1943 book, Craik (4, p. 61) was perhaps the first to suggest that organisms maintain internal representations of the external world and to provide a rationale for their use:

If the organism carries a “small-scale model” of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, use the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer, and more competent manner to the emergencies that face it.

In this cognitive view of prospective simulation, an internal model allows an organism to contemplate the consequences of actions from its current state without actually committing itself to those actions. Since Craik’s initial proposal, internal models have become widely implicated in various brain subsystems with a diverse range of applications in biological control. Beyond facilitating the rapid and flexible modification of control policies in the face of changes in the environment, internal models provide an extraordinary range of advantages to a control system, from increasing the robustness of feedback corrections to distinguishing between self-generated and externally generated sensory input. However, there tends to be confusion as to what exactly constitutes an internal model. This confusion has likely arisen because the internal model hypothesis has independently emerged in distinct areas of neuroscientific research prompted by disparate computational motivations. Furthermore, there are intricate interactions between various types of internal models maintained by the brain. Here, we aim to provide a unifying account of biological internal models, review their adaptive benefits, and evaluate the empirical support for their use in the brain.

In order to accomplish this, we describe various conceptions of internal models within a common computational formalism based on the principle of rationality. This principle posits that an agent will endeavor to act in the most appropriate manner according to its objectives and the “situational logic” of its environment (5, p. 147; 6) and can be formally applied to any control task and data set. It provides a parsimonious framework in which to study the nervous system and the mechanisms by which solutions to sensorimotor tasks are generated. In particular, probabilistic inference (7) and optimal feedback control (OFC) (8) together provide concise computational accounts for many sensory and motor processes of biological control. In Section 2, we describe how these theories characterize optimal perception and action across a wide variety of scenarios. Recently, technical work has integrated these two theories into a common probabilistic framework by developing and exploiting a deeper theoretic equivalence (9, 10). This framework will provide the mathematical architecture necessary to integrate putative internal modeling mechanisms across a range of research areas, from sensorimotor control to behavioral psychology and cognitive science. In Section 3, we review theoretical arguments and experimental evidence supporting the contribution of internal models to the ability of nervous systems to produce adaptive behavior in the face of noisy and changing environmental conditions at many spatiotemporal scales of control.

2. INTERNAL MODELS IN THE PROBABILISTIC FRAMEWORK

Bayesian inference and optimal control have become mainstream theories of how the brain processes sensory information and controls movement, respectively (11). Their common theme is that behavior can be understood as an approximately rational solution to a problem defined by

task objectives and a characterization of the external environment, sensory pathways, and musculoskeletal dynamics—that is, they are normative solutions. In this section, we contextualize these theories in each of their respective domains of perception and action and review the experimental techniques employed to acquire evidence supporting their implementation in the nervous system.

2.1. Bayesian Inference in the Brain

In Bayesian inference, probabilities are assigned to each possible value of a latent state variable z one wishes to estimate, reflecting the strength of the belief that a given value represents the true state of the world (7). It is hypothesized that the brain encodes a prior $p(z)$ reflecting its beliefs regarding the state z before any sensory information has been received, as well as a probabilistic internal model describing the dependency of sensory signals y on the latent state z , known as a generative model in computational neuroscience (12). On receiving sensory information y , this probabilistic internal model can be used to compute a likelihood $p(y|z)$ that quantifies the probability of observing the signals y if a particular state z is true. Using these probabilistic representations of state uncertainty, Bayes's rule prescribes how the prior $p(z)$ and likelihood $p(y|z)$ are combined in a statistically optimal manner to produce the posterior probability distribution $p(z|y)$:

$$p(z|y) = \frac{p(y|z)p(z)}{p(y)}, \quad 1.$$

where $p(y) = \sum_z p(y|z)p(z)$ is known as the evidence for the observation y . In the context of sensory processing, Bayesian inference is proposed as a rational solution to the problem of estimating states of the body or environment from sensory signals afflicted by a variety of sources of uncertainty (**Figure 1a**). Sensory signaling is corrupted by noise at many points along the neural pathway, including transduction, action potential generation, and synaptic transmission (13). Furthermore, relevant state variables are typically not directly observable and therefore must be inferred from stochastic, statistically dependent observations drawn from multiple sensory modalities.

Several lines of behavioral evidence suggest that humans and other animals learn an internal representation of prior statistics and integrate this representation with knowledge of the noise in their sensory inputs in order to generate state estimates through probabilistic inference. First, many studies have exhibited a stimulus prior (e.g., the location of an object or duration of a tone) to a subject performing a task and shown that the prior is internalized and reflected in behavior (14–17). Importantly, as predicted by Bayes's rule, this prior bias is greater when the stimulus signal is less reliable and thus more uncertain. Second, other studies have assumed a reasonable prior so as to explain a range of phenomena and illusions as rational inferences in the face of uncertainty. For example, a prior over the direction of illumination of a scene (18–20) or over the speed of object motion (21) can explain several visual phenomena, such as how we extract shape from shading or perceive illusory object motion.

Beyond the sensorimotor domain, Bayesian methods have also been successful in explaining human reasoning. In the cognitive domain, the application of Bayesian principles using relatively complex probabilistic models has provided normative accounts of how humans generalize from few samples of a variable (22), make inferences regarding the causal structure of the world (23), and derive abstract rules governing the relationships between sets of state and sensory variables (24). Behavioral analyses that estimate high-dimensional cognitive prior representations from low-dimensional (e.g., binary) responses have been used to demonstrate that humans maintain a prior representation for faces and that this naturalistic prior is conserved across tasks (25).

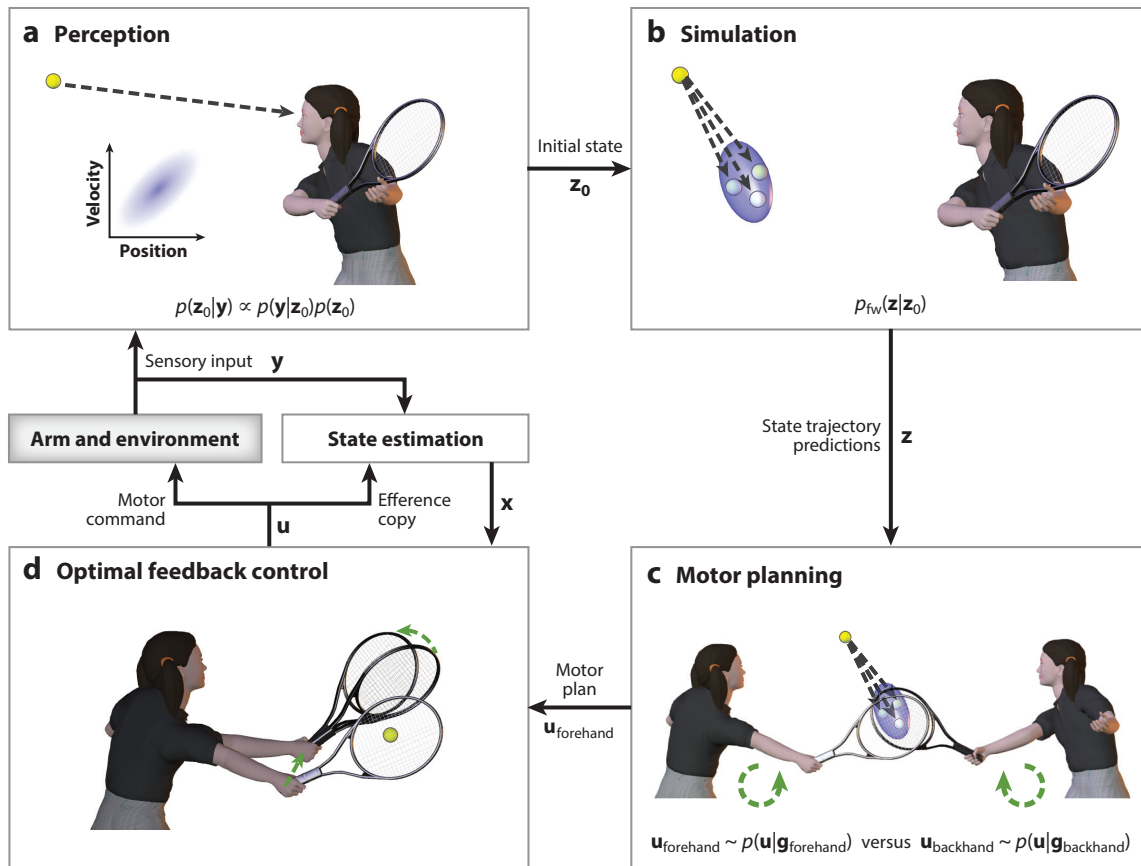


Figure 1

The roles of internal models in sensorimotor control. (*a*) Perception. Sensory input \mathbf{y} is used to estimate the ball's state \mathbf{z}_0 , which is uncertain due to noise along the sensory pathway and the inability to directly observe the full state of the ball (e.g., its spin and velocity). Bayes's rule is used to calculate the posterior; the inset shows an example of a posterior over one component of position and velocity. (*b*) Simulation. An internal dynamical model p_{fw} simulates the forward trajectory \mathbf{z} of the ball. At short timescales, this internal modeling is necessary to overcome delays in sensory processing, while at longer timescales, the predictive distribution $p_{\text{fw}}(\mathbf{z}|\mathbf{z}_0)$ of the ball's trajectory can be used for planning. (*c*) Motor planning. An internal simulation of the ball's trajectory and prospective movements is evaluated in order to generate an action plan. The player may have to decide between body reorientations in order to play a forehand or backhand. (*d*) Optimal feedback control. Once a motor plan has been specified, motor commands \mathbf{u} are generated by an optimal feedback controller that uses a state estimator to combine sensory feedback and forward sensory predictions (based on an efference copy of the motor command) in order to correct motor errors online in task-relevant dimensions (*green arrows*).

2.1.1. Bayesian forward modeling. Bayesian computations can be performed with respect to the current time or used to predict future states, as hypothesized by Craik. Consider the problem of tracking a ball during a game of tennis (see **Figure 1b**). The response of any given photoreceptor in a player's retina can provide only delayed, noisy signals regarding the position \mathbf{y} of the ball at a given time. From the probabilistic point of view, this irreducible uncertainty in the reported ball position is captured by a distribution $p(\mathbf{y})$. Since a complete characterization of the state \mathbf{z} of the tennis ball, including its velocity, acceleration, and spin, is not directly observable, this information must be inferred from position samples transduced from many photoreceptors at different time points in concert with the output of an internal model. Given a previously inferred

posterior $p(z_t|y_{:t})$ over possible ball states z_t based on previous sensory input $y_{:t}$ up to time t , an internal forward model $p_{\text{fw}}(z_{t+1}|z_t)$ can be used to predict the state of the ball at the future time step $t + 1$:

$$p(z_{t+1}|y_{:t}) = \int_{z_t} p_{\text{fw}}(z_{t+1}|z_t)p(z_t|y_{:t})dz. \quad 2.$$

The internal forward dynamical model p_{fw} must take physical laws, such as air resistance and gravity, into account. From a perceptual point of view, new sensory information y_{t+1} can then be integrated with this predictive distribution in order to compute a new posterior distribution at time $t + 1$:

$$p(z_{t+1}|y_{:t+1}) \propto p(y_{t+1}|z_{t+1})p(z_{t+1}|y_{:t}). \quad 3.$$

This iterative algorithm, known as Bayesian filtering, can be used to track states z_t, z_{t+1}, \dots of the body or the environment in the presence of noisy and delayed signals for the purposes of state estimation (see Section 3.2). The extrapolation of latent states over longer timescales can be used to predict states further into the future for the purposes of planning movement (see Section 3.3). The results of such computations are advantageous to the tennis player. On a short timescale, they enable the player to predictively track the ball with pursuit eye movements, while on a longer timescale, the player can plan to move into position well in advance of the ball's arrival in order to prepare the next shot.

In the brain, the dichotomy between the prediction step (based on a forward model) and the observation step is reflected, at least partially, in dissociated neural systems. With respect to velocity estimation, a detailed analysis of retinal circuitry has revealed a mechanism by which target velocity can be estimated at the earliest stages of processing (26). Axonal conductance delays endow retinal cells with spatiotemporal receptive fields that integrate information over time and fire in response to a preferred target velocity. Furthermore, the retina contains a rudimentary predictive mechanism based on the gain control of retina ganglion cell activity, whereby the initial entry of an object into a cell's receptive field causes it to fire, but the activity is then silenced (27). By contrast, more complex predictions (e.g., motion under gravity) require higher-order cortical processing.

2.1.2. Neural implementation. Theories have been developed regarding how neuronal machinery could perform the requisite Bayesian calculations. These theories fall into two main classes: population coding mechanisms in feedforward network architectures (28–31) and recurrently connected dynamical models (32–34). In the former, neural receptive fields are proposed to tile the sensory space of interest such that their expected firing rates encode the probability [or log probability (29)] of a particular value of the encoded stimulus. For example, this implies that each neuron in a population would stochastically fire within a limited range of observed positions of a reach target and fire maximally for its preferred value. Importantly, the variability in neural activity can then be directly related to the uncertainty regarding the precise stimulus values that generated the input in a manner consistent with Bayesian theory (28). Thus, across neurons, the population activity would reflect the posterior probability distribution of the target position given sensory input. This neural representation can then be fed forward to another layer of the network to produce a motor response. Such population codes are able to implement Bayes's rule in parsimonious network architectures and account for empirical neural activity statistics during sensorimotor transformations (30), Bayesian decision-making (35), and sensory computations such as cue integration (28), filtering (36), and efficient stimulus coding (31).

Although the functional implications of population codes can be directly related to Bayesian calculations, they do not incorporate the rich dynamical interactions between neurons in cortical

circuits or model the complex temporal profiles of neural activity that follow transient stimulus input (37, 38). These considerations have motivated the development of dynamical models of cortex with recurrent connectivity that approximate Bayesian inference (32, 34), though the characterization of this class of models from a computational point of view remains an ongoing challenge (39). In contrast to the probabilistic population coding approach, it has been postulated that neural variability across time reflects samples from a probability distribution based on a direct coding representation (40). In this model, population activity encodes sensory variable values (as opposed to the probability of a particular variable value) such that the variability of neural activity across time reflects the uncertainty in the stimulus representation. When sensory input is received, neural circuits generate samples from the posterior distribution of inferred input features. In the absence of external input, spontaneous activity corresponds to sampling from the prior distribution, which serves as an internal model of the sensory statistics of the environment. In support of this theory, the change in spontaneous visual cortical activity during development has been shown to be consistent with the gradual learning of a generative internal model of the visual environment, whereby spontaneous activity adapted to reflect the average statistics of all visual input (41).

2.2. Optimal Feedback Control

Bayesian inference is the rational mathematical framework for perception and state estimation based on noisy and uncertain sensory signals. Analogously, optimal control has been a dominant framework in sensorimotor control to derive control laws that optimize behaviorally relevant criteria and thus rigorously comply with the principle of rationality (11) (**Figure 1d**). Understanding how natural motor behavior arises from the combination of a task and the biomechanical characteristics of the body has driven the theoretic development of optimal control models in the biological context (42, 43). Initially, models were developed that posited that, for a given task, planning specified either the desired trajectory or the sequence of motor commands to be generated. These models typically penalized lack of smoothness, such as the time derivative of hand acceleration (known as jerk) (44) or joint torques (45). The role of any feedback was, at best, to return the system to the desired trajectory. These models aimed to provide a normative explanation for the approximately straight hand paths and bell-shaped speed profiles of reaching movements. However, these models are accurate only for movement trajectories averaged over many trials and do not account for the richly structured trial-to-trial variability observed in human motor coordination (8).

A fundamental characteristic of biological control is that the number of effector parameters to be optimized far exceeds the dimensionality of the task requirements. For example, infinitely many different time series of hand positions and joint angles can be used to achieve a task such as picking up a cup. Despite the plethora of possible solutions, motor behavior is relatively stereotyped both across a population and for an individual person, suggesting that the nervous system selects actions based on a prudent set of principles. How the brain chooses a particular form of movement out of the many possible is known as the degrees-of-freedom problem in motor control (46). A ubiquitous empirical observation in goal-directed motor tasks is that effector states tend to consistently covary in a task-dependent manner (8, 47–50). In particular, these covariances tend to be structured in such a way as to minimize movement variance along task-relevant dimensions while allowing variability to accumulate in task-irrelevant dimensions.

OFC was introduced (8, 11) in the motor control context in order to provide a normative solution to the degrees-of-freedom problem of motor coordination and, in particular, to develop a broad account of effector covariance structure and motor synergy as a function of task requirements. In this class of control laws, the core distinction with respect to optimal (feedforward or

desired trajectory) control is that sensory feedback is integrated into the production of motor output. OFC policies continually adapt to stochastic perturbations [for example, due to noise within the motor system (51)] and therefore predict temporal patterns of motor variability that have been widely tested in behavioral experiments. An emergent property of OFC, known as the minimum intervention principle, explains the correlation structures of task-oriented movements (8). Simply put, as movements deviate from their optimal trajectories due to noise, OFC specifically predicts that only task-relevant deviations will be corrected (8). For example, when reaching to a target that is either narrow or wide, subjects tend to make straight-line movements to the nearest point on the target (Figure 2a). However, when the hand is physically perturbed early in the movement, corrections are seen only when reaching toward the narrow target, not when reaching toward

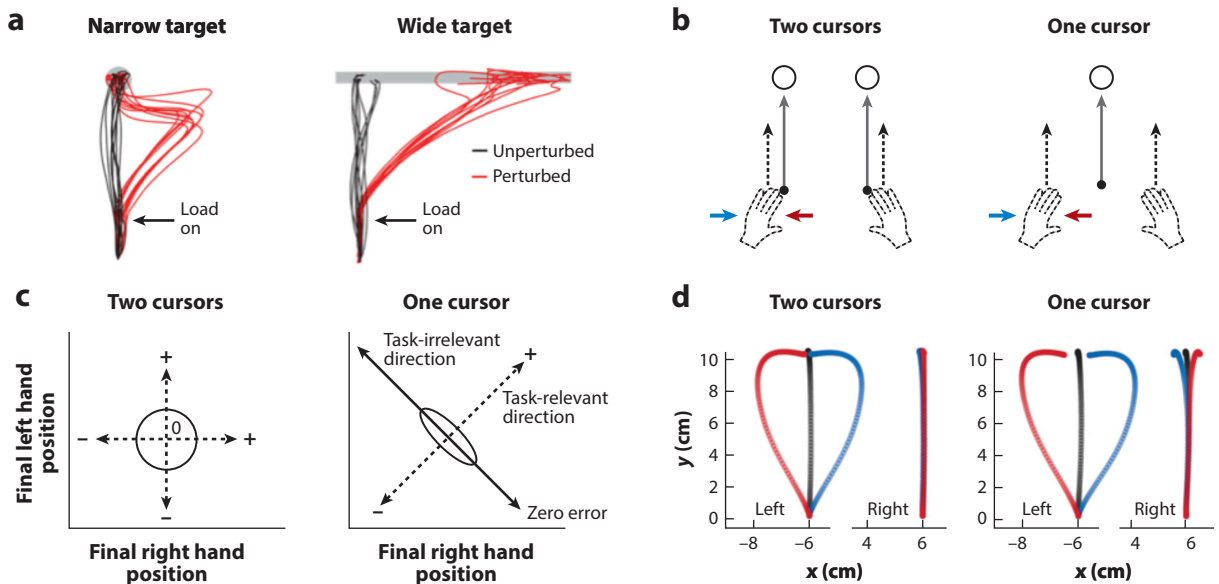


Figure 2

The minimum intervention principle and exploitation of redundancy. (a) Unperturbed movements (*black traces*, showing individual hand movement paths) to narrow or wide targets tend to be straight and to move to the closest point on the target. Hand paths during the application of mechanical loads (*red traces*, in response to a force pulse that pushes the hand to the right) delivered immediately after movement onset, which disrupt the execution of the planned movement, obey the principle of minimum intervention: For a narrow target (*left*), the hand paths correct to reach the target, whereas for a wide target (*right*), there is no correction, and the hand simply reaches to another point on the target. (b) Participants make reaching movements to targets using cursors. In a two-cursor condition, each hand moves its own cursor (*black dots*) to a separate target. In a one-cursor condition, the cursor is displayed at the average location of the two hands, and participants reach with both hands to move this common cursor to a single target. During the movement, the left hand could be perturbed with a leftward (*red*) or rightward (*blue*) force field or could remain unperturbed. (c) When each hand controls its own cursor, there is only one combination of final hand positions for which there is no error (*center of circle*). Optimal feedback control predicts that there will be no correlation between the endpoint positions (the *black circle* shows a schematic distribution of errors). When the two hands control the position of a single cursor, many combinations of final hand positions give zero error (*black diagonal line*, task-irrelevant dimension). Optimal feedback control predicts correction in one hand to deviations in the other, leading to negative correlations between the final locations of the two hands, so that if one hand is too far to the left, the other compensates by moving to the right (*black ellipse*). (d) This panel shows the movement trajectories for the left and right hands in response to the perturbations shown in panel b (one-cursor condition). The response of the right hand to perturbations of the left hand shows compensation only for the one-cursor condition, in accordance with the predictions of optimal feedback control. In addition, negative correlations in final hand positions can be seen in unperturbed movements for the one-cursor condition but not for the two-cursor condition (not shown). Panel a adapted from Reference 52 with permission; panels b–d adapted from Reference 53.

the wide target (**Figure 2a**); because the perturbation does not affect task success in the latter case, there is no reason to intervene. Intervening would actually be counterproductive, because it typically requires more energy and adds noise into the reach.

In sensorimotor control, the specification of a particular behavioral task begins with a definition of what constitutes the relevant internal state x (which may include components corresponding to the state of the arm and external environment) and control signals u . In general, the state variables should include all the variables, which, together with the equations of motion describing the system dynamics and the motor commands, are sufficient to predict future configurations (in the absence of noise). A discrete-time stochastic dynamics model can then be specified that maps the current state x_t and control inputs u_t to future states x_{t+1} . This model is characterized by the conditional probability distribution $p_{\text{env}}(x_{t+1} | x_t, u_t)$. For reaching movements, for example, the state x could correspond to the hand position, joint angles, and angular velocities, and the control signals u might correspond to joint torques. Given these dynamics, the aim of optimal control is to minimize a cost function that includes both control and state costs. The state cost Q rewards states that successfully achieve a task (such as placing the hand on a target), while R represent an energetic cost such as that required to contract muscles (for a discussion of cost function specification in the biological context, see the sidebar titled Costs, Rewards, Priors, and Parsimony). To make predictions regarding motor behavior, a control policy π [a mapping from states to control signals $u_t = \pi(x_t)$] is optimized to minimize the total cumulative costs expected to be incurred. This objective $V_\pi(x_t)$ is known as the cost-to-go of a control policy (in control

COSTS, REWARDS, PRIORS, AND PARSIMONY

Critics of optimal control theories of motor control point out that one can always construct a cost function to explain any behavioral data (at the extreme, the cost can be the deviations of the movement from the observed behavior). Therefore, to be a satisfying model of motor control, it is crucial that the assumed costs, rewards, and priors be well motivated and parsimonious. Initial work on optimal motor control used cost functions that did not correspond to ecologically relevant quantities. For example, extrinsic geometric smoothness objectives such as jerk (44) or the time derivative of joint torque (45) do not straightforwardly relate to biophysically important variables. By contrast, OFC primarily penalizes two components in the cost. The first is an energetic or effort cost. Such costs are widespread in animal behavior modeling and provide well-fitting cost functions when simulating muscle contractions (54) and walking (55, 56), suggesting that such movements tend to minimize metabolic energy expenditure. By representing effort as an energetic cost discounted in time, one can account for both the choices animals make and the vigor of their movements (57). The second penalized component, task success, is typically represented by a cost on inaccuracy.

When experimenters place explicit costs or rewards on a task (such as movement target points), people are usually able to adapt their control to be close to optimal in terms of optimizing such explicit objectives (58–60). The parsimony and the experimental benefits of a model where the experimenter specifies costs at the task level are not present in oracular motor control models, which requires an external entity to provide a detailed prescription for motor behavior. Early theories of biological movement were often inspired by industrial automation. Research tended to focus on how reference trajectories for a particular task were executed rather than planned. For any given task, there are infinitely many trajectories that reach a desired goal and infinitely many others that do not, and the problem of selecting one is off-loaded to a trajectory oracle, reminiscent of industrial control engineers serving as the *deus ex machina*. As a theory of biological movement, this is problematic. Oracles can select movement trajectories not necessarily to solve the task in an optimal manner (as would be the goal in industrial automation) but rather to fit movement data, which leads to an overfitting problem (7).

theory) or value function (in reinforcement learning, where it typically quantifies cumulative expected rewards rather than costs):

$$V_{\pi}(x_t) = Q(x_t) + R[\pi(x_t)] + \mathbb{E}_{x_{t+1} \sim p_{\text{env}}[\cdot|x_t, \pi(x_t)]} [V_{\pi}(x_{t+1})]. \quad 4.$$

This characterization of the cost-to-go function, known as a Bellman equation, intuitively implies that the optimal controller balances the instantaneous costs in the current state x_t with the minimization of expected future cumulative costs in the subsequent state x_{t+1} .

This formulation is quite general. When applied to motor behavior, costs are often modeled as a quadratic function of states and control signals, while the dynamics model $p_{\text{env}}(x_{t+1}|x_t, u_t)$ typically takes the form of a linear equation with additive Gaussian noise (43). Furthermore, the noise term is adapted to scale with the magnitude of the control input, as found in the nervous system (51). This signal-dependent noise arises through the organization of the muscle innervation. The force that a single motor neuron can command is directly proportional to the number of muscle fibers that it innervates. When small forces are generated, motor neurons that innervate a small number of muscle fibers are active. When larger forces are generated, additional motor neurons that innervate a larger number of muscle fibers are also active. This is known as Henneman's size principle. Recruiting a larger number of muscle fibers from a single alpha motoneuron (the final neuronal output of the motor system) increases the variability of the output, leading to variability in the force that is proportional to the average force produced by that muscle (61, 62). This OFC problem formulation provides a reasonable balance between capturing the essential features of the sensorimotor task and enabling the accurate computation of optimal control policies; linear-quadratic-Gaussian problems with signal-dependent noise can be solved by the iteration of two matrix equations that converge exponentially fast (43).

Variants of this OFC model have been tested in many experiments involving a variety of effectors, task constraints, and cost functions (48, 49, 63–66). For example, studies have examined tasks in which a person's hands either move separate cursors to individual targets or together move a single cursor (whose location is the average position of the two hands) to a single target (**Figure 2b**). The predictions of OFC differ for these two scenarios (**Figure 2c**). In the former, perturbations to each arm can be corrected only by that arm, so a perturbation to one arm should be corrected only by that arm. However, in the latter situation, both arms could contribute to the control of the cursor, so perturbations to one arm should also be corrected by the other arm. Indeed, force perturbations of one hand result in corrective responses in both hands, consistent with an implicit motor synergy, as predicted by OFC (**Figure 2d**). Moreover, in a directed force production task, a high-dimensional muscle space controls a low-dimensional finger force. Electromyography recordings revealed task-structured variability in which the task-relevant muscle space was tightly controlled and the task-irrelevant muscle space showed much greater variation, again confirming predictions of OFC (64).

OFC is also a framework in which active sensing can be incorporated. Although engineering models typically assume state-independent noise, in the motor system the quality of sensory input can vary widely. For example, a person's ability to localize their hand proprioceptively varies substantially over the reaching workspace. Including state-dependent noise in OFC means that the quality of sensory input will depend on the actions taken. The solution to OFC leads to a trade-off between making movements that allow one to estimate the state accurately and task achievement. The predictions of the optimal solution match those seen in human participants when they are exposed to state-dependent noise (67).

Recent work has focused on the adaptive feedback responses within an OFC framework. One way to measure the magnitude of the visuomotor response (positional gain) is to apply lateral visual perturbations to the hand during a reaching movement. On such a visually perturbed trial,

a robotic interface is typically used to constrain the hand within a simulated mechanical channel so that the forces into the channel are a reflection of the visuomotor reflex gain. Such studies have shown that the reflex gains are sensitive to the task and that the gains increase when the perturbation is task relevant and decrease when it is not (63). Moreover, the reflex gain varies throughout a movement in a way that qualitatively agrees with the predictions of OFC (66). Reflexive responses due to muscle stretch caused by mechanical perturbation can be decomposed into short-latency (<50 ms) and long-latency (50–100 ms) components, both of which occur before the onset of volitional control (>100 ms) (68). Short-latency components are generated by a spinal pathway (i.e., the transformation of proprioceptive feedback into motor responses occurs at the level of the spinal cord), while long-latency components are transcortical in nature (i.e., the cortex is involved in modulating the reflex). The long-latency response specifically can be voluntarily manipulated based on the behavioral context (69), and it has been suggested that this task-based flexibility is consistent with an optimal feedback controller operating along a pathway through the primary motor cortex (70). Neural activity in the primary motor cortex reflects both low-level sensory and motor variables (71) while also being influenced by high-level task goals (72). This diversity of encoding is precisely what one would expect from an optimal feedback controller (73). Further evidence in favor of this hypothesis includes the fact that primary motor cortex neurons appear to encode the transformation of shoulder and elbow perturbations into feedback responses (74).

2.3. Duality Between Bayesian Inference and Optimal Control

Classically, a control policy $u = \pi(x)$ deterministically maps states to control signals. However, in the probabilistic framework, it is more natural to consider stochastic policies $p(u|x)$ representing distributions over possible control commands conditioned on a given state. Furthermore, it is impossible for the brain to represent a deterministic quantity with perfect precision; therefore, probabilistic representations may be a more appropriate technical language in the sensorimotor control context (75). This probabilistic perspective allows us to review a general duality between control and inference.

It has long been recognized that certain classes of Bayesian inference and optimal control problems are mathematically equivalent or dual. Such an equivalence was first established between the Kalman filter and the linear-quadratic regulator (76) and has recently been generalized to nonlinear systems (9, 77). The intuition is as follows. Suppose a person is performing a goal-directed reaching movement and wants to move their hand to a target. The problem of identifying the appropriate motor commands can be characterized as the minimization of a cost-to-go function (Equation 4). However, an alternative but equivalent approach can be considered: The person could imagine their hand successfully reaching the target at some point in the future and infer the sequence of motor commands that was used to get there. The viewpoint transforms the control problem into an inference problem.

More technically, the duality can be described using trajectories of states $\mathbf{x} := (x_0, \dots, x_T)$ and control signals $\mathbf{u} := (u_0, \dots, u_{T-1})$ up to a horizon T . Consider the conditional probability defined by $p(\mathbf{g}|\mathbf{x}) \propto \exp[-Q(\mathbf{x})]$, where $Q(\mathbf{x}) := \sum_{i=0}^T Q(x_i)$ is the above-mentioned state-dependent cost encoding the desired outcome (Equation 4). The variable \mathbf{g} can be thought of as a pseudo-observation of a successfully completed goal-directed task. The task is considered to be more likely to be successful if less state costs are incurred. The control cost $R(\mathbf{u}) := \sum_{i=0}^{T-1} R(u_i)$ can be absorbed in a prior over control signals $p(\mathbf{u}) \propto \exp[-R(\mathbf{u})]$, with more costly control commands [large $R(\mathbf{u})$] being more unlikely a priori. Bayesian inference can then be employed to compute the joint probability of motor outputs \mathbf{u} and state trajectories \mathbf{x} given the observation of

a successful task completion \mathbf{g} :

$$p(\mathbf{x}, \mathbf{u}|\mathbf{g}) \propto p_{\text{env}}(\mathbf{x}|\mathbf{x}_0, \mathbf{u})p(\mathbf{g}|\mathbf{x})p(\mathbf{u}) = p_{\text{env}}(\mathbf{x}|\mathbf{x}_0, \mathbf{u})e^{-Q(\mathbf{x})}e^{-R(\mathbf{u})}. \quad 5.$$

The posterior probabilities of control signals \mathbf{u} that are most likely to lead to a successful completion of the task \mathbf{g} along a particular state trajectory \mathbf{x} are proportional to the expected cumulative costs, as in the optimal control perspective (Equation 4). By marginalizing over state trajectories \mathbf{x} , one obtains the posterior $p(\mathbf{u}|\mathbf{g})$ as a sum-over-paths of the costs incurred (78). This perspective has led to theoretic insights within a class of control problems known as Kullback-Leibler control (10) or linearly solvable Markov decision processes (79), where the control costs take the form of a Kullback-Leibler divergence. In particular, this class of stochastic optimal control problems is formally equivalent to graphical model inference problems (10) and is a relaxation of deterministic optimal control (80). Thus, approximate inference methods, which have provided inspiration for neural and behavioral models of the brain's perceptual processes, may also underpin the algorithms used by the brain during planning (see Section 3.3).

2.4. What Constitutes an Internal Model in the Nervous System?

In neuroscience, neural representations of a person's body or environment—that is, internal models—are conceptualized in a wide range of theories regarding how the brain interprets, predicts, and manipulates the world. Most generally, one may consider a representation of the joint distribution $p(\mathbf{x}, \mathbf{z}, \mathbf{y}, \mathbf{u})$ between time series of sensory inputs \mathbf{y} , latent states \mathbf{z} , internal states \mathbf{x} , and motor signals \mathbf{u} . Together, the latent external states \mathbf{z} and internal states \mathbf{x} reflect the state of the world from the point of view of the nervous system, but we separate them conceptually to reflect a separation between environmental and bodily states. This probabilistic representation can be considered a complete internal model. Such a formulation contains within it various characterizations of internal models from different disciplines of neuroscience as conditional densities. Therefore, the phrase internal model can be used for markedly different processes, and we suggest that it is important for researchers to be explicit about what type of internal model they are investigating in a given domain. Here, we attempt to nonexhaustively categorize the elements that can be considered part of an internal model in sensorimotor control:

- **Prior models:** These models comprise priors over sensory signals, $p(\mathbf{y})$; motor signals, $p(\mathbf{u})$; and states of the world, $p(\mathbf{z})$ and $p(\mathbf{x})$. The world is far from homogeneous, and numerous studies have shown that people are adept at learning the statistical regularities of sensory inputs and the distributions of latent states (for a review, see Reference 40).
- **Perceptual inference models:** A class of internal models known in computational neuroscience as recognition models compute latent world states (such as objects) given sensory input, $p(\mathbf{z}|\mathbf{y})$, and are postulated to be implemented along higher-order sensory pathways culminating in the temporal lobes. Generative models, by contrast, are models that describe processes that generate sensory data. A generative model may be captured by the joint distribution between sensory input and latent variables, $p(\mathbf{y}, \mathbf{z})$, or computed from the product of a state prior and the conditional distribution of sensory inputs given latent world states, $p(\mathbf{y}|\mathbf{z})$. Given sensory input, the generative model can be inverted via Bayes's rule to compute the probabilities over the latent states that may have generated the observed input. Further uses of such generative models are predictive coding (81) and refference cancellation (see Section 3.1).
- **Sensory and motor noise models:** The brain is sensitive to the noise characteristics and reliability of our sensory and motor apparatus (13). On the sensory side, to calculate $p(\mathbf{y}|\mathbf{z})$ involves not only a transformation but also knowledge of the noise on the sensory signal \mathbf{y}

itself. On the motor side, the control output \mathbf{u} is also corrupted by noise, and knowledge of this noise can be used to refine the probability distribution of future states \mathbf{x} . Maintaining such noise models aids the nervous system in accurately planning and implementing control policies that are robust to sensory and motor signal corruption (82).

- Forward dynamical models: In general, we think of a forward dynamical model as a neural circuit that can take the present estimated state, \mathbf{x}_0 , and predict states in the future. This could model the passive dynamics, $p(\mathbf{x}|\mathbf{x}_0)$, of the system or also use the current motor output to predict the state evolution, $p(\mathbf{x}|\mathbf{x}_0, \mathbf{u})$.
- Cognitive maps, latent structure representation, and mental models: Abstract relational structures between state variables (possibly pertaining to distinct objects in the world) may be compactly summarized in the conditional probability distributions $p(\mathbf{z}^n|\mathbf{z}^1, \dots, \mathbf{z}^{n-1})$ of a graphical model. Such representations can also be embedded in continuous internal spaces such that a metric on the space encodes the relational strength between variables. These models can be recursively organized in hierarchies, thus facilitating the low-dimensional encoding of control policies and the transfer of learning across contexts (for a review of latent structure learning in the context of motor control, see Reference 83).

The probabilistic formalism allows one to relate internal models across a range of systems within the brain. However, it leaves many aspects of the internal models unspecified. Internal models can be further defined by a structural form that links inputs and outputs. For example, they may capture linear or nonlinear relationships between motor outputs and sensory inputs, as in the relationship between joint torques and future hand positions. They may contain free parameters that can be quickly adapted in order to adapt to contextual variations, such as the length and inertia of limbs during development. They can also be further specified by the degree of approximation in the model implementation. Consider the problem of predicting the future from the past. At one extreme, one can generate simulations from a rich model containing internal variables that directly reflect physically relevant latent states, such as gravitational forces and object masses. On the other hand, a mapping from current to future states can be learned directly from experience without constructing a rich latent representation. Such mappings can be encapsulated compactly in simple heuristic rules, which may provide a good trade-off between generalizability and efficiency. Finally, internal models span a range of spatiotemporal resolutions. Some internal models, such as those involved in state estimation, compute forward dynamics on very short spatiotemporal scales, such as centimeters and milliseconds (see Section 3.2), while others, such as those used during planning, simulate over timescales that are orders of magnitude longer, such as kilometers and days (see Section 3.3).

2.5. Probabilistic Forward and Inverse Models

In the sensorimotor context, internal models are broadly defined as neural systems that mimic musculoskeletal or environmental dynamical processes (84, 85). An important feature of putative internal models in sensorimotor control is their dynamical nature, which distinguishes internal models from other neural representations of the external world that the brain maintains, such as recognition models, as studied in perception. This dynamical nature is reflected in the brain computations associated with internal models. Whether contributing to state estimation, refference cancellation, or planning, internal forward and inverse models relate world states across a range of temporal scales. In the tennis example described above, internal models may be used to make anticipatory eye movements in order to overcome sensory delays in tracking the ball. By incorporating a motor response, internal models can be used to simulate the ballistic trajectory of a tennis ball after it has been struck. This leads to a classical theoretic dissociation of internal

models into different classes (85). Internal models that represent future states of a process (ball trajectories) given motor inputs (racquet swing) are known as forward models. Conversely, models that compute motor outputs (the best racquet swing) given the desired state of the system at a future time point (a point-winning shot) are known as inverse models.

In the probabilistic formalism, the internal forward model p_{fw} can be encapsulated by the distribution over possible future states x_{t+1} given the current state x_t and control signals u_t :

$$p_{fw}(x_{t+1}|x_t, u_t). \quad 6.$$

A prediction regarding a state trajectory $\mathbf{x} := (x_0, \dots, x_T)$ can be made by repeatedly applying the forward model $p_{fw}(\mathbf{x}|\mathbf{x}_0, \mathbf{u}) = \prod_{i=1}^T p_{fw}(x_i|x_{i-1}, u_{i-1})$. By combining a forward model p_{fw} and a prior over controls $p(\mathbf{u})$, the inverse model p_{inv} can be described in the probabilistic formalism using Equation 5. Consider the problem of computing the optimal control signals that implement a movement toward a desired goal state \mathbf{g} . This state could be, for example, the valuable target position of a reach movement. An inverse model is then a mapping from this desired state to a control policy \mathbf{u}^* that can be identified with the posterior probability distribution computed via control inference (Equation 5):

$$p_{inv}(\mathbf{u}|\mathbf{g}) \propto \int_{\mathbf{x}} p_{fw}(\mathbf{x}|\mathbf{x}_0, \mathbf{u}) p(\mathbf{g}|\mathbf{x}) p(\mathbf{u}) d\mathbf{x}, \quad 7.$$

$$\mathbf{u}^* = \underset{\mathbf{u}}{\operatorname{argmax}} p_{inv}(\mathbf{u}|\mathbf{g}). \quad 8.$$

Typically, in the sensorimotor control literature, a mapping from desired states at each point in time to the control signals \mathbf{u}^* is described as an inverse model. This mapping requires the explicit calculation of a desired state trajectory \mathbf{x}^* . This perspective can be embedded within the probabilistic framework by setting $p(\mathbf{g}|\mathbf{x}^*) = 1$ and $p(\mathbf{g}|\mathbf{x}) = 0$ for all other state trajectories $\mathbf{x} \neq \mathbf{x}^*$. By contrast, in OFC and reinforcement learning, motor commands are generated based on the current state without the explicit representation of a desired state trajectory. Alternatively, motor commands may depend on previous control signals independent of the current state. Such approaches to policy representation can serve as models of motor chunking in sensorimotor control.

3. THE ROLES OF INTERNAL MODELS IN BIOLOGICAL CONTROL

3.1. Sensory Reafference Cancellation

Sensory input can be separated into two streams: afferent information, which is information that comes from the external world, and reafferent information, which is sensory input that arises from our own actions. From a sensory receptors point of view, these sources cannot be separated. However, it has been proposed that forward models are a key mechanism that allows us both to determine whether the sensory input we receive is a consequence of our own actions and to filter out the components arising from our own actions so as to be more attuned to external events, which tend to be more behaviorally important (86). To achieve this, a forward model receives a signal of the outgoing motor commands and uses this so-called efference copy to calculate the expected sensory consequences of an ongoing movement (87). This predicted reafferent signal (known as the corollary discharge in neurophysiology, although this term is now often used synonymously with efference copy) can then be removed from incoming sensory signals, leaving only sensory signals due to environment dynamics.

This mechanism plays an important role in stabilizing visual perception during eye movements. When the eyes make a saccade to a new position, the sensory representation of the world shifts across the retina. In order for the brain to avoid concluding that the external world has been

displaced based on this retinal flow, a corollary discharge is generated from outgoing motor commands and integrated into the visual processing of the sensory input (88). A thalamic pathway relays signals about upcoming eye movements from the superior colliculus to the frontal eye fields, where it causally shifts the spatial receptive fields of target neurons in order to cancel the displacement due to the upcoming saccade (89). Furthermore, the resulting receptive field shifts are time locked by temporal information pertaining to the timing of the upcoming saccade carried by the corollary discharge.

Perhaps the best-worked-out example of the neural basis of such a predictive model is in the cerebellum-like structure of the weakly electric fish (90). These animals generate pulses (or waves) of electrical discharge into the water and can then sense the field that is generated to localize objects. However, the field depends on many features that the fish controls, such as the timing of the discharge and the movement and posture of the fish. The cerebellum-like structure learns to predict the sensory consequences (i.e., the field) based on both sensory input and the motor command and remove this from the signal so that any remaining signal reflects an unexpected input that pertains to objects in the environment. A recent review (91) elucidated the detailed mechanism of synaptic modulation (anti-Hebbian learning) and the manner in which the sensory prediction is built up from a set of basis functions.

3.2. Forward State Estimation for Robust Control

An estimate of the current state of an effector is necessary for both motor planning and control. There are only three sources of information that can be used for state estimation: sensory inputs, motor outputs, and prior knowledge. In terms of sensory input, the dominant modality for such state estimation is proprioceptive input (i.e., from receptors in the skin and muscles). While blind and deaf people have close to normal sensorimotor control, the rare patients with loss of proprioceptive input are severely impaired in their ability to make normal movements (92, 93). The motor signals that generate motion can also provide information about the likely state of the body. However, to link the motor commands to the ensuing state requires a mapping between the motor command and the motion—that is, a forward dynamic model (2)—in an analogous fashion to many observer models in control theory. There are at least two key benefits of such an approach. First, the output of the internal model can be optimally combined with sensory inflow via Bayesian integration (Section 2.1), minimizing state estimation variance due to noise in sensory feedback (94). Second, using the motor command (which is available in advance of the change in state) with the internal model makes movement more robust with respect to errors introduced by the unavoidable time delays in the sensorimotor loop. Feedback-based controllers with delayed feedback are susceptible to destabilization since control input optimized for the system state at a previous time point may increase, rather than decrease, the motor error when applied in the context of the current (unknown) state (85). Biological sensorimotor loop delays can be on the order of 80–150 ms for proprioceptive to visual feedback (68). However, a forward model that receives an efferent copy of motor outflow and simulates upcoming states can contribute an internal feedback loop to effect feedback control before sensory feedback is available (2, 3).

3.2.1. State estimation and sensorimotor control. Predictive control is essential for the rapid movements commonly observed in dexterous behavior. Indeed, this predictive ability can be demonstrated easily with the so-called waiter task. If you hold a weighty book on the palm of your hand with an outstretched arm and use your other hand to remove the book (like a waiter removing objects from a tray), the supporting hand remains stationary. This shows our ability to anticipate events caused by our own movements in order to generate the appropriate and

exquisitely timed reduction in muscle activity necessary to keep the supporting hand still. By contrast, if someone else removes the book from your hand, even with vision of the event, it is close to impossible to keep the hand stationary even if the removal is entirely predictable (95).

Object manipulation also exhibits an exquisite reliance on anticipatory mechanisms. When an object is held in a precision grip, enough grip force must be generated to prevent the object from slipping. The minimal grip force depends on the object load (i.e., weight at rest) and the frictional properties of the surface. Subjects tend to maintain a small safety margin so that if the object is raised, the acceleration causes an increase in the load force, requiring an increase in the grip force to prevent slippage. Recordings of the grip and load force in such tasks show that the grip force increases with no lag compared with the load force even in the initial phase of movement, thus ruling out the possibility that grip forces were adapted based on sensory feedback (96, 97). Indeed, such an anticipatory mechanism is very general, with no lag in grip force modulation seen when a person jumps up and down while holding the object. By contrast, if the changes in load force are externally generated, then compensatory changes in grip force lag by approximately 80 ms, suggesting a reactive response mechanism (98).

In contrast to internal models that estimate the state of the body based on efferent copies, internal models of the influence of external environmental perturbations are also utilized in state estimation. An analysis of postural responses to mechanical perturbations showed that long-latency feedback corrections were consistent with a rapid Bayesian updating of estimated state based on forward modeling of delayed sensory input (99). Furthermore, trial-to-trial changes in the motor response suggested that the brain rapidly adapted to recent perturbation statistics, reflecting the ability of the nervous system to flexibly alter its internal models when exposed to novel environmental dynamics. Although forward modeling can be based on both proprioceptive and visual information, the delays in proprioceptive pathways can be several tens of milliseconds shorter than those in visual pathways. During feedback control, the brain relies more heavily on proprioceptive information than on visual information (independent of the respective estimation variances), consistent with an optimal state estimator based on multisensory integration (100).

Certain actions can actually make state estimation easier, and there is evidence that people may expend energy to reduce the complexity of state estimation. For example, in a task analogous to sinusoidally translating a coffee cup without spilling its contents, people choose to move in a way that makes the motion of the contents more predictable, despite the extra energetic expense that this requires (101). Such a strategy could potentially minimize the computational complexity of internal forward modeling and thereby reduce errors in state estimation.

3.2.2. Neural substrates. Extensive research has been conducted with the aim of identifying the neural loci of putative forward models for sensorimotor control. Two brain regions in particular have been implicated: the cerebellum and the parietal cortex. It has long been established that the cerebellum is important for motor coordination. Although patients with cerebellar damage can generate movement whose gross structure matches that of a target movement, their motions are typically ataxic and characterized by dysmetria (typically the overshooting or undershooting of target positions during reaching) and oscillations when reaching (intention tremor) (102). In particular, these patients experience difficulty in controlling the inertial interactions among multiple segments of a limb, which results in greater inaccuracy of multijoint movements compared with single-joint movements. An integrative theoretic account (2, 103) suggested that these behavioral deficits could be caused by a lack of internal feedback and thus that the cerebellum may contain internal models that play a critical role in stabilizing sensorimotor control. A range of investigations across multiple disciplines has supported this hypothesis, including electrophysiology (104–106), neuroimaging (97), lesion analysis (103, 107), and noninvasive stimulation (108). In particular, the

above-mentioned ability of humans to synchronize grip force with lift, which provided indirect behavioral evidence of an internal forward model, is impaired in patients with cerebellar degeneration (107). Optimal control models have enabled researchers to estimate impairments of the forward dynamic models in cerebellar patients making dysmetric reaching movements (109). In this study, hypermetric patients appeared to overestimate arm inertia, leading them to overshoot the target, while hypometric patients tended to underestimate arm inertia, resulting in the opposite pattern of deviations from optimality. The authors were therefore able to compute dynamic perturbations that artificially increased (for hypermetric patients) or decreased (for hypometric patients) arm inertia, thus compensating for the idiosyncratic biases of individual patients. This study highlights the contribution of optimal control and internal models to a detailed understanding of a particular movement disability and the possibility of therapeutic intervention.

The parietal cortex has also been implicated in representing forward state estimates. A subregion of the superior parietal lobule known as the posterior parietal cortex contains neural activity consistent with forward state estimation signaling (110), which may be utilized for visuomotor planning (111). Indeed, transcranial magnetic stimulation of this region, resulting in transient inhibition of cortical activity, impaired the ability of subjects to error correct motor trajectories based on forward estimates of state (112). In another study, following intracranial electrical stimulation of the posterior parietal cortex, subjects reported that they had made various physical movements even though they had not actually done so and electromyography had detected no muscle activity (113). This illusory awareness of movement is consistent with the activation of a forward state representation of the body. A study based on focal parietal lesions in monkeys reported a double dissociation between visually guided and proprioceptively guided reach movement impairments and lesions of the inferior and superior parietal lobules, respectively (114). This finding suggests that forward representations of state are localized to different areas of the posterior parietal cortex depending on the sensory source of state information.

3.3. Learning and Planning Novel Behaviors

The roles of internal models described thus far operate on relatively short timescales and do not fit Craik's original conception of their potential contribution to biological control, which concerned the internal simulation of possible action plans over longer timescales in order to predict and evaluate contingent outcomes. Through the computational lens of optimal control, Craik's fundamental rationale for internal modeling falls within the broad domain of algorithms by which the brain can acquire new behaviors, which we review in this section.

3.3.1. Reinforcement learning and policy optimization. Control policies can be optimized using a range of conceptually distinct but not mutually exclusive algorithms, including reinforcement learning (115) and approximate inference (116). Reinforcement learning provides a suite of iterative policy-based and value-based optimization methods that have been applied to solve OFC problems. Indeed, initial inspiration for reinforcement learning was derived from learning rules developed by behavioral psychologists (117). Theoretical and empirical analyses of reinforcement learning methods indicate that a key algorithmic strategy that can aid policy optimization is to learn estimates of the cost-to-go function V_π introduced in Section 2.2. Once V_π is known, the optimal controls $u^*(x_t)$ are easily computed without explicit consideration of the future costs [by selecting the control output that is most likely to lead to the subsequent state x_{t+1} with minimal $V_\pi(x_{t+1})$]. A related and even more direct method is to learn and cache value estimates (known as Q-values) associated with state-action combinations (115). Thus, value estimates are natural

quantities for the brain to represent internally, as they are the long-term rationales for being in a given state and define optimized policies.

In many reinforcement learning algorithms, a key signal is the prediction error, which is the difference between expected and actual rewards or costs. This signal can be used to iteratively update an estimate of the cost-to-go and is guaranteed to converge to the correct cost-to-go values (although the learning process may take a long time) (115). Neural activity in the striatum of several mammalian species (including humans) appears to reflect the reinforcement learning of expected future reward representations (118, 119). Indeed, reward-related neurons shift their firing patterns in the course of learning, from signaling reward directly to signaling the expected future reward based on cues associated with later reward, consistent with a reward prediction error based on temporal differences (118).

The main shortcoming of such model-free methods for learning optimal control policies is that they are prohibitively slow. When these methods are applied to naturalistic motor control tasks with high-dimensional, nonlinear, and continuous state spaces (corresponding to the roughly 600 muscles controlled by the nervous system), potentially combined with complex object manipulation, it becomes clear that human motor learning is unlikely to be based on these methods alone due to the time required to produce control policies with human-level performance. Furthermore, environment dynamics can transform unexpectedly, and the goals of an organism may change depending on a variety of factors. Taken together, all of this suggests that humans and animals must integrate alternative algorithms in order to flexibly and rapidly adapt their behavior. In particular, internal forward models can be used to predict the performance of candidate control strategies without actually executing them, as originally envisaged by Craik (4) (**Figure 1c**). These internal model simulations and evaluations (which operate over relatively long timescales compared with the internal forward models discussed above) can be integrated with reinforcement learning (115) and approximate inference methods (120). Thus, motor planning may be accomplished more quickly and robustly using internal forward models. Indeed, trajectory rollouts (121) and local searches (122) form key components of many state-of-the-art learning systems.

3.3.2. Prediction for planning. Planning refers to the process of generating novel control policies internally rather than learning favorable motor outputs from repeated interactions with the environment (**Figure 1c**). Internal forward modeling on timescales significantly longer than those implemented in state estimation contributes significantly at this point in the sensorimotor control process. Ultimately, once a task has been specified and potential goals identified, the brain needs to generate a complex spatiotemporal sequence of muscle activations. Planning this sequence at the level of muscle activations is computationally intractable due to the curse of dimensionality (123). Specifically, the number of states (or volume, in the case of a continuous control problem) that must be evaluated scales exponentially with the dimensionality of the state space. This issue similarly afflicts the predictive performance of forward dynamic models, where state-space dimensionality is determined by the intricate structure and nonstationarity of the musculoskeletal system and the wider external world. Biological control hierarchies have been described across the spectrum of behavioral paradigms, from movement primitives and synergies in motor control (124) to choice fragments in decision-making (125). From a computational efficiency perspective, these hierarchies allow low-level, partially automated components to be learned separately but also flexibly combined in order to generate broader solutions in a hierarchical fashion, thus economizing control by enabling the nervous system to curtail the number of calculations it needs to make (126). For example, one learns to play the piano not by going through music note by note, but rather by practicing layers and segments of music in isolation before combining these fluent chunks together (127).

Given the hierarchical structure of the motor system, motor commands may be represented, and thus planned, at multiple levels of abstraction. Different levels of abstraction are investigated in distinct fields of neuroscience research that focus on partially overlapping subsystems. However, here we take a holistic view and do not focus on arbitrary divisions between components of an integrated control hierarchy. At the highest level, if multiple possible goals are available, a decision may be made regarding which is to be the target of movement. Neuroimaging (128) and single-unit recordings (129) suggest that scalar values associated with goal states are encoded in an area of the brain known as the ventromedial prefrontal cortex. Comparing such value signals allows a target to be established. Selection among food options is often used to study neural value representation since food is a primary reinforcer. In such an experiment, when confronted with novel goals that have never been encountered before, the brain synthesizes value predictions from memories of related goals in order to make a decision (130). The precise mechanism by which this is accomplished is still under investigation, but these results require an internal representation that is sensitive to the relational structure among food items, possibly embedded in a feature space of constituent nutrients, and a generalization mechanism with which new values can be constructed.

This internal representation and mechanism can be embedded in the probabilistic framework described here. Let x be a vector of goal features. The value v can then be modeled as the latent variable to be inferred, and a value model $p(v|x)$ can be learned using experienced goal-value pairs and used to infer the value of a novel item. Analogously, in the example of tennis, a player who has scored points from hitting to the backhand and also by performing drop shots may reasonably infer that a drop shot to the backhand will be successful. In psychology and neuroscience, the process by which decision variables in value-based and perceptual decision-making are retrieved and compared is described mechanistically by evidence integration or sequential sampling models (131). Within the probabilistic framework elaborated in Section 2, these models can be considered iterative approximate inference algorithms (132). There is both neural (36) and behavioral (133) evidence for their implementation in the brain. These sampling processes have been extended to tasks that require sequential actions over multiple states of control (134). A network of brain structures, primarily localized to prefrontal cortical areas, has been hypothesized to encode an internal model of the environment at the task level that relates relatively abstract representations of states, actions, and goals (135, 136). From a probabilistic perspective (see Section 2.5), this internal model can then be inverted via Bayesian inference to compute optimal actions (132). One heuristic strategy to accomplish this computation is to simply retrieve memories of past environment experiences based on state similarity as a proxy for internal forward modeling. In the human brain, this process appears to be mediated by the hippocampus (137).

Once a goal has been established, the abstract kinematic structure of a movement and the final state of the end effector (e.g., a hand) may be planned, a stage that may be referred to as action selection. One line of evidence for the existence of such motor representations comes from studies of the hand path priming effect (138). In these studies, participants are required to make obstacle-avoiding reaching movements. However, when cued to do so in the absence of obstacles, the participants appear to take unnecessarily long detours around the absent obstacles. Such suboptimal movements are inconsistent with OFC but are thought to be due to the efficient reuse of the abstract spatiotemporal form of the previously used movements. When such representations are available in the nervous system (as in the hand path priming experiments), it is possible that they may be reused in forward modeling simulations during motor planning. When combined with sampling strategies (120), the retrieval of abstract motor forms could provide a

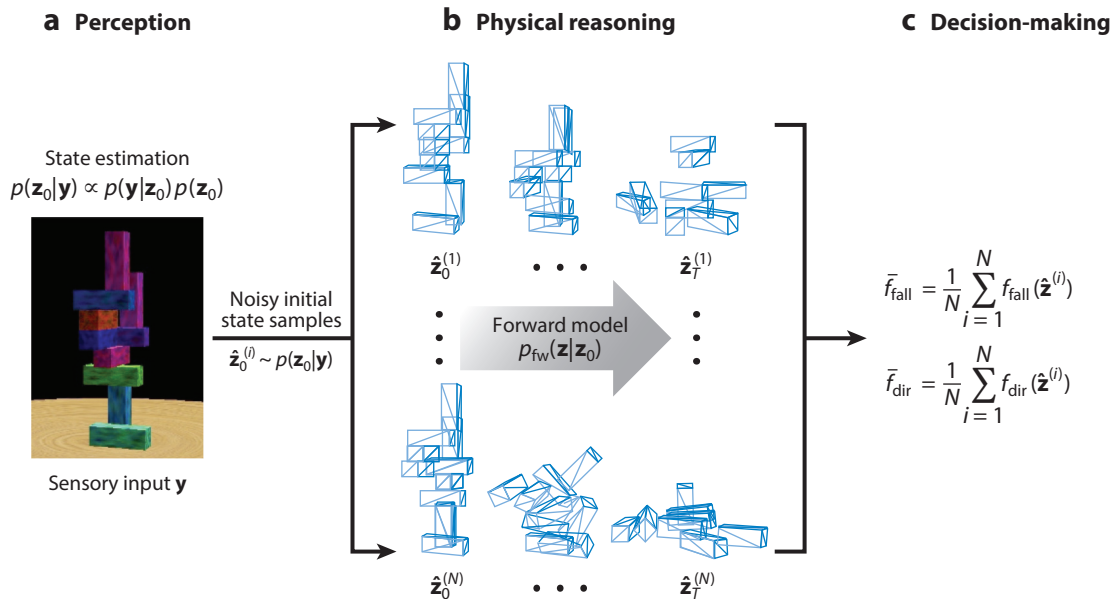


Figure 3

Physical reasoning. Participants must decide whether a complex scene of blocks will fall and, if so, the direction of the fall. A model of their performance combines perception, physical reasoning, and decision-making. (a) A Bayesian model of perception uses the sensory input \mathbf{y} to estimate a participant's belief $p(\mathbf{z}_0|\mathbf{y})$ regarding the initial environment state, including the position, geometry, and mass of the blocks. (b) Stochastic simulations based on samples from the posterior are performed using a noisy and approximate model of the physical properties of the world. The simulations use a forward model to sample multiple state trajectories (superscripts) over time (subscripts): $\hat{\mathbf{z}}^{(i)} = (\hat{\mathbf{z}}_0^{(i)}, \dots, \hat{\mathbf{z}}_T^{(i)})$. (c) The outputs of this intuitive physics engine can then be processed to make judgments, such as the proportion of the tower that will fall (\bar{f}_{fall}) and the direction of the fall (\bar{f}_{dir}). Experiments have indicated that humans are adept at making rapid judgments regarding the dynamics of such complex scenes, and these judgments are consistent with predictions generated using this model, which includes approximate Bayesian methods combined with internal forward models. Figure adapted from Reference 139.

computational foundation for the mental rehearsal of movement, which could be relatively efficient if applied at a high level of abstraction in the motor hierarchy.

In tasks involving complex object interactions, it may be particularly important to internally simulate the impact of different control strategies on the environment dynamics in order to avoid catastrophic outcomes, as envisaged by Craik. Humans are able to make accurate judgments regarding the dynamics of various visual scenes involving interacting objects under the influence of natural physical forces (Figure 3). This putative intuitive physics engine (139), which combines an internal model approximating natural physics with Monte Carlo sampling procedures, could be directly incorporated into motor planning within the probabilistic framework. Consider, for example, the problem of carrying a tray piled high with unstable objects. By combining internal simulations of the high-level features of potential movement plans with physical reasoning about the resulting object dynamics, one would be able to infer that it is more stable to grip the tray on each side rather than in the center and thus avoid having the objects fall to the floor. Thus, internal forward models can make a crucial contribution at the planning stage of control by simulating future state trajectories conditional on motor commands. It may be necessary to implement this processing at a relatively high level of the motor hierarchy in order to do so efficiently, given the complexity of the simulations. In the context of the tray example, the critical feature of the motor

movement in evaluating the stability of the objects is the manner in which the tray is gripped. Thus, simulating the large number of possible arm trajectories that move the hand into position is irrelevant to the critical success of the internal modeling. Identifying the essential abstract features of movement to input into a forward modeling process may be a crucial step in planning complex and novel movements.

4. CONCLUSIONS AND FUTURE DIRECTIONS

We have presented a formal integration of internal models with the rationality frameworks of Bayesian inference and OFC. In doing so, we have used the probabilistic formalism to review the various applications of internal models across a range of spatiotemporal scales in a unified manner. OFC provides a principled way in which a task can be associated with a cost, leading to an optimal control law that takes into account the dynamics of the body and the world as well as the noise processing involved in sensing and actuation. The theory is consistent with a large body of behavioral data. OFC relies on state estimation, which itself relies on internal models that are also of general use in a variety of processes and for which there is accumulating behavioral and neurophysiological evidence.

Major hurdles remain in understanding OFC in biology. First, it is unclear how a task specifies a cost function. While for a simple reaching movement it may be easy to use a combination of terminal error and energy, the links to cost are much less transparent in many real-world tasks. For example, when a person needs to remove keys from a pocket or tie shoelaces, it is difficult to calculate the cost involved. Indeed, recent work in robotics and machine learning has sought to learn abstract goal representations for use during planning and control rather than relying on a cost function (140). Second, although OFC can consider arbitrarily long (even infinite) horizons, people clearly plan their actions under finite-horizon assumptions by establishing a task-relevant temporal context. It is unclear how the brain temporally segments tasks and the extent to which each task is solved independently (126). Third, the representation of state is critical for OFC, but how state is constructed and used is largely unknown, though there are novel theories, with some empirical support, regarding how large state spaces could be modularized to make planning and policy encoding efficient (75). Fourth, even given a cost function or goal state specification, fully solving OFC in a reasonable time for a complex system such as a human body is intractable. The brain must use approximations to the optimal solution that are still unknown, although a variety of probabilistic machine learning methods (141) may provide inspiration for such investigations. Finally, the neural basis of both OFC and internal models is still in its infancy. However, the elaboration of OFC within the brain will take advantage of new techniques for dissecting neural circuitry (such as optogenetics), which have already delivered new insights into the neural basis of feedback-based sensorimotor control (142, 143).

Although many aspects of the computations underpinning processes such as sensory reafference cancellation and state estimation are well understood, the motor planning process remains poorly understood at a computational level. Some behavioral signatures and neural correlates of the computational principles by which plans are formed have been identified, but this has occurred primarily in tasks containing relatively small state and action spaces, such as sequential decision-making and spatial navigation. By contrast, the processes by which biological control solutions spanning large and continuous state spaces are constructed remain relatively unexplored. Future investigations may need to embed rich dynamical interactions between object dynamics and task goals in novel and complex movements. Such task manipulations may generate new insights into motor planning since the planning process may then depend on significant cognitive input, and so may reveal a more integrative form of planning across the sensorimotor hierarchy.

SUMMARY POINTS

1. Optimal feedback control and Bayesian estimation are rational principles for understanding human sensorimotor processing.
2. Internal models are necessary to facilitate dexterous control.
3. Forward models can assist in sensory filtering, state estimation, and planning.
4. Efficient internal models can mitigate the curse of dimensionality in sensorimotor control.

FUTURE ISSUES

1. Given a motor task, how are a state representation and cost function constructed?
2. What are the neural algorithms by which the solution to optimal feedback control is approximated?
3. How are internal models structured?
4. Are similar circuit mechanisms implemented across different prediction timescales?

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank the Wellcome Trust (Sir Henry Wellcome Fellowship to D.M. and Investigator Award to D.M.W.) and the Royal Society (Noreen Murray Professorship in Neurobiology to D.M.W.). We thank Emo Todorov for useful discussions.

LITERATURE CITED

1. Conant R, Ashby R. 1970. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* 1:89–97
2. Miall RC, Wolpert DM. 1996. Forward models for physiological motor control. *Neural Netw.* 9:1265–79
3. Kawato M. 1999. Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9:718–27
4. Craik KJW. 1943. *The Nature of Explanation*. Cambridge, UK: Cambridge Univ. Press
5. Popper K. 1957. *The Poverty of Historicism*. Boston: Beacon
6. Newell A. 1982. The knowledge level. *Artif. Intell.* 18:87–127
7. MacKay D. 2003. *Information Theory, Inference, and Learning Algorithms*. Cambridge, UK: Cambridge Univ. Press
8. Todorov E, Jordan MI. 2002. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5:1226–35
9. Todorov E. 2008. General duality between optimal control and estimation. In *2008 47th IEEE Conference on Decision and Control*, pp. 4286–92. New York: IEEE
10. Kappen HJ, Gómez V, Opper M. 2012. Optimal control as a graphical model inference problem. *Mach. Learn.* 87:159–82

11. Todorov E. 2004. Optimality principles in sensorimotor control. *Nat. Neurosci.* 7:907–15
12. Hinton GE, Dayan P, Frey BJ, Neal RM. 1995. The “wake-sleep” algorithm for unsupervised neural networks. *Science* 268:1158–61
13. Faisal A, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9:292–303
14. Kording K, Ku S, Wolpert D. 2004. Bayesian integration in force estimation. *J. Neurophysiol.* 92:3161–65
15. Kording K, Wolpert D. 2004. Bayesian integration in sensorimotor learning. *Nature* 427:244–47
16. Jazayeri M, Shadlen M. 2010. Temporal context calibrates interval timing. *Nat. Neurosci.* 13:1020–26
17. Tassinari H, Hudson T, Landy M. 2006. Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci.* 26:10154–63
18. Adams W, Graf E, Ernst M. 2004. Experience can change the ‘light-from-above’ prior. *Nat. Neurosci.* 7:1057–58
19. Kersten D, Mamassian P, Yuille A. 2004. Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55:271–304
20. Stone J, Kerrigan IS, Porrill J. 2009. Where is the light? Bayesian perceptual priors for lighting direction. *Proc. R. Soc. B* 276:1797–804
21. Weiss Y, Simoncelli EP, Adelson EH. 2002. Motion illusions as optimal percepts. *Nat. Neurosci.* 5:598–604
22. Tenenbaum JB, Griffiths TL. 2001. Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* 24:629–40
23. Griffiths TL, Tenenbaum JB. 2009. Theory-based causal induction. *Psychol. Rev.* 116:661–716
24. Tenenbaum JB, Griffiths TL, Kemp C. 2006. Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10:309–18
25. Houlshby NM, Huszár F, Ghassemi MM, Orbán G, Wolpert DM, Lengyel M. 2013. Cognitive tomography reveals complex, task-independent mental representations. *Curr. Biol.* 23:2169–75
26. Kim J, Greene M, Zlateski A, Lee K, Richardson M, et al. 2014. Space-time wiring specificity supports direction selectivity in the retina. *Nature* 509:331–36
27. Berry MJ II, Brivanlou IH, Jordan TA, Meister M. 1999. Anticipation of moving stimuli by the retina. *Nature* 398:334–38
28. Ma WJ, Beck JM, Latham PE, Pouget A. 2006. Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9:1432–38
29. Jazayeri M, Movshon JA. 2006. Optimal representation of sensory information by neural populations. *Nat. Neurosci.* 9:690–96
30. Beck JM, Latham PE, Pouget A. 2011. Marginalization in neural circuits with divisive normalization. *J. Neurosci.* 31:15310–19
31. Ganguli D, Simoncelli E. 2014. Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. *Neural Comput.* 26:2103–34
32. Rao RP. 2004. Bayesian computation in recurrent neural circuits. *Neural Comput.* 16:1–38
33. Deneve S. 2008. Bayesian spiking neurons I: inference. *Neural Comput.* 20:91–117
34. Aitchison L, Lengyel M. 2016. The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLOS Comput. Biol.* 12:e1005186
35. Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, et al. 2008. Probabilistic population codes for Bayesian decision making. *Neuron* 60:1142–52
36. Yang J, Scholz J, Latash M. 2007. The role of kinematic redundancy in adaptation of reaching. *Exp. Brain Res.* 176:54–69
37. Vogels T, Rajan K, Abbott L. 2005. Neural network dynamics. *Annu. Rev. Neurosci.* 28:357–76
38. Murphy B, Miller K. 2009. Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron* 61:635–48
39. Denève S, Machens CK. 2016. Efficient codes and balanced networks. *Nat. Neurosci.* 19:375–82
40. Fiser J, Berkes P, Orbán G, Lengyel M. 2010. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* 14:119–30
41. Berkes P, Orbán G, Lengyel M, Fiser J. 2011. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331:83–87

42. Todorov E, Li W. 2003. Optimal control methods suitable for biomechanical systems. In *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Vol. 2, pp. 1758–61. New York: IEEE
43. Todorov E. 2005. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput.* 17:1084–108
44. Flash T, Hogan N. 1985. The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* 5:1688–703
45. Uno Y, Kawato M, Suzuki R. 1989. Formation and control of optimal trajectory in human multijoint arm movement. *Biol. Cybernet.* 61:89–101
46. Bernstein NI. 1967. *The Coordination and Regulation of Movement*. Oxford, UK: Pergamon
47. Latash ML, Scholz JP, Schönér G. 2002. Motor control strategies revealed in the structure of motor variability. *Exerc. Sport Sci. Rev.* 30:26–31
48. Diedrichsen J. 2007. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr. Biol.* 17:1675–79
49. Liu D, Todorov E. 2007. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J. Neurosci.* 27:9354–68
50. Kutch JJ, Kuo AD, Bloch AM, Rymer WZ. 2008. Endpoint force fluctuations reveal flexible rather than synergistic patterns of muscle cooperation. *J. Neurophysiol.* 100:2455–71
51. Harris C, Wolpert D. 1998. Signal-dependent noise determines motor planning. *Nature* 394:780–84
52. Scott SH. 2012. The computational and neural basis of voluntary motor control and planning. *Trends Cogn. Sci.* 16:541–49
53. Diedrichsen J. 2007. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr. Biol.* 17:1675–79
54. Hatze H, Buys JD. 1977. Energy-optimal controls in the mammalian neuromuscular system. *Biol. Cybernet.* 27:9–20
55. Anderson FC, Pandy MG. 2001. Dynamic optimization of human walking. *J. Biomech. Eng.* 123:381–90
56. Srinivasan M, Ruina A. 2006. Computer optimization of a minimal biped model discovers walking and running. *Nature* 439:72–75
57. Shadmehr R, Huang HJ, Ahmed AA. 2016. A representation of effort in decision-making and motor control. *Curr. Biol.* 26:1929–34
58. Wolpert D, Landy M. 2012. Motor control is decision-making. *Curr. Opin. Neurobiol.* 22:996–1003
59. Trommershauser J, Maloney L, Landy M. 2008. Decision making, movement planning and statistical decision theory. *Trends Cogn. Sci.* 12:291–97
60. Diamond J, Wolpert D, Flanagan J. 2017. Rapid target foraging with reach or gaze: The hand looks further ahead than the eye. *PLOS Comput. Biol.* 13:e1005504
61. Jones K, Hamilton A, Wolpert D. 2002. Sources of signal-dependent noise during isometric force production. *J. Neurophysiol.* 88:1533–44
62. Hamilton A, Jones K, Wolpert D. 2004. The scaling of motor noise with muscle strength and motor unit number in humans. *Exp. Brain Res.* 157:417–30
63. Franklin DW, Burdet E, Tee KP, Osu R, Chew CM, et al. 2008. CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.* 28:11165–73
64. Valero-Cuevas FJ, Venkadesan M, Todorov E. 2009. Structured variability of muscle activations supports the minimal intervention principle of motor control. *J. Neurophysiol.* 102:59–68
65. Nagengast AJ, Braun DA, Wolpert DM. 2010. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. *PLOS Comput. Biol.* 6:e1000857
66. Dimitriou M, Wolpert DM, Franklin DW. 2013. The temporal evolution of feedback gains rapidly update to task demands. *J. Neurosci.* 33:10898–909
67. Yeo SH, Franklin DW, Wolpert DM. 2016. When optimal feedback control is not enough: Feedforward strategies are required for optimal control with active sensing. *PLOS Comput. Biol.* 12:e1005190
68. Scott SH. 2016. A functional taxonomy of bottom-up sensory feedback processing for motor actions. *Trends Neurosci.* 39:512–26
69. Rothwell JC, Traub MM, Marsden CD. 1980. Influence of voluntary intent on the human long-latency stretch reflex. *Nature* 286:496–98

70. Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH. 2012. Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* 478:387–90
71. Shen L, Alexander GE. 1997. Neural correlates of a spatial sensory-to-motor transformation in primary motor cortex. *J. Neurophysiol.* 77:1171–94
72. Scott S, Kalaska J. 1997. Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *J. Neurophysiol.* 77:826–52
73. Scott SH. 2004. Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* 5:532–46
74. Pruszynski J, Scott S. 2012. Optimal feedback control and the long-latency stretch response. *Exp. Brain Res.* 218:341–59
75. McNamee D, Wolpert DM, Lengyel M. 2016. Efficient state-space modularization for planning: theory, behavioral and neural signatures. In *Advances in Neural Information Processing Systems 29*, ed. DD Lee, M Sugiyama, UV Luxburg, I Guyon, R Garnett, pp. 4511–19. Red Hook, NY: Curran
76. Kalman RE. 1960. A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82:35
77. Mitter S, Newton N. 2003. A variational approach to nonlinear estimation. *SIAM J. Control Optim.* 42:1813–33
78. Todorov E. 2009. Compositionality of optimal control laws. In *Advances in Neural Information Processing Systems 22*, ed. Y Bengio, D Schuurmans, JD Lafferty, CKI Williams, A Culotta, pp. 1856–64. Red Hook, NY: Curran
79. Todorov E. 2007. Linearly-solvable Markov decision problems. In *Advances in Neural Information Processing Systems 19*, ed. B Schölkopf, JC Platt, T Hoffman, pp. 1369–76. Red Hook, NY: Curran
80. Rawlik K, Toussaint M, Vijayakumar S. 2013. On stochastic optimal control and reinforcement learning by approximate inference (extended abstract). In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pp. 3052–56. Palo Alto, CA: AAAI Press
81. Rao RP, Ballard DH. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2:79–87
82. van Beers RJ. 2009. Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63:406–17
83. Braun DA, Nagengast AJ, Wolpert DM. 2011. Risk-sensitivity in sensorimotor control. *Front. Hum. Neurosci.* 5:1
84. Kawato M, Furukawa K, Suzuki R. 1987. A hierarchical neural-network model for control and learning of voluntary movement. *Biol. Cybernet.* 57:169–85
85. Jordan MI. 1996. Computational aspects of motor control and motor learning. In *Handbook of Perception and Action*, Vol. 2: *Motor Skills*, ed. H Heuer, SW Keele, pp. 71–120. San Diego, CA: Academic
86. von Holst E, Mittelstaedt H. 1950. Das Reafferenzprinzip: Wechselwirkungen zwischen Zentralnervensystem und Peripherie. *Naturwissenschaften* 37:464–76
87. Crapse TB, Sommer MA. 2008. Corollary discharge across the animal kingdom. *Nat. Rev. Neurosci.* 9:587–600
88. Haarmeier T, Bunjes F, Lindner A, Berret E, Thier P. 2001. Optimizing visual motion perception during eye movements. *Neuron* 32:527–35
89. Sommer MA, Wurtz RH. 2006. Influence of the thalamus on spatial visual processing in frontal cortex. *Nature* 444:374–77
90. Bell C. 2001. Memory-based expectations in electrosensory systems. *Curr. Opin. Neurobiol.* 11:481–87
91. Sawtell NB. 2017. Neural mechanisms for predicting the sensory consequences of behavior: insights from electrosensory systems. *Annu. Rev. Physiol.* 79:381–99
92. Rothwell JC, Traub MM, Day BL, Obeso JA, Thomas PK, Marsden CD. 1982. Manual motor performance in a deafferented man. *Brain* 105:515–42
93. Gordon J, Ghilardi MF, Ghez C. 1995. Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy. *J. Neurophysiol.* 73:361–72
94. Wolpert D, Ghahramani Z, Jordan M. 1995. An internal model for sensorimotor integration. *Science* 269:1880–82
95. Diedrichsen J, Verstynen T, Hon A, Lehman SL, Ivry RB. 2003. Anticipatory adjustments in the unloading task: Is an efference copy necessary for learning? *Exp. Brain Res.* 148:272–76

96. Flanagan JR, Wing AM. 1993. Modulation of grip force with load force during point-to-point arm movements. *Exp. Brain Res.* 95:131–43
97. Kawato M, Kuroda T, Imamizu H, Nakano E, Miyauchi S, Yoshioka T. 2003. Internal forward models in the cerebellum: fMRI study on grip force and load force coupling. *Prog. Brain Res.* 142:171–88
98. Johansson RS, Riso R, Häger C, Bäckström L. 1992. Somatosensory control of precision grip during unpredictable pulling loads. I. Changes in load force amplitude. *Exp. Brain Res.* 89:181–91
99. Crevecoeur F, Scott S. 2013. Priors engaged in long-latency responses to mechanical perturbations suggest a rapid update in state estimation. *PLOS Comput. Biol.* 9:e1003177
100. Crevecoeur F, Munoz DP, Scott SH. 2016. Dynamic multisensory integration: Somatosensory speed trumps visual accuracy during feedback control. *J. Neurosci.* 36:8598–611
101. Nasseroleslami B, Hasson C, Sternad D. 2014. Rhythmic manipulation of objects with complex dynamics: predictability over chaos. *PLOS Comput. Biol.* 10:e1003900
102. Diedrichsen J, Bastian A. 2014. Cerebellar function. In *The Cognitive Neurosciences*, ed. MS Gazzaniga, GR Mangun, pp. 451–60. Cambridge, MA: MIT Press. 5th ed.
103. Therrien AS, Bastian AJ. 2015. Cerebellar damage impairs internal predictions for sensory and motor function. *Curr. Opin. Neurobiol.* 33:127–33
104. Ebner TJ, Pasalar S. 2008. Cerebellum predicts the future motor state. *Cerebellum* 7:583–88
105. Hewitt AL, Popa LS, Pasalar S, Hendrix CM, Ebner TJ. 2011. Representation of limb kinematics in Purkinje cell simple spike discharge is conserved across multiple tasks. *J. Neurophysiol.* 106:2232–47
106. Brooks JX, Carriot J, Cullen KE. 2015. Learning to expect the unexpected: rapid updating in primate cerebellum during voluntary self-motion. *Nat. Neurosci.* 18:1310–17
107. Müller F, Dichgans J. 1994. Dyscoordination of pinch and lift forces during grasp in patients with cerebellar lesions. *Exp. Brain Res.* 101:485–92
108. Miall RC, Christensen LOD, Cain O, Stanley J. 2007. Disruption of state estimation in the human lateral cerebellum. *PLOS Biol.* 5:2733–44
109. Bhanpuri NH, Okamura AM, Bastian AJ. 2014. Predicting and correcting ataxia using a model of cerebellar function. *Brain* 137:1931–44
110. Mulliken GH, Musallam S, Andersen RA. 2008. Forward estimation of movement state in posterior parietal cortex. *PNAS* 105:8170–7
111. Buneo CA, Andersen RA. 2006. The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia* 44:2594–606
112. Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST. 1999. Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat. Neurosci.* 2:563–67
113. Desmurget M, Reilly KT, Richard N, Szathmari A, Mottolese C, Sirigu A. 2009. Movement intention after parietal cortex stimulation in humans. *Science* 324:811–13
114. Rushworth MF, Nixon PD, Passingham RE. 1997. Parietal cortex and movement. I. Movement selection and reaching. *Exp. Brain Res.* 117:292–310
115. Sutton R, Barto A. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press
116. Toussaint M, Storkey A. 2006. Probabilistic inference for solving discrete and continuous state Markov decision processes. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 945–52. New York: ACM
117. Dolan RJ, Dayan P. 2013. Goals and habits in the brain. *Neuron* 80:312–25
118. Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science* 275:1593–99
119. O’Doherty JP, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. 2004. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304:452–54
120. Kappen HJ, Ruiz HC. 2016. Adaptive importance sampling for control and inference. *J. Stat. Phys.* 162:1244–66
121. Theodorou E, Buchli J, Schaal S. 2010. A generalized path integral control approach to reinforcement learning. *J. Mach. Learn. Res.* 11:3137–81
122. Silver D, Veness J. 2010. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems 23*, ed. JD Lafferty, CKI Williams, J Shawe-Taylor, RS Zemel, A Culotta, pp. 2164–72. Red Hook, NY: Curran

123. Bellman R. 1957. *Dynamic Programming*. Princeton, NJ: Princeton Univ. Press
124. D'Avella A, Saltiel P, Bizzi E. 2003. Combinations of muscle synergies in the construction of a natural motor behavior. *Nat. Neurosci.* 6:300–8
125. Huys QJ, Lally N, Faulkner P, Eshel N, Seifritz E, et al. 2015. Interplay of approximate planning strategies. *PNAS* 112:3098–103
126. Botvinick M, Niv Y, Barto A. 2009. Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition* 113:262–80
127. Newell A, Rosenbloom PS. 1981. Mechanisms of skill acquisition and the law of practice. In *Cognitive Skills and Their Acquisition*, ed. JR Anderson, pp. 1–56. Cambridge, MA: MIT Press
128. McNamee D, Rangel A, O'Doherty J. 2013. Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex. *Nat. Neurosci.* 16:479–85
129. Padoa-Schioppa C, Assad JA. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–26
130. Barron HC, Dolan RJ, Behrens TEJ. 2013. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat. Neurosci.* 16:1492–98
131. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113:700
132. Solway A, Botvinick MM. 2012. Goal-directed decision making as probabilistic inference. *Psychol. Rev.* 119:120–54
133. Resulaj A, Kiani R, Wolpert DM, Shadlen MN. 2009. Changes of mind in decision-making. *Nature* 461:263–66
134. Solway A, Botvinick MM. 2015. Evidence integration in model-based tree search. *PNAS* 112:11708–13
135. Daw ND, Niv Y, Dayan P. 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8:1704–11
136. Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. 2014. Orbitofrontal cortex as a cognitive map of task space. *Neuron* 81:267–79
137. Bornstein AM, Norman KA. 2017. Reinstated episodic context guides sampling-based decisions for reward. *Nat. Neurosci.* 20:997–1003
138. Jax S, Rosenbaum D. 2007. Hand path priming in manual obstacle avoidance: evidence that the dorsal stream does not only control visually guided actions in real time. *J. Exp. Psychol. Hum. Percept. Perform.* 33:425–41
139. Battaglia PW, Hamrick JB, Tenenbaum JB. 2013. Simulation as an engine of physical scene understanding. *PNAS* 110:18327–32
140. Srinivas A, Jabri A, Abbeel P, Levine S, Finn C. 2018. Universal planning networks: learning generalizable representations for visuomotor control. In *Proceedings of the 35th International Conference on Machine Learning*, ed. J Dy, A Krause, pp. 4739–48. Proc. Mach. Learn. Res. Vol. 80. N.p.: PMLR
141. Ghahramani Z. 2015. Probabilistic machine learning and artificial intelligence. *Nature* 521:452–59
142. Esposito M, Capelli P, Arber S. 2014. Brainstem nucleus MdV mediates skilled forelimb motor tasks. *Nature* 508:351–56
143. Azim E, Jiang J, Alstermark B, Jessell TM. 2014. Skilled reaching relies on a V2a propriospinal internal copy circuit. *Nature* 508:357–63



Contents

A Century of Robotic Hands <i>C. Piazza, G. Grioli, M.G. Catalano, and A. Bicchi</i>	1
Escaping Oz: Autonomy in Socially Assistive Robotics <i>Caitlyn Clabaugh and Maja Matarić</i>	33
Modular Reconfigurable Robotics <i>Jungwon Seo, Jamie Paik, and Mark Yim</i>	63
Control Across Scales by Positive and Negative Feedback <i>R. Sepulchre, G. Drion, and A. Franci</i>	89
Formal Methods for Control Synthesis: An Optimization Perspective <i>Calin Belta and Sadra Sadraddini</i>	115
Discrete Event Systems: Modeling, Observation, and Control <i>Stéphane Lafortune</i>	141
From Visual Understanding to Complex Object Manipulation <i>Judith Bütepage, Silvia Cruciani, Mia Kokic, Michael Welle, and Danica Kragic</i>	161
Robotic Micromanipulation: Fundamentals and Applications <i>Zhuoran Zhang, Xian Wang, Jun Liu, Changsheng Dai, and Yu Sun</i>	181
Microrobotics and Microorganisms: Biohybrid Autonomous Cellular Robots <i>Yunus Alapan, Oncay Yasa, Berk Yigit, I. Ceren Yasa, Pelin Erkoc, and Metin Sitti</i>	205
Toward Autonomy in Sub-Gram Terrestrial Robots <i>Ryan St. Pierre and Sarah Bergbreiter</i>	231
A Tour of Reinforcement Learning: The View from Continuous Control <i>Benjamin Recht</i>	253
System Identification: A Machine Learning Perspective <i>A. Chiuso and G. Pillonetto</i>	281

A Perspective on Incentive Design: Challenges and Opportunities <i>Lillian J. Ratliff, Roy Dong, Shreyas Sekar, and Tanner Fiez</i>	305
Internal Models in Biological Control <i>Daniel McNamee and Daniel M. Wolpert</i>	339
Agricultural Robotics <i>Stavros G. Vougioukas</i>	365
Modeling and Estimation for Advanced Battery Management <i>Xinfan Lin, Youngki Kim, Shankar Mohan, Jason B. Siegel, and Anna G. Stefanopoulou</i>	393
Cyber-Physical Manufacturing Systems <i>Dawn M. Tilbury</i>	427
The Engineering of Climate Engineering <i>Douglas G. MacMartin and Ben Kravitz</i>	445

Errata

An online log of corrections to *Annual Review of Control, Robotics, and Autonomous Systems* articles may be found at <http://www.annualreviews.org/errata/control>