

The Synergy Between Neuroscience and Control Theory: The Nervous System as Inspiration for Hard Control Challenges

Manu S. Madhav¹ and Noah J. Cowan²

¹Mind/Brain Institute, Johns Hopkins University, Baltimore, Maryland 21218, USA;
email: manusmad@jhu.edu

²Department of Mechanical Engineering, Johns Hopkins University, Baltimore,
Maryland 21218, USA

Annu. Rev. Control Robot. Auton. Syst. 2020.
3:243–67

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

<https://doi.org/10.1146/annurev-control-060117-104856>

Copyright © 2020 by Annual Reviews.
All rights reserved

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

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

Keywords

biological inspiration, system identification, feedback, closed-loop neuroscience

Abstract

Here, we review the role of control theory in modeling neural control systems through a top-down analysis approach. Specifically, we examine the role of the brain and central nervous system as the controller in the organism, connected to but isolated from the rest of the animal through insulated interfaces. Though biological and engineering control systems operate on similar principles, they differ in several critical features, which makes drawing inspiration from biology for engineering controllers challenging but worthwhile. We also outline a procedure that the control theorist can use to draw inspiration from the biological controller: starting from the intact, behaving animal; designing experiments to deconstruct and model hierarchies of feedback; modifying feedback topologies; perturbing inputs and plant dynamics; using the resultant outputs to perform system identification; and tuning and validating the resultant control-theoretic model using specially engineered robophysical models.

1. INTRODUCTION

Biologists have long recognized the remarkable similarity between certain biological processes and synthetic feedback regulation. In fact, Wallace, along with Darwin, directly referenced the centrifugal governor as being like the evolutionary principle, in that it “checks and corrects any irregularities almost before they become evident” (1, p. 62). This analogy is intriguing but imperfect since there is no “evolver” that senses errors and makes said corrections. Closer to its engineering control system analog is the idea of biological homeostasis developed by the French physiologist Claude Bernard. The internal environment of the body stabilizes itself against external influences, creating what he called *la milieu interieur* (2, 3). Thus, it has long been recognized that feedback regulation exists within an individual organism’s body (if not, precisely speaking, at the macroscale through evolution and natural selection). In broad terms, feedback regulation refers to the idea that, in order for a system to achieve a desired output, the actual output is compared with the desired output, fed through some computation, and redirected to a system input, thus closing a feedback loop. This comparison (the error signal) allows a controller to generate the appropriate input to the rest of the system, termed the plant, so as to drive the error signal as close to zero as possible. Feedback is responsible for the regulation of fluid pressure within the eye (4), the viscoelastic reflexive properties of muscles (5), the thermal regulation of blood flow (6), and so on.

Feedback regulation is as ubiquitous in engineering as it has always been in biology, and control theory provides the basis for a duality between the two. Control theory enables the quantification of the dynamics of a biological system within a quantitative engineering framework (7) and provides a means to develop appropriate abstractions that enable transfer of knowledge between the fields. It also provides a common language to describe and analyze biological feedback as well as design experiments to manipulate feedback topologies and understand the underlying mechanisms (8) through the process of system identification.

In this review, we take a comparative approach, drawing on examples from across animal taxa—from humans to cockroaches, from fish to fruit flies. We do not intend this review to provide a comprehensive investigation of any particular animal or behavior. Rather, the organisms and behaviors selected here highlight a finding, methodology, or perspective that we feel is of broad interest to the control engineer or neuroscientist. In Section 2, we make a case that the brain and nervous system merit being modeled separately from the body [though no more importantly (9)] as the control system of the animal. Then, in Section 3, we describe some (but not all) of the key differences between biological control systems and their engineering analogs, in order to highlight opportunities for learning from biology. We also provide some caution that not all features of neural control are worth copying (although they must nevertheless be addressed). Finally, in Section 4, we provide a broad road map for how to investigate neural control using engineering tools, as a first step toward translating neural control algorithms to engineering and robotics.

2. MAKING THE CASE FOR THE BRAIN AS CONTROLLER

The parameterization and translation of biological systems into the language of control is beneficial not only to increase our understanding of biology but also to the field of control theory itself. However, designing a synthetic system from a corresponding biological one is not easy.

In the fifteenth century, Leonardo da Vinci designed flying machines based on studies of the anatomy of birds. However, these innovative and beautiful designs did not result in humans flapping and soaring above Renaissance Italy. In fact, it took another 400 years for a heavier-than-air vehicle to take flight. The innovation that made this possible was not merely in materials and aerodynamics—it was in control. The Wright brothers purposefully designed an inherently unstable aircraft (10) in closed loop with an intelligent, adaptive controller—the pilot—making it

possible to keep the aircraft aloft where other, less controllable designs had failed. They traded off stability for controllability and maneuverability, finding the sweet spot where the aircraft was not too unstable while retaining sufficient control affordance.

From an analytical standpoint, the nervous system and musculoskeletal system are “just” dynamical systems that operate in closed loop with one another. The nervous system was not designed to control a preexisting plant, as can be the case in engineering applications; rather, the nervous system and the body coevolved. Furthermore, the brain is certainly not the only feedback system in an animal.

Despite this apparent entanglement of mechanics, sensing, and neural computation, we argue that the nervous system stands apart in both form and function from the rest of the body, making it analogous to the robotic controller. This viewpoint derives, in part, from the directionality of the flow of information into and out of the neural substrate (see **Figure 1**). The transduction of sensory inputs into the language of the brain—i.e., electrical impulses in neuronal tissue—is, with few exceptions, high impedance in the sense that sensory transduction has a minimal retroactive (11) effect on the environmental variables it is sensing. Indeed, sensory receptors transform signals (displacement, forces, light fluxes, sound waves, etc.) into neural activity, and once in the currency of neural spikes, that sensory information does not reemerge into the physical world until motor units discharge into muscles. Likewise, this neuromuscular interface is low output: The descending motor commands are not affected by the dynamics of the muscles or the forces they apply to the skeletal system (except, of course, via neural feedback from proprioceptors and other sensors). In this way, the nervous system gathers information and influences the environment only through these one-way sensory and motoric portals into and out of the nervous system.

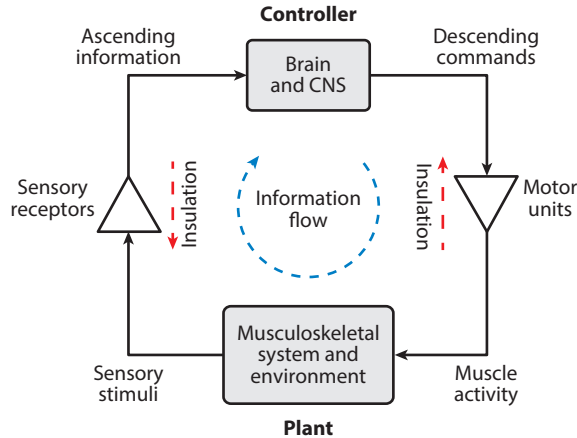


Figure 1

Schematic of the flow of information between an organism and its environment. The plant (a combination of the organism’s musculoskeletal system and the environment) receives descending commands from the brain and central nervous system (CNS) through motor units. This interaction is unidirectional since environmental perturbations do not return to the brain through motor units. Instead, measurements enter the brain via transduction by sensory receptors. This process is also unidirectional since the brain cannot affect sensory stimuli except through motor action. The brain and CNS are thus highly insulated in the direction of both sensory afferents and motor commands. This insulation or low retroactivity (10, 11) is indicated by the symbol of an amplifier (*triangle*). The resultant unidirectional flow of information isolates the CNS from the plant it controls, much as in engineering control systems.

The nervous system stands apart in a modular fashion from the environment, isolated through unidirectional inputs and outputs. The interfaces to the nervous system thus enjoy the properties of an insulation device, similar to an operational amplifier in an electrical circuit, where input and output influences are masked by high-gain feedback loops (11). By contrast, cellular mechanisms for signaling may not always enjoy such modularity, because high-impedance transduction in cellular and molecular circuits may be energetically costly (12). For more than 100 years, developments in control theory have intrinsically relied on these one-way couplings between controller and plant, rendering the tools of control systems analysis and design remarkably well suited to understanding neural control.

This modular relationship between the nervous system and the musculoskeletal system it controls, coupled with the extraordinary computational flexibility of the nervous system, establishes the brain and nervous system as the controller, not as a matter of convenience of parlance, but as a deep architectural analogy afforded by said modularity. While this distinction is of course fuzzier in biological systems than in their engineered counterparts, it seems clear that the role of the brain is to transform sensory input into motor output. This does not make the biomechanical plant any less important than the computations undertaken by the neural control system (13), but rather illuminates the appropriateness of drawing a distinction between various levels of feedback outside the realm of neural spikes—e.g., dynamics within the body and body–environment interactions—and sensorimotor feedback control. The latter feedback enters the nervous system through sensors and discharges from the nervous system through muscles. While a well-designed mechanical system can greatly simplify control (14, 15), the clear dichotomy of the brain and body as controller and plant, respectively, positions feedback control theory as the ideal computational and analytical framework to examine the brain and nervous system, decode the computations that result in behavior, and examine these behaviors in the context of feedback from the environment in which they are performed.

In general, neural controllers seem to be more robust and to operate over more diverse environmental conditions than their engineering counterparts. Sensors and actuators in engineering are incontrovertibly more accurate and precise than their biological analogs, and yet organisms manage to process and infer knowledge from their uncertain and noisy sensory information and manipulate the environment through noisy actuation with greater ease than the best robots. There are many differences between neuronal computations and their engineering analogs—the challenge lies in identifying which biological features confer an advantage.

3. BUG OR FEATURE? CHARACTERISTICS OF NEURAL CONTROLLERS

Even complex engineering control systems are designed in a modular fashion, with individual components and subassemblies designed to have a well-defined set of characteristics in order to simplify integration. This typically means that the control engineer focuses on a relatively simple feedback control topology, with at most a few nested feedback loops. Biological feedback, by contrast, is a multipartite hierarchy of feedback loops (16) interwoven with complex mechanics at multiple levels (17).

In addition, nature often violates the rather religious partition of sensing, control, and actuation observed in human-made systems. Rather than having separate components designed for a particular purpose and then integrated, an organism evolves as a single unit, and thus its components are free from the constraint of having one specific function. Sensory organs often perform computation (18), and sensory feedback is integrated into motor systems (19).

How do we translate this multifunctional architecture to the generally modular, unifunctional design space of robotics and control systems? As a first step, it is essential to address which features are bugs of the biological wetware and which are features worthy of being copied. It is possible that differences between biological and artificial computation are simply a side effect of biochemistry and biophysics—the necessities of using biological wetware. However, there may be control principles hidden within how biological systems relax engineering assumptions that are valuable to engineers as we seek to discover generalizable paradigms in robot and control design.

In the following sections, we highlight some ways in which the brain and its interface with the organism's body and the environment set up challenges that are rarely found in contemporary engineered systems. However, decoding nature's way of solving these challenges is going to be critical in moving toward more robust and general-purpose robotic and control systems. Our goal in this section is not to provide an exhaustive list but rather to examine a few concrete but important examples of features of neural control systems that are generally different from those in engineering.

3.1. Active Sensing

For an engineering system, sensing and actuation are largely decoupled. On a robot, sensors are typically fixed in location, range, and sensing volume relative to the robot. However, this is radically different in the biological world. Sensing and actuation are tightly coupled, with actuation often employed to enhance sensing in addition to providing locomotion. Envision entering and trying to find one's way around a dark room as compared with a well-lit room. We employ very different sensing strategies in these two cases: In the dark, we tend to choose a path closer to boundaries and use our hands to actively explore the boundaries.

Here, we use the term active sensing to refer to the process of using motor actions in order to improve sensory acquisition (20). Active sensing is a fundamental and widespread nonlinearity in biological sensorimotor control and is employed across the animal kingdom. Active sensing has been reported in many biological systems. Rats change the timing and velocity of whisking in order to increase the duration and number of whisker contacts during exploration (21). Flies alter their visual processing during flight such that their motion-sensitive neurons are better tuned to the higher velocities associated with flight (22). Humans employ different manipulation strategies to measure different physical features of the same object (23), employing what has been known in the psychological literature as active touch (24).

The weakly electric fish *Eigenmannia vivescens* inhabits freshwater environments in Central and South America and can perceive the world through electricity; an electric field generated in the fish's tail is sensed back through electroreceptors distributed on its skin, forming an image much like retinal photoreceptors do for light. These fish also use their mechanosensory lateral line (flow and pressure sensing) and vision to sense their surroundings. In addition, they readily perform an untrained refuge-tracking behavior in the laboratory, which is thought to facilitate hiding within aquatic foliage in the wild in order to protect themselves from predation (25). The fish are able to stabilize the image—visual, electrosensory, and mechanosensory—of the refuge relative to their bodies such that the error between the refuge and body position can be maintained close to zero (26).

In complete darkness, these fish supplement their regular tracking behavior with higher-frequency oscillations (27, 28). They also oscillate their bodies next to prey prior to the actual prey capture maneuver (29). The spatiotemporal flow of the electrosensory image in electric fish depends heavily on their movement (30). In *E. vivescens*, Stamper et al. (27) discovered that active sensing movements increase with increased conductivity (low electrosensory salience). Following

up on this work, Biswas et al. (31) developed an experimentally closed-loop apparatus (see Section 4.6) in which refuge movement was controlled in real time in response to fish movement, thus shaping the sensory feedback available to the fish. They found that, on top of maintaining position with respect to the refuge, these fish use an additional layer of nonlinear feedback control to maintain a consistent level of sensory feedback caused by the active sensing movements. The control algorithm remains unknown, but Kunapareddy & Cowan (32) developed a nonlinear controller, inspired by the electric fish, that uses active sensing movements to enhance nonlinear observability. Research on weakly electric fish continues to expand our understanding of how animals can maintain stable sensing and perform tasks by utilizing variable reliance on sensory inputs and performing active sensing maneuvers when necessary.

The tight coupling between sensing and actuation in active sensing goes against a long-standing principle of design in feedback control: the separation principle. Under the separation principle, designing a state estimator for a system can be decoupled from designing a controller, and for a linear time-invariant system, the eigenvalues of the closed-loop dynamics are just a union of the eigenvalues of the estimator and controller. Indeed, it is well known that linear quadratic Gaussian control combines a Kalman filter and a linear quadratic regulator, both designed independently. Thus, while in some cases biological closed loops appear to be at least approximately optimal (33, 34), there is a fundamental difference in the coupling of estimation and control between synthetic and biological systems. Indeed, active sensing even occurs during the single-degree-of-freedom task of refuge tracking in *E. virescens* (27, 31, 35) despite its simple (approximately linear time-invariant) plant (14) (see Section 4.4).

3.2. Neural Representations of the Plant

The brain often contains within itself representations of mechanisms present in the rest of the body or of interactions between the body and the environment (the plant, in control-theoretic parlance). These plant models are beneficial in several ways. They allow the organism to simulate and evaluate the consequences of an action before actually performing the action (36). In addition, these models can help researchers distinguish between self-generated (expected) motions and environmental (unexpected) perturbations (37) and study the estimation and regulation of internal states, comparison and learning of control policies, overcoming of feedback delays, prediction of future states (19), and planning of long-term actions (38).

Within systems neuroscience, the motor learning community has long embraced the idea of forward models in the brain. Specifically, the cerebellum is the region of the brain thought to contain internal models of the motor apparatus (39) and of perceptual processes (40). Deficits in cerebellar function may be captured—at least to some extent—as if there is a mismatch in model parameters (41). There is evidence to indicate that these internal models perform state estimation in a probabilistic framework, combining prior knowledge with current information from the environment acquired through sensory input (42). Bayesian integration can aid sensorimotor learning (43), and the evidence suggests that this is the underlying process that allows control strategies to be tuned through optimal feedback control (33). (For a review of internal models in the probabilistic framework, see Reference 36.)

In control theory, plant models are closely related to state estimation and adaptive control. State estimators use the measured plant output to fully or partially reconstruct its internal states. Adaptive control utilizes estimators in order to estimate internal states as well as dynamical parameters of the plant itself. In particular, indirect adaptive controllers estimate an explicit representation of the plant in the process of implementing the controller (44). Model-reference adaptive controllers regulate both the internal model and the plant using two controllers that learn from output error

such that, under persistent excitation, the model parameters converge to those of the plant (45). In fact, it can be argued that any good regulator of a system must theoretically include a model of that system (46). (For a review of the parallels between adaptive controllers and internal models in the nervous system, see Reference 47.)

Thus, internal models are not just a good idea but necessary for proper functioning of a controller. However, it is unclear whether these models need to be explicitly instantiated in a mechanism distinct from the rest of the controller, as in indirect adaptive controllers or in the mammalian cerebellum. The resolution at which these internal models need to be maintained is also unknown—i.e., whether they need to provide only high-level templates (48) for guiding control strategies or need to actually model the specifics of the mechanism. These factors determine the complexity of the controller and the amount of resources that a controller devotes to simulating the plant dynamics. As a complex controller controlling a complex plant, the nervous system could have implemented sensorimotor control strategies in several different ways; however, internal models are clearly a core aspect of the neural control strategy. Thus, the nervous system provides an opportunity to explore questions about the role of plant models in control.

3.3. Neural Representations of the Environment and Exogenous Signals

In contrast to plant models, environment models in the brain refer to representations of environmental variables or states of the organism in relation to the environment. These models allow the formation of short- and long-term memory as well as planning and control strategies that take into account reference frames external to the organism.

Learning the features of exogenous sensory input—i.e., signals generated external to the animal—can allow a controller to use its historical states to predict future states. Accurate and reliable prediction can dramatically improve the performance of a controller and has been used in controller design techniques such as model predictive control and generalized predictive control (47, 49). In humans, prediction is extensively used in motor planning (19), for example, in adjusting grip strength in response to anticipated changes in load (50) and in anticipating reversals in ocular pursuit (51).

In addition to anticipating and performing corrective motor movements, animals also use prediction of exogenous signals to effect higher-level control strategies. Above, we described the untrained behavior of the weakly electric fish *E. virescens* to track a refuge using the combined information from vision, electrosensation, and mechanosensation. Although it is a simple, single-degree-of-freedom, approximately linear time-invariant behavior (14), refuge tracking exhibits a marked nonlinearity, showing improved tracking performance for (more predictable) single-sine stimuli as opposed to components of sum-of-sines stimuli (52), thus failing to satisfy the superposition property of linear systems in a critical way. Such data are evidence for the internal model principle (53–55).

One of the best examples of environmental models in the nervous system is the cognitive map in the mammalian brain. The existence of cognitive maps was first postulated by Tolman (56). Subsequently, evidence for them was found through the activity of neurons of the rodent hippocampal formation. Most notably, place cells are neurons in the hippocampus (57) that fire action potentials when a rat is in a particular location in its environment. The firing rates of head-direction cells of the thalamus (58) and other associated brain regions display precise tuning to the orientation of the animal's head with respect to its environment. Grid cells of the medial entorhinal cortex display firing in a characteristic hexagonal lattice that tiles the environment (59). Border cells of the entorhinal cortex fire at environmental boundaries (60, 61).

These categories of neurons fire with respect to the animal's pose in allocentric space (locations or orientations with respect to the environment or an absolute reference frame) as opposed to egocentric space (locations or orientations with respect to the animal or a relative reference frame). However, there are neurons that code for egocentric variables (62) and vector relationships to landmarks (63), boundaries (64), and objects (65).

The neurons of the cognitive map do not directly respond to specific sensory input. For instance, place cells still fire in the dark (66), but their spatial specificity decreases due to accumulated error (67). These findings and other extensive experimental evidence indicate that the cognitive map is an abstract, high-level representation of space. However, the space that is encoded is not always trivial; some place cells that encode the x - y location of a rat running around in a square arena encode temporal intervals during a task in which rats are trained to run on a treadmill for a certain amount of time (68). They fire at specific auditory frequencies in a task where rats are trained to release a lever at a particular value of a changing frequency stimulus (69). In bats, mammals that are similar in size to rodents but operate in a three-dimensional environment, neurons of the hippocampal formation appear to encode three-dimensional space (70). In bats and rats observing a fellow animal performing a task for a reward, these neurons encode the locations of both the viewing animal and the behaving animal (71, 72). These experiments point to an inherent plasticity in computation that results in a flexible choice of the coordinate system of the cognitive map. These coordinates, in turn, are thought to be the framework for the other major role of the hippocampus—the formation of episodic memory. The hippocampus is thought to form memories by anchoring them to the framework of the cognitive map (73). It is as though the map forms the key to the dictionary of mammalian memory.

The cognitive map is thought to be constructed through the integration of several sensory modalities, and yet in terms of connectivity, the hippocampal formation is many synapses away from primary sensory afferents. This separation affords this region the opportunity to be flexible, abstract, and deliberate in how sensory information is integrated into the cognitive map. In addition, the internal network dynamics in these regions sustain and update firing patterns, giving rise to computations such as path integration (74). How sensory cues interact with internal dynamics to update the cognitive map is largely unknown. When rats forage in a virtual-reality apparatus where the relationship between cues can be modulated, hippocampal place cells are more influenced by visual landmarks (75, 76), whereas entorhinal grid cells are more influenced by self-motion cues (75, 77). In addition, these modulations have an aftereffect on the path integration once visual landmarks are turned off; the rate of update of position in the cognitive map is recalibrated by the past history of landmark movement (76).

Even though the cognitive map in mammals has been the subject of intense scientific exploration for decades, much of its role in behavior remains unknown. In robotics, environmental representations are used as a precursor for navigation. Typically, the coordinates of these representations are predetermined, and information is filled in by the robot through a combination of codependent iterations of mapping (placing sensory features within the coordinates) and localization (determining self-position within the coordinates) (78). We do not know how exactly animals solve this problem of simultaneous localization and mapping (SLAM). Evidence from the cognitive map indicates that they go about it differently than robots. Sensory integration is performed simultaneously but with different weightings in multiple networks linked through feedback connections, and the coordinates of these representations are dependent on the task. Traditional robotic SLAM and navigation might benefit greatly from an understanding of the advantages of flexible, parallel representations. Ultimately, neuroscientists would also benefit from control theorists examining hippocampal computations with an eye toward understanding how these representations might influence closed-loop behavior.

3.4. Spatially Distributed Computation

Even though we assigned the role of controller to the brain and central nervous system, that certainly does not mean that computation is solely the domain of neural dynamics. Computation can be performed just as well by dynamics in physical substrates toward the extremities: in sensory systems, motor systems, and regulatory mechanisms. For example, the hydrodynamics within the semicircular canals of the ear cause an output signal that integrates head angular acceleration into angular velocity and rejects low- and high-frequency components (79, 80). The mammalian retinal circuitry increases the visual contrast gain (81), adapts to changing intensity (82), decorrelates the statistics of natural images (83), and can even dynamically change these computations depending on visual input (84). Antagonist muscle groups in the limbs coactivate to change the mechanical impedance of the limb (85), resulting in zero-delay reflexive responses to perturbations (5).

These computations are sometimes a side effect of how these mechanisms are physically manifest; they are simply the dynamics of the mechanism. However, evolution does not recognize our distinction between dynamics and computation. The cockroach antenna is an excellent example. The American cockroach, *Periplaneta americana*, can follow surfaces in its environment at speeds of up to 80 cm/s using antennae as tactile sensors (86). Behavioral data suggest that a controller using the antenna as input would need access to both the distance to the wall (position) and its rate of change (velocity) with respect to the cockroach, and correlates of both position and velocity were indeed found in activity recorded in the antennal nerve (87). Further investigation revealed that individual mechanoreceptors in the antennae encode the velocity signal, and the position signal arises in their combined population output due to differential latency and filtering (18). The exponentially decreasing stiffness profile of the antenna segments increases the look-ahead distance and rejects large deflections, thus keeping the antenna in contact with the wall (88). Together, these computations performed by the antennal system allow the cockroach to respond to antennal touch within the window of conduction delays along the flagellum (18, 86).

In robotics, computation is typically not an intended design feature of the sensory or actuation systems. Computation is relegated to the controller, which is often a central processor that in fact tries to compensate for the dynamics imposed by sensing and actuation. Such modularity certainly simplifies—but may also limit—engineering design. Control-theoretic modeling of these animal sensorimotor systems allows us to understand when it might be useful to delegate some of the computation to the extremities and actually use the dynamics of sensing and actuation to our advantage.

3.5. Nested Feedback Loops on Multiple Levels

Organisms have feedback inherent at almost all levels of organization—at the cellular (89), tissue (90), organ (91), system (80), and organismal (92, 161, 162) scales. It is possible that these seemingly endless levels of feedback are simply a by-product of the messy process of evolution by natural selection. If this were true, then a single centralized controller in a synthetic system that has access to all inputs and authority over all outputs should perform equally as well as a system with a hierarchy of feedback. It is likely, however, that nested feedback loops give an organism an advantage in terms of redundancy, robustness, and responsiveness. On the flip side, such redundancy can inhibit inquiry, since certain types of perturbations (such as sensory ablation) can have little effect on behavioral output (93).

Sequential composition of feedback controllers in robotics allows complex behaviors to emerge by appropriately stacking controllers such that the goal state of a higher-level controller falls within the domain of attraction of lower-level controllers (94, 95), and analogous processes in

biology may provide new design ideas. In humans, for example, a hierarchical structure of sensorimotor control could allow a task to be performed through different control strategies based on the nature of the perturbations (17). The control of sensorimotor behavior by the nervous system can be thought of as a computation abstracted through levels of hierarchical feedback (96).

When studying nested feedback loops, it is crucial not to get drawn unnecessarily into the complexities of an organic implementation. Rather, it is useful to ask how we might simplify the design. Are all the levels of feedback truly necessary? The use of top-down control-theoretic analysis (Section 4.3) paired with experimental system identification (Section 4.7) allows us to probe the function of each level of the hierarchy and determine whether they are required in a synthetic implementation of the control system (8).

3.6. Feedback Delays

It is a definitional trait of dynamical systems that past states influence future states. Thus, in general, dynamical systems will propagate, loosely speaking, time-delayed information. However, dynamical systems theory establishes a distinction between lags caused by dynamics and those caused by delays. An ideal delay would take an infinite number of states to encode. Notwithstanding this nuanced distinction, delays and other phase lags can have a significant effect on an organism performing a feedback-regulated behavior, causing instabilities, changes in equilibria, and oscillations (97). Different components of the behavior—the sensory systems, brain and central nervous system, motor systems, and environment—produce different amounts of delay. Organisms exist across a staggering range of scales, from bacteria ($\sim 10^{-6}$ m) to blue whales ($\sim 10^1$ m). Even within the relatively narrow scope of terrestrial mammals, size scales can range from the shrew ($\sim 10^{-2}$ m, $\sim 10^{-3}$ kg) to the elephant ($\sim 10^0$ m, $\sim 10^3$ kg).

The total delay in terrestrial mammals increases with animal mass (M) proportional to $M^{0.21}$, such that large mammals could experience nearly 20 times the delay of small mammals, but they fortunately enjoy slower timescales of locomotor instability and control (98). Delays are the clearest example of a bug of biological computation—there is no meaningful way in which delay can help in sensorimotor control, and engineered systems can easily outperform their biological analogs in terms of delay. That said, there is a silver lining: Because delay can be such a tremendous bottleneck for sensorimotor control performance in animals, it can be an advantage experimentally and computationally when trying to understand the principles of sensorimotor control (99–101).

3.7. Computational Noise

Noise is an inescapable and permeating feature of real-world dynamical systems. In a sensorimotor loop, noise is induced at the sensory periphery, in terms of the variability of the neural signal transduced from identical sensory input. The data-processing inequality theorem imposes a fundamental constraint on data inference from the world: The information lost due to noisy acquisition cannot be recovered through physical computation (102). Noise also enters the sensorimotor loop at the motor periphery, through variability of actuation on the environment. Muscular force output depends on the impedance of the environment. Even when adaptive controllers are referenced to internal models of environment impedance and muscular force generation, the output could contain variability, and these effects are compounded by factors such as fatigue and response delay.

Sensory and actuation noise are indeed issues in robotic actuation as well. Robotic control, however, being typically executed through digital computation, is deterministic (barring pseudorandom software implementations such as stochastic planning algorithms). This fact is a fundamental architectural difference between *in vivo* and *in silico* computation. Neural computation is inherently stochastic, and each step of computation adds its own layer of noise and

stochasticity to the information being transferred. Noise is thus injected at every stage of the sensorimotor loop and is a major cause of variability of behavior across trials (103). Neural noise can induce stochastic phase locking and cycle skipping and can alter the shape of tuning curves (104). It can also shift neural dynamics across bifurcation points, causing noise-induced transitions (105).

Biologically inspired controllers that account for the presence of internal noise have been developed in the optimal control framework for both sensorimotor control (106) and estimation of time (107, 108). Noise can confer advantages to stochastic computing, such as raising signals above threshold through stochastic resonance (109, 110) and decorrelating and whitening sensory signals (111, 112).

4. A ROAD MAP FOR NEURO-INSPIRED CONTROL

Clearly, the organismal nervous system has much to teach us about control theory. How do we go about decoding these secrets of neural control? Here are a few suggestions.

4.1. Choose a Champion Organism

The term champion organism is credited to the eminent neuroethologist Walter Heiligenberg. Such an animal has superior performance in a particular sensorimotor behavior and has neuronal structures that “incorporate and optimize particular neuronal designs that may be less conspicuous in organisms lacking these superior capabilities” (113, p. 249). The organism performs the behavior readily, robustly, and better than most others, in its natural environment. This indicates that superiority in this behavior is likely important to the survival and reproduction of the organism, and the necessary evolutionary capital has been spent toward its optimization. This also increases the chances that we can identify useful characteristics of this behavior by experimental exploration. For this purpose, the behavior of the champion organism must also be experimentally tractable: It should be reliably produced in a controlled laboratory setting, and the experimenter should be able to measure relevant outputs, control and perturb relevant inputs, and constrain unnecessary dimensions (see **Figure 2**).

Finding a champion organism for a particular behavior is not easy. For this reason, neuroethologists often investigate species that are more exotic than the traditional model organisms of biology—flies, worms, rodents, and primates. We also advocate this approach for the aspiring neuro-inspired control theorist. For example, weakly electric fish, in addition to having an exotic sensory system, are champions of using sensorimotor integration (117) and active sensing (27, 31) to perform an untrained, experimentally tractable, one-dimensional task: following a moving refuge (25, 26).

4.2. Relate the Behavior to Ethology

As mentioned above, it also helps to understand why the behavior relates to the organism in its ecological context. An ethologically irrelevant behavior may be interesting, but we run the risk of investigating a controller that is not operating in the conditions for which it evolved. Sponberg et al. (118) demonstrated that the hawk moth *Manduca sexta* incurs a performance deficit in dim light while tracking artificial flowers moving at frequencies greater than 2 Hz because it needs to integrate longer temporal durations to obtain the necessary sensitivity. The flowers that the moths feed from in the wild, however, sway well below the frequency where the luminance-dependent visual processing produces a significant decrease in sensorimotor tracking performance. Rather than a curious coincidence that tracking bandwidth adequately captures natural flower movement,

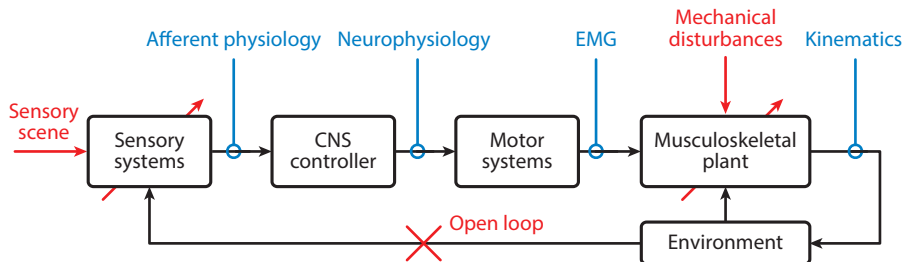


Figure 2

Simplified schematic of a task-level neuromechanical control. There are two categories of signals: applied perturbations (*red*) and measured responses (*blue*). Many systems admit direct movement of the sensory scene, facilitating system identification; mechanical disturbances provide another class of perturbations (114). One can systematically alter subsystem properties (*red diagonal arrows*); mechanical modulations can be achieved, for example, by changing flow conditions in an air or water tunnel or adding mass to the body, and sensory modulations can be achieved, for example, by dimming the lights or, in the special case of electric fish, by changing conductivity (27). Tethering animals (115, 116) enables the opening of the reafferent feedback loop (*red X*), one of a hierarchy of experimental topologies (8). Kinematics, forces, afferent signals, central nervous system (CNS) neurons, and electromyography (EMG) signals are accessible in a wide variety of animal models.

the researchers suggested that not only did the flowers and moths coevolve for odor–olfaction and appearance–vision, but in fact their performance in tracking motion dynamics may have coevolved to foster pollination. Relating ethological context to sensorimotor control is not limited to the control of movement; for example, the control of social electrogenesis is more sophisticated in species of electric fish that are found in groups in their natural habitat (119).

4.3. Use Top-Down, Integrative Modeling as a Starting Point

In control theory, a system (be it biological or mechanical) is described by its dynamics—the transformation between inputs and outputs. Such models are commonly described in terms of time-domain ordinary differential equations, allowing rich representations of nonlinear dynamics. In many cases of interest, the dynamics can be linearized in a neighborhood of an equilibrium, affording the use of powerful representations, such as impulse response functions or frequency response functions, but this must be done cautiously, since even for small signals around equilibria, biological systems can display interesting nonlinearities, such as a failure of superposition despite adhering to homogeneity (52). Whatever the modeling paradigm (differential equations, transfer functions, etc.), such models are agnostic to their physical realization, i.e., the mechanism by which these mathematical transformations manifest in the real world. Thus, the mechanism responsible for the vestibulo-ocular reflex—the compensatory movement of the eyes in response to head movements—has essentially the same dynamics as a mass connected to a spring and damper (80), both of which can be modeled as second-order transfer functions. In fact, a set of variables whose evolution in time represents the dynamics of the model, its states, need not have any direct physical meaning at all.

The direction of scientific inquiry classifies the process of modeling an unknown system into two categories: top down and bottom up. The decision of which process to use is also determined by the type of data available, the assumptions and analytic techniques used, and the types of models that emerge from these analyses.

Bottom-up modeling typically refers to synthesizing a complex system from simpler models based on lower-level physical, chemical, or biological principles. This process typically utilizes mechanistic models, governed by the lower-level laws that dictate their behavior, such as the mechanisms at the level of individual neurons (120). Higher-level models (i.e., networks and behavior) are synthesized through the composition of these mechanistic models. The mechanistic models are analytically formulated or computationally simulated, and experimental data are used to validate them at the level of either mechanisms or higher compositions. Bottom-up modeling is also referred to as reductionism, and the underlying philosophy is that the whole is the sum of its parts. Ultimately, however, practical (e.g., computational intractability) and conceptual (e.g., symmetry breaking) factors preclude the possibility of using bottom-up modeling alone (121).

Top-down modeling, on the other hand, refers to deconstructing mechanisms from their compositions and is critical to understanding neural systems and behavior (122). Experimental data at the highest level (e.g., a behavior of interest) are used to construct abstract, often idealized models. Investigation then proceeds down the hierarchy, producing lower-level models that quantitatively describe mechanisms and requiring increasingly specialized experiments and analyses. Top-down modeling is inherently data driven, requiring experimental data to both construct the models and verify them.

The ensuing application of control theory is largely insensitive to the procedure that produced the model. However, we advocate a top-down approach when applying control theory to understand and derive engineering inspiration from biological systems, for the following reasons.

4.3.1. Top-down models constrain mechanistic models. Mechanistic models, as described above, are derived from a specific hypothesis about the architecture and function of a mechanism. When one uses mechanistic models to synthesize a high-level model, the parametric space can be high dimensional and covariant. Top-down modeling can help define the scope of investigation for these parameters, which define the dynamic responses that the richer mechanistic models need to achieve. A top-down model provides the necessary responses for a mechanism to instantiate, and a bottom-up model provides sufficient laws for its implementation. Both approaches are valuable for a complete understanding of a given biological system, but it is our assertion that a top-down approach can make the process more efficient.

Consider the neural circuitry responsible for the vestibulo-ocular reflex, which is incredibly complex even though the reflex itself is simple by biological standards. The hydrodynamics of the endolymph in the semicircular canals in the ear cause deflections of the hair cells of the crista. Primary afferent neurons relay this information to secondary vestibular neurons, which in turn stimulate first the premotor neurons of the visual system and then the motor neurons responsible for lateral movements of the eye, which in turn results in the mechanical movement of the eye itself. There are mechanistic models that physiologists and computational biologists have painstakingly developed for each of these constituent components. However, the composition of these models must respond like a mass connected to a spring and damper—the second-order dynamical system model of the vestibulo-ocular reflex (80). Knowing this fact ahead of time can accelerate our search for mechanistic models.

A concrete example of how modeling the behavior can assist in quantifying its components can be found by considering again the weakly electric fish, *E. virescens* (see **Figure 3a,b**). The behavioral transfer function of refuge tracking was approximated as a second-order system (much like the vestibulo-ocular reflex) by analyzing responses to perturbation of the refuge through sinusoidal trajectories (123). However, the behavior is clearly a composition of the plant (body and fin) and controller (sensors and brain). Knowing that the behavior is second order sets up the relationship between the plant and controller dynamics. If the plant were kinematic, i.e., the relationship

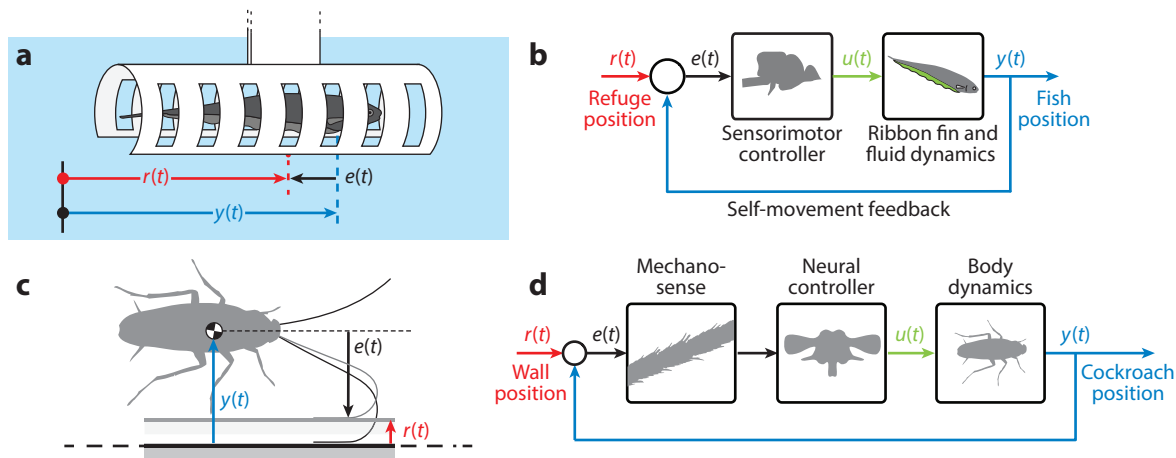


Figure 3

Illustration of top-down control systems modeling on two champion model systems. (a) Refuge tracking in *Eigenmannia virescens*. The fish moves back and forth [position = $y(t)$] to maintain a constant position with respect to a moving refuge [position = $r(t)$] (123). (b) Control-theoretic abstraction of the behavior. Through its sensory systems, the fish perceives the error signal, $e(t)$, between its own position and that of the refuge. The neural controller applies a control input, $u(t)$, to the biomechanical or environmental plant, which in turn produces the output, the fish's position $y(t)$ (52). (c) Wall following in *Periplaneta americana*. As it runs, the cockroach maintains a constant offset or error, $e(t)$, between its own position, $y(t)$, and the wall position, $r(t)$ (87). (d) The error signal is measured as deflection of the antennal mechanosensors (18) and processed by the neural controller (99) to generate a control signal to the body dynamics. The interaction of the body dynamics and the environment results in the output, the cockroach position $y(t)$. Figure adapted from Reference 7, with portions created by Eatai Roth and Eric Fortune.

between motor commands and body position were purely integrative, then the corresponding controller would have to be low pass—a categorically opposite prediction. By contrast, if the plant were mechanical, i.e., the inertia of the fish has a significant role in determining its position, then the controller would have to be high pass (123). Later experiments specifically targeting the ribbon fin locomotor system came to the conclusion that the plant is, in fact, best described as mechanics based (14), implicating a high-pass controller. Thus, there is a critical interdependence between the components that lead to a behavior, and a top-down model can leverage this dependence to transfer information about one component to a better understanding of another, subject to assumptions about the system's feedback topology (8).

4.3.2. Top-down models incorporate emergent dynamics. It is logical to expect that the composition of quantitatively modeled components would yield predictable results. However, in practice, it is well known that this is not always the case. The interaction of a large number of subsystems invariably results in so-called emergent dynamics—behavior that is not easily predicted from the constituents (121), especially in feedback-regulated systems (7). Emergent dynamics can be due to the stochasticity in the components, the strength of their relative interactions, their timing and delays, differences in their initial conditions, and, critically, the nature of their feedback topology.

A canonical example of emergent dynamics in biology is attractor dynamics in the brain. A network of interconnected neurons can produce dynamics similar to a bump of activity moving along a (physically nonexistent) line or sheet (124, 125). These dynamics are thought to be crucial to the animal's representation of space and time and the formation of memory, and can perform computations such as path integration (126, 127). Inputs to the network as well as intrinsic dynamics

of the network can cause the bump to move around, increasing or decreasing the probability of firing of neurons in the network. This emergent low-dimensional abstraction is not apparent at a cellular or synaptic level, even if one were to painstakingly model the properties of the neurons individually and set up the tens of thousands of synaptic connections that hippocampal neurons make with each other (128). We propose that constructing the higher-level description is a crucial first step for investigating the lower-level ones.

4.4. Choose a Simple Plant

At first glance, it may seem counterintuitive to investigate the control of simple plants in neuroscience; after all, control engineers have little trouble controlling such systems. Indeed, we are investigating animals to hopefully discover new control strategies to regulate complex mechanisms in complex and dynamic environments. Would not a simple plant have an equally simple controller? This is certainly true in engineering, where we have the luxury of designing a controller for a preexisting plant, and we try to formulate a solution that is only complicated enough to do the job.

The nervous system faces a completely different set of constraints. We believe that, in many cases, neural mechanisms that were learned or evolved to perform one task are co-opted to control a different task or plant. For example, in weakly electric fish, the electric fields in most species are myogenic—i.e., generated from modified, nonmoving muscle tissue. It is a motor region of the brain acting through a motor neuron that activates the electric organ (129). In fact, some species of electric fish have gone one step further: The electric field is neurogenic—i.e., produced directly from neural tissue. In these species, juveniles retain the myogenic organ, whereas mature adults have a neurogenic organ (130). Thus, a neural network that evolved to control movement now activates muscle tissue (in juveniles) and neural tissue (in adults)—each with its own dynamics—to produce electric fields in the surrounding water.

Plant dynamics often operate around hyperbolic equilibria, i.e., equilibria for which the local linearization provides a faithful instantiation of the qualitative behavior near the equilibrium, per the Hartman–Grobman theorem. Biological systems often control such almost-linear plants with nonlinear, adaptive controllers, so that the resultant closed-loop system can operate at the confluence of stability and maneuverability. In the above-mentioned champion of refuge tracking, *E. virescens*, the fore–aft position dynamics can be modeled as a second-order lumped-parameter system. In this system, the intersection point of the forward- and backward-propagating portion of its longitudinal ribbon fin, termed the so-called nodal point, is used as the control input (14). And yet these fish exhibit interesting nonlinearities, such as active sensing and adaptive prediction of refuge motion—nonlinearities that almost certainly extend to more complex plants but are much more easily characterized in the context of a simple, known plant model, especially given the sensitivity of controller estimates to plant dynamics (123).

Around the equilibrium of hovering flight, the relationship between the abdominal angle and body pitch rotation of the hawk moth *M. sexta* is captured by linear, second-order dynamics (115). However, the closed-loop system is very nearly unstable, giving the animal easy access to an unstable but more maneuverable regime (115, 131). The discovery of how the animal uses its flexible airframe for flight control was facilitated by reduction to a single pitch degree of freedom.

Choosing these simple, linearizable plants allows one to focus on identifying the highly nonlinear properties of the neural controller. Complex plant dynamics make this problem considerably more challenging, although advances in aeromechanical modeling, for instance, can still be used to shed light on neural control mechanisms even in the context of complex flight mechanics (132). Furthermore, simplified plant modeling leaves many unanswered questions—such as how

the animal actually manages a high-dimensional plant (48)—that can be addressed only when we embrace such complexity.

4.5. Determine the Ability to Perturb Inputs, Modulate Plant Dynamics, and Measure Outputs

A system identification experiment is critically dependent on the experimental topology—how the inputs and outputs of the biological system interconnect with those of the experimental system (8). For a biological system, the experimenter must determine what the inputs and outputs are. Inputs need not always be sensory; both direct neural stimulation and a change in reward amount are viable inputs. The experimenter must also determine what output variables to measure from the organism, such as kinematics, forces, behavioral metrics, or neural measurements. Similar to inputs, outputs can be constrained or unconstrained—for example, an animal can be restricted to locomote in a one-dimensional track.

Perturbation of plant dynamics is another intervention that can be valuable to understand the effect of the plant and the controller–plant interaction. For example, removing the antennal flagella disrupts the flight stability of hawk moths, while reattaching them restores flight control (133). In a less drastic intervention, a biased force field was applied to humans performing a reaching task and used to test hypotheses about the structure of the underlying adaptive controller (134).

When deciding on inputs, outputs, and manipulations of plant dynamics, one must be aware of the benefits and consequences of each choice. Every constrained degree of freedom pushes an organism out of its normal ethological operating range. Just as we advocate for experiments on intact animal models, we also strongly encourage intact behavioral paradigms. Applying minimal constraints and perturbations ensures that the controller is also operating within its nominal dynamic range and is not creating responses that are outside its evolutionary specifications. Depending on the system at hand, a natural place to probe a given behavior when identifying neural control algorithms is at the levels of self-movement kinematics (measured by sensors) and electromyograms (motor outputs) (135, 136); doing so enables so-called joint input–output system identification and can be used to separately identify the plant and feedback.

4.6. Determine the Experimental Topology

In the real world, an animal is in a closed feedback loop with its environment: Its motor outputs affect the body and environment, which in turn affect the animal's sensory inputs. A well-controlled experiment is, in a very real sense, a simulated environment for the organism in which we determine the rules of the world. This is what the experimental topology refers to—the nature of the perturbations and inputs applied to the biological system and how the outputs of the system affect its inputs.

For an open-loop topology, the outputs of the biological system do not influence its inputs, as would be the case in a tethered preparation (115, 137). Open-loop experiments can reveal important information about the behavior and control computations. However, there can be trial-to-trial variability due to internal dynamics, initial conditions, or uncontrolled inputs (116). In addition, dynamical systems can have equilibria that are unstable (138) or bifurcations that are not amenable to effective perturbation by open-loop stimulation (139).

In an experimentally closed-loop topology, a subset of the behavioral outputs of the animal determine its inputs through an experimental feedback law (31, 140). This approach is effectively modifying environmental feedback, typically in a way that is subtle enough that the animal remains in the same behavioral regime. In the fruit fly *Drosophila melanogaster*, these closed-loop

manipulations have been extensively used to study the impact of visual processing on flight control. The difference between the left and right wing-beat amplitudes of a tethered fly is used to control visual stimuli in the virtual-reality arena surrounding it (141). These experiments have been used for decades to understand the impact of visual cues (142) and optic flow (143) on flight control. Magnets attached to fruit flies allowed instantaneous perturbations in the roll axis, which showed that flies use a proportional–integral controller to correct these perturbations rapidly (144). An apparatus that allows two-photon calcium imaging while head-fixed flies walk on an air-supported ball (145) enabled the discovery of a neural correlate of the hypothesized ring-attractor angular integration system for encoding orientation in the central complex of flies (146, 147).

Behaviorally closed-loop manipulations can also be used to completely alter the dynamics of the underlying behavior to enable easier system identification. In *E. virescens*, the jamming avoidance response—in which fish adjust their electric field frequency away from those of interfering conspecifics—has been extensively studied as an escape behavior. Even though the neural circuitry underlying the jamming avoidance response was uncovered decades ago (148), the dynamics of the behavior was difficult to characterize because escape behaviors are inherently unstable: The animal avoids the stimulus, in this case by changing its frequency. Madhav et al. (138) overcame this problem by creating a closed-loop experimental system that stabilizes these unstable dynamics. This system enabled perturbation experiments to characterize the local linear and global nonlinear behavior and predict a saddle-node bifurcation in its dynamics.

A rather recent paradigm is to use neurally closed-loop topologies, where the experimenter gains access to the animal’s nervous system during a behavioral experiment and stimulates the nervous system in response to organismal outputs (149, 150), uses states of the nervous system to modulate organismal inputs (151), or stimulates the nervous system using its own states (152). These experiments are more challenging to perform, and estimation and stimulation of neural states can be imprecise. Nevertheless, neurally closed-loop topologies offer a powerful new way to directly modulate the controller and quantify its impact on sensorimotor behavior.

4.7. Apply System Identification

To complete the experimental design, one must also determine the input signals to be applied. System identification dictates that a dynamical system be perturbed using inputs of specific characteristics. The inputs are designed to be sufficiently rich (153) to persistently excite the dynamics in the range of interest. There are limitless choices for signal design, but common choices are step inputs (140), band-limited noise (137), chirps (138), and sums of sinusoids (52, 133). Here, we advocate a simple approach of using sums of sinusoids, not because they are “best” in some mathematical sense [indeed, noise may be better for a variety of signal-processing reasons (154)] but because the responses provide an easy-to-read signature in the frequency domain that indicates whether the animal is performing the desired behavior. But this is merely a design choice, and the animal or behavior under study needs to be taken into account. Indeed, a critical test of linearity is whether the response to one class of inputs (say, a step input) can be used to fit a model that predicts the response to another class (say, sums of sinusoids) (155). In fact, the failure of responses from single sinusoidal inputs to predict the response to sum-of-sine inputs led to the identification of critical tracking nonlinearities in weakly electric fish (52).

These inputs, paired with the corresponding outputs of the system, form the nonparametric input–output model. Nonparametric models are often redundant and high dimensional and can be fitted to parsimonious, parametric models through time-domain or frequency-domain techniques (156)—for example, transfer functions, autoregressive models, or state-space models (157).

In behaviorally or neurally closed-loop topologies, identification can also be performed at different dynamical equilibria of the combined plant–controller–experiment closed loop. Clamping the inputs to different behavioral or neural variables can cause the system to change equilibria or exhibit completely different dynamics (138). However, one must be cautious about not causing the system to veer beyond its normal, ethologically relevant dynamical range. Doing so may still result in quantifiable dynamics but will provide very little useful information about the role of the controller in shaping the behavior.

When applied to a biological system, the goal of system identification is to distill the essential dynamics of the behavior—the transformation from sensory input to motor output mediated through the brain and central nervous system—into a low-dimensional parametric state space. This state space, combined with system identification of the plant, allows us to construct low-dimensional models of the neural controller (115).

4.8. Tune Control Parameters

A parametric model of the organismal controller can, of course, be realized in synthetic hardware or computation. However, it is important to realize that the particular parameters of the controller are tuned to that plant and may be wildly irrelevant to the engineered control system that we desire to deploy. Thus, the next step is to scale and tune the parameters of the controller to fit the dynamics of the engineering plant.

This can be at first done through simulations. However, to observe and improve the impact of the controller in the real world, it becomes necessary to implement it in a synthetic plant model, a so-called robophysical prototype (158). Robophysical models incorporate the essential feature of the biological plant dynamics that the controller coevolved with while being subject to the different scale, substrate, and dynamics of an engineered system. The cockroach wall-following system (**Figure 3c,d**) provides a useful example; a synthetic version of a cockroach antenna placed on a robot allowed it to follow walls, and the proportional and derivative gains of distance to the wall (the key parameters of the original cockroach controller) could be tuned (87). A more intricate version of the antenna with tunable stiffness between segments (159) allowed testing of hypotheses regarding the role of exponentially decreasing antenna stiffness and the role of distally pointing hairs on the antennae as a mechanical stabilizer to simplify control (15). Robophysical prototypes are often designed with tunable plant parameters that, along with the tunable control parameters, allow the control theorist to both better quantify the role of the original biological controller and modify its dynamics to suit the needs of the engineering applications.

5. OUTLOOK

While this review has focused on centralized regulatory mechanisms (i.e., negative feedback), closed-loop mechanisms can also destabilize a dynamical system and/or make it more sensitive to its inputs (i.e., create positive feedback) (120). Also not addressed—but every bit as important—is understanding where a given controller lies on the feedback-versus-feedforward and centralized-versus-decentralized planes (160). This point underscores the need to dig deeper toward a component-level understanding of control. Indeed, while in this review we advocate a top-down, data-driven approach as a starting point for many problems in neural control, an important goal in reverse engineering a biological control system is to understand it in terms of bottom-up neuromechanics. So one should not misinterpret our course as being dismissive of approaches that synthesize fundamental principles at lower levels into a top-down understanding of control. On the contrary, ignoring biological details can be a major mistake when translating

biological control principles from nature to artifact. In doing so, tackling the system dynamics at multiple levels of organization is essential (48).

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 Chris Yang for critical neural feedback. This work was supported in part by the Army Research Office under the SLICE Multidisciplinary University Research Initiative (MURI) Program, award W911NF1810327 (N.J.C.), and a Kavli Postdoctoral Fellowship (M.S.M.).

LITERATURE CITED

1. Wallace AR. 1858. On the tendency of varieties to depart indefinitely from the original type. *Proc. Linn. Soc. Lond.* 3:53–62
2. Gross CG. 1998. Claude Bernard and the constancy of the internal environment. *Neuroscientist* 4:380–85
3. Milhorn HT. 1966. *Application of Control Theory to Physiological Systems*. Philadelphia: Saunders
4. Acott TS, Kelley MJ, Keller KE, Vranka JA, Abu-Hassan DW, et al. 2014. Intraocular pressure homeostasis: maintaining balance in a high-pressure environment. *J. Ocul. Pharmacol. Ther.* 30:94–101
5. Loeb G. 1995. Control implications of musculoskeletal mechanics. In *Proceedings of 17th International Conference of the Engineering in Medicine and Biology Society*, Vol. 2, pp. 1393–94. Piscataway, NJ: IEEE
6. Savage MV, Brengelmann GL. 1996. Control of skin blood flow in the neutral zone of human body temperature regulation. *J. Appl. Physiol.* 80:1249–57
7. Cowan NJ, Ankarali MM, Dyhr JP, Madhav MS, Roth E, et al. 2014. Feedback control as a framework for understanding tradeoffs in biology. *Integr. Comp. Biol.* 54:223–37
8. Roth E, Sponberg S, Cowan NJ. 2014. A comparative approach to closed-loop computation. *Curr. Opin. Neurobiol.* 25:54–62
9. Chiel HJ, Beer RD. 1997. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends. Neurosci.* 20:553–57
10. Hooven FJ. 1978. The Wright brothers' flight-control system. *Scientific American*, Nov. 1, pp. 166–85
11. Vecchio DD, Sontag ED. 2009. Engineering principles in bio-molecular systems: from retroactivity to modularity. *Eur. J. Control* 15:389–97
12. Del Vecchio D, Ninfa AJ, Sontag ED. 2008. Modular cell biology: retroactivity and insulation. *Mol. Syst. Biol.* 4:161
13. Tytell ED, Holmes P, Cohen AH. 2011. Spikes alone do not behavior make: why neuroscience needs biomechanics. *Curr. Opin. Neurobiol.* 21:816–22
14. Sefati S, Neveln ID, Roth E, Mitchell T, Snyder JB, et al. 2013. Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability. *PNAS* 110:18798–803
15. Mongeau JM, Demir A, Lee J, Cowan NJ, Full RJ. 2013. Locomotion- and mechanics-mediated tactile sensing: antenna reconfiguration simplifies control during high-speed navigation in cockroaches. *J. Exp. Biol.* 216:4530–41
16. Alcocer-Cuarón C, Rivera AL, Castaño VM. 2014. Hierarchical structure of biological systems: a bio-engineering approach. *Bioengineered* 5:73–79
17. Loeb GE, Brown IE, Cheng EJ. 1999. A hierarchical foundation for models of sensorimotor control. *Exp. Brain. Res.* 126:1–18
18. Mongeau JM, Sponberg SN, Miller JP, Full RJ. 2015. Sensory processing within cockroach antenna enables rapid implementation of feedback control for high-speed running maneuvers. *J. Exp. Biol.* 218:2344–54
19. Wolpert DM, Flanagan J. 2001. Motor prediction. *Curr. Biol.* 11:R729–32
20. Bajcsy R. 1988. Active perception. *Proc. IEEE* 76:996–1005

21. Grant RA, Mitchinson B, Fox CW, Prescott TJ. 2009. Active touch sensing in the rat: anticipatory and regulatory control of whisker movements during surface exploration. *J. Neurophysiol.* 101:862–74
22. Jung SN, Borst A, Haag J. 2011. Flight activity alters velocity tuning of fly motion-sensitive neurons. *J. Neurosci.* 31:9231–37
23. Lederman SJ, Klatzky RL. 1987. Hand movements: a window into haptic object recognition. *Cogn. Psychol.* 19:342–68
24. Gibson J. 1962. Observations on active touch. *Psychol. Rev.* 69:477–91
25. Bastian J. 1982. Vision and electroreception: integration of sensory information in the optic tectum of the weakly electric fish *Apteronotus albifrons*. *J. Comp. Physiol. A* 147:287–97
26. Rose GJ, Canfield JG. 1993. Longitudinal tracking responses of *Eigenmannia* and *Sternopygus*. *J. Comp. Physiol. A* 173:698–700
27. Stamper SA, Roth E, Cowan NJ, Fortune ES. 2012. Active sensing via movement shapes spatiotemporal patterns of sensory feedback. *J. Exp. Biol.* 215:1567–74
28. Uyanik I, Stamper SA, Cowan NJ, Fortune ES. 2019. Sensory cues modulate smooth pursuit and active sensing movements. *Front. Behav. Neurosci.* 13:59
29. Nelson ME, MacIver MA. 1999. Prey capture in the weakly electric fish *Apteronotus albifrons*: sensory acquisition strategies and electrosensory consequences. *J. Exp. Biol.* 202:1195–203
30. Hofmann V, Sanguinetti-Scheck JI, Kunzel S, Geurten B, Gomez-Sena L, Engemann J. 2013. Sensory flow shaped by active sensing: sensorimotor strategies in electric fish. *J. Exp. Biol.* 216:2487–500
31. Biswas D, Arend LA, Stamper SA, Vágvölgyi BP, Fortune ES, Cowan NJ. 2018. Closed-loop control of active sensing movements regulates sensory slip. *Curr. Biol.* 28:4029–36.e4
32. Kunapareddy A, Cowan NJ. 2018. Recovering observability via active sensing. In *2018 American Control Conference*, pp. 2821–26. Piscataway, NJ: IEEE
33. Liu D, Todorov E. 2007. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J. Neurosci.* 27:9354–68
34. Todorov E, Jordan MI. 2002. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5:1226–35
35. Jun JJ, Longtin A, Maler L. 2016. Active sensing associated with spatial learning reveals memory-based attention in an electric fish. *J. Neurophysiol.* 115:2577–92
36. McNamee D, Wolpert DM. 2019. Internal models in biological control. *Annu. Rev. Control Robot. Auton. Syst.* 2:339–64
37. 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
38. Miall R, Wolpert D. 1996. Forward models for physiological motor control. *Neural Netw.* 9:1265–79
39. Wolpert DM, Miall R, Kawato M. 1998. Internal models in the cerebellum. *Trends. Neurosci.* 2:338–47
40. Baumann O, Borra RJ, Bower JM, Cullen KE, Habas C, et al. 2015. Consensus paper: the role of the cerebellum in perceptual processes. *Cerebellum* 14:197–220
41. Bhanpuri NH, Okamura AM, Bastian AJ. 2014. Predicting and correcting ataxia using a model of cerebellar function. *Brain* 137:1931–44
42. Körding KP, Ku Sp, Wolpert DM. 2004. Bayesian integration in force estimation. *J. Neurophysiol.* 92:3161–65
43. Körding KP, Wolpert DM. 2004. Bayesian integration in sensorimotor learning. *Nature* 427:244–7
44. Aström KJ, Wittenmark B. 2013. *Adaptive Control*. North Chelmsford, MA: Courier
45. Sun J. 2014. Model reference adaptive control. In *Encyclopedia of Systems and Control*, ed. J Baillieul, T Samad. London: Springer. https://doi.org/10.1007/978-1-4471-5102-9_116-1
46. Conant RC, Ross Ashby W. 1970. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* 1:89–97
47. Tin C, Poon CS. 2005. Internal models in sensorimotor integration: perspectives from adaptive control theory. *J. Neural Eng.* 2:1–37
48. Full RJ, Koditschek DE. 1999. Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *J. Exp. Biol.* 202:3325–32
49. Clarke D, Mohtadi C, Tuffs P. 1987. Generalized predictive control—part I. The basic algorithm. *Automatica* 23:137–48

50. Johansson RS, Cole KJ. 1992. Sensory-motor coordination during grasping and manipulative actions. *Curr. Opin. Neurobiol.* 2:815–23
51. Collins CJS, Barnes GR. 2009. Predicting the unpredictable: Weighted averaging of past stimulus timing facilitates ocular pursuit of randomly timed stimuli. *J. Neurosci.* 29:13302–14
52. Roth E, Zhuang K, Stamper SA, Fortune ES, Cowan NJ. 2011. Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*. *J. Exp. Biol.* 214:1170–80
53. Cutlip S, Freudenberg J, Cowan N, Gillespie RB. 2019. Haptic feedback and the internal model principle. In *2019 IEEE World Haptics Conference*, pp. 568–73. Piscataway, NJ: IEEE
54. Huang J, Isidori A, Marconi L, Mischiati M, Sontag E, Wonham W. 2018. Internal models in control, biology and neuroscience. In *2018 IEEE Conference on Decision and Control*, pp. 5370–90. Piscataway, NJ: IEEE
55. Francis BA, Wonham WM. 1976. The internal model principle of control theory. *Automatica* 12:457–65
56. Tolman EC. 1948. Cognitive maps in rats and men. *Psychol. Rev.* 55:189–208
57. O’Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34:171–75
58. Taube JS, Muller RU, Ranck JB. 1990. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.* 10:420–35
59. Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. 2005. Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–6
60. Savelli F, Yoganarasimha D, Knierim JJ. 2008. Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus* 18:1270–82
61. Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI. 2008. Representation of geometric borders in the entorhinal cortex. *Science* 322:1865–68
62. Wang C, Chen X, Lee H, Deshmukh SS, Yoganarasimha D, et al. 2018. Egocentric coding of external items in the lateral entorhinal cortex. *Science* 362:945–49
63. Deshmukh SS, Knierim JJ. 2013. Influence of local objects on hippocampal representations: landmark vectors and memory. *Hippocampus* 23:253–67
64. Lever C, Burton S, Jeewajee A, O’Keefe J, Burgess N. 2009. Boundary vector cells in the subiculum of the hippocampal formation. *J. Neurosci.* 29:9771–77
65. Høydal ØA, Skytøen ER, Andersson SO, Moser MB, Moser EI. 2019. Object-vector coding in the medial entorhinal cortex. *Nature* 568:400–4
66. Quirk GJ, Muller RU, Kubie JL. 1990. The firing of hippocampal place cells in the dark depends on the rat’s recent experience. *J. Neurosci.* 10:2008–17
67. Zhang S, Schönfeld F, Wiskott L, Manahan-Vaughan D. 2014. Spatial representations of place cells in darkness are supported by path integration and border information. *Front. Behav. Neurosci.* 8:222
68. MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. 2011. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* 71:737–49
69. Aronov D, Nevers R, Tank DW. 2017. Mapping of a non-spatial dimension by the hippocampal entorhinal circuit. *Nature* 543:719–22
70. Yartsev MM, Ulanovsky N. 2013. Representation of three-dimensional space in the hippocampus of flying bats. *Science* 340:367–72
71. Danjo T, Toyozumi T, Fujisawa S. 2018. Spatial representations of self and other in the hippocampus. *Science* 359:213–18
72. Omer DB, Maimon SR, Las L, Ulanovsky N. 2018. Social place-cells in the bat hippocampus. *Science* 359:218–24
73. Derdikman D, Knierim JJ, eds. 2014. *Space, Time and Memory in the Hippocampal Formation*. New York: Springer
74. Savelli F, Knierim JJ. 2019. Origin and role of path integration in the cognitive representations of the hippocampus: computational insights into open questions. *J. Exp. Biol.* 222:jeb188912
75. Chen G, Lu Y, King JA, Cacucci F, Burgess N. 2019. Differential influences of environment and self-motion on place and grid cell firing. *Nat. Commun.* 10:630

76. Jayakumar RP, Madhav MS, Savelli F, Blair HT, Cowan NJ, Knierim JJ. 2019. Recalibration of path integration in hippocampal place cells. *Nature* 566:533–37
77. Campbell MG, Ocko SA, Mallory CS, Low II, Ganguli S, Giocomo LM. 2018. Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nat. Neurosci.* 21:1096–106
78. Dissanayake M, Newman P, Clark S, Durrant-Whyte H, Csorba M. 2001. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Trans. Robot. Automat.* 17:229–41
79. Goldberg JM, Fernandez C. 1971. Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. I. Resting discharge and response to constant angular accelerations. *J. Neurophysiol.* 34:635–60
80. Robinson DA. 1981. The use of control systems analysis in the neurophysiology of eye movements. *Annu. Rev. Neurosci.* 4:463–503
81. Shapley R, Enroth-Cugell C. 1984. Visual adaptation and retinal gain control. *Prog. Retin. Res.* 3:263–346
82. Demb JB. 2008. Functional circuitry of visual adaptation in the retina. *J. Physiol.* 586:4377–84
83. Atick JJ, Redlich AN. 1992. What does the retina know about natural scenes? *Neural Comput.* 4:196–210
84. Rivlin-Etzion M, Grimes WN, Rieke F. 2018. Flexible neural hardware supports dynamic computations in retina. *Trends Neurosci.* 41:224–37
85. Hogan N. 1984. Adaptive control of mechanical impedance by coactivation of antagonist muscle. *IEEE Trans. Autom. Control* 29:681–90
86. Camhi JM, Johnson EN. 1999. High-frequency steering maneuvers mediated by tactile cues: antennal wall-following in the cockroach. *J. Exp. Biol.* 202:631–43
87. Lee J, Sponberg SN, Loh OY, Lamperski AG, Full RJ, Cowan NJ. 2008. Templates and anchors for antenna-based wall following in cockroaches and robots. *IEEE Trans. Robot.* 24:130–43
88. Mongeau JM, Demir A, Dallmann CJ, Jayaram K, Cowan NJ, Full RJ. 2014. Mechanical processing via passive dynamic properties of the cockroach antenna can facilitate control during rapid running. *J. Exp. Biol.* 217:3333–45
89. Brandman O, Meyer T. 2008. Feedback loops shape cellular signals in space and time. *Science* 322:390–95
90. Shraiman BI. 2005. Mechanical feedback as a possible regulator of tissue growth. *PNAS* 102:3318–23
91. Röder PV, Wu B, Liu Y, Han W. 2016. Pancreatic regulation of glucose homeostasis. *Exp. Mol. Med.* 48:e219
92. Powers WT. 1973. Feedback: beyond behaviorism. *Science* 179:351–56
93. Roth E, Hall RW, Daniel TL, Sponberg S. 2016. Integration of parallel mechanosensory and visual pathways resolved through sensory conflict. *PNAS* 113:12832–37
94. Burrige RR, Rizzi AA, Koditschek DE. 1999. Sequential composition of dynamically dexterous robot behavior. *Int. J. Robot. Res.* 18:534–55
95. Kallem V, Komoroski AT, Kumar V. 2011. Sequential composition for navigating a nonholonomic cart in the presence of obstacles. *IEEE Trans. Robot.* 27:1152–59
96. Ballard DH. 2015. *Brain Computation as Hierarchical Abstraction*. Cambridge, MA: MIT Press
97. an der Heiden U. 1979. Delays in physiological systems. *J. Math. Biol.* 8:345–64
98. More HL, Donelan JM. 2018. Scaling of sensorimotor delays in terrestrial mammals. *Proc. R. Soc. B* 285:20180613
99. Cowan NJ, Lee J, Full RJ. 2006. Task-level control of rapid wall following in the American cockroach. *J. Exp. Biol.* 209:1617–29
100. Liu P, Cheng B. 2017. Limitations of rotational manoeuvrability in insects and hummingbirds: evaluating the effects of neuro-biomechanical delays and muscle mechanical power. *J. R. Soc. Interface* 14:20170068
101. Elzinga MJ, Dickson WB, Dickinson MH. 2012. The influence of sensory delay on the yaw dynamics of a flapping insect. *J. R. Soc. Interface* 9:1685–96
102. Beaudry NJ, Renner R. 2011. An intuitive proof of the data processing inequality. arXiv:1107.0740 [quant-ph]
103. Faisal AA, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9:292–303

104. Longtin A. 2003. Effects of noise on nonlinear dynamics. In *Nonlinear Dynamics in Physiology and Medicine*, ed. A Beuter, L Glass, M Mackey, M Titcombe, pp. 149–89. New York: Springer
105. Horsthemke W, Lefever R. 1984. Noise-induced transitions in physics, chemistry, and biology. In *Noise-Induced Transitions: Theory and Applications in Physics, Chemistry, and Biology*, pp. 164–200. Berlin: Springer
106. Todorov E. 2005. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput.* 17:1084–108
107. Carver SG, Fortune ES, Cowan NJ. 2013. State-estimation and cooperative control with uncertain time. In *2013 American Control Conference*, pp. 2990–95. Piscataway, NJ: IEEE
108. Lamperski A, Cowan NJ. 2016. Optimal control with noisy time. *IEEE Trans. Autom. Control* 61:319–33
109. Bulsara A, Jacobs EW, Zhou T, Moss F, Kiss L. 1991. Stochastic resonance in a single neuron model: theory and analog simulation. *J. Theor. Biol.* 152:531–55
110. Gluckman BJ, Netoff TL, Neel EJ, Spano WL, Spano ML, Schiff SJ. 1996. Stochastic resonance in a neuronal network from mammalian brain. *Phys. Rev. Lett.* 77:4098–101
111. Pozzorini C, Naud R, Mensi S, Gerstner W. 2013. Temporal whitening by power-law adaptation in neocortical neurons. *Nat. Neurosci.* 16:942–48
112. Huang CG, Zhang ZD, Chacron MJ. 2016. Temporal decorrelation by SK channels enables efficient neural coding and perception of natural stimuli. *Nat. Commun.* 7:11353
113. Heiligenberg W. 1991. The neural basis of behavior: a neuroethological view. *Annu. Rev. Neurosci.* 14:247–67
114. Jindrich DL, Full RJ. 2002. Dynamic stabilization of rapid hexapedal locomotion. *J. Exp. Biol.* 205:2803–23
115. Dyhr JP, Morgansen KA, Daniel TL, Cowan NJ. 2013. Flexible strategies for flight control: an active role for the abdomen. *J. Exp. Biol.* 216:1523–36
116. Roth E, Reiser MB, Dickinson MH, Cowan NJ. 2012. A task-level model for optomotor yaw regulation in *Drosophila melanogaster*: a frequency-domain system identification approach. In *2012 IEEE 51st IEEE Conference on Decision and Control*, pp. 3721–26. Piscataway, NJ: IEEE
117. Sutton EE, Demir A, Stamper SA, Fortune ES, Cowan NJ. 2016. Dynamic modulation of visual and electrosensory gains for locomotor control. *J. R. Soc. Interface* 13:20160057
118. Sponberg S, Dyhr JP, Hall RW, Daniel TL. 2015. Luminance-dependent visual processing enables moth flight in low light. *Science* 348:1245–48
119. Stamper SA, Madhav MS, Cowan NJ, Fortune ES. 2012. Beyond the jamming avoidance response: Weakly electric fish respond to the envelope of social electrosensory signals. *J. Exp. Biol.* 215:4196–207
120. Sepulchre R, Drion G, Franci A. 2019. Control across scales by positive and negative feedback. *Annu. Rev. Control Robot. Auton. Syst.* 2:89–113
121. Anderson PW. 1972. More is different. *Science* 177:393–96
122. Krakauer JW, Ghazanfar AA, Gomez-Marín A, MacIver MA, Poeppel D. 2017. Neuroscience needs behavior: correcting a reductionist bias. *Neuron* 93:480–90
123. Cowan NJ, Fortune ES. 2007. The critical role of locomotion mechanics in decoding sensory systems. *J. Neurosci.* 27:1123–8
124. Amari S. 1977. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* 27:77–87
125. Amit DJ, Tsodyks M. 1991. Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate—spikes, rates and neuronal gain. *Netw. Comput. Neural. Syst.* 2:259–73
126. Samsonovich A, McNaughton BL. 1997. Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* 17:5900–20
127. Burak Y, Fiete IR. 2009. Accurate path integration in continuous attractor network models of grid cells. *PLOS Comput. Biol.* 5:e1000291
128. Megías M, Emri Z, Freund T, Gulyás A. 2001. Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience* 102:527–40
129. Metzner W. 1993. The jamming avoidance response in *Eigenmannia* is controlled by two separate motor pathways. *J. Neurosci.* 13:1862–78

130. Kirschbaum F. 1983. Myogenic electric organ precedes the neurogenic organ in apteronotid fish. *Naturwissenschaften* 70:205–7
131. Dickinson MH, Muijres FT. 2016. The aerodynamics and control of free flight manoeuvres in *Drosophila*. *Philos. Trans. R. Soc. B* 371:20150388
132. Zhang C, Hedrick TL, Mittal R. 2018. An integrated study of the aeromechanics of hovering flight in perturbed flows. *ALAA J.* 57:3753–64
133. Sane SP, Dieudonne A, Willis Ma, Daniel TL. 2007. Antennal mechanosensors mediate flight control in moths. *Science* 315:863–66
134. Bhushan N, Shadmehr R. 1999. Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern.* 81:39–60
135. Kiemel T, Zhang Y, Jeka JJ. 2011. Identification of neural feedback for upright stance in humans: stabilization rather than sway minimization. *J. Neurosci.* 31:15144–53
136. van der Kooij H, van Asseldonk E, van der Helm FCT. 2005. Comparison of different methods to identify and quantify balance control. *J. Neurosci. Methods* 145:175–203
137. Theobald JC, Ringach DL, Frye MA. 2010. Dynamics of optomotor responses in *Drosophila* to perturbations in optic flow. *J. Exp. Biol.* 213:1366–75
138. Madhav MS, Stamper SA, Fortune ES, Cowan NJ. 2013. Closed-loop stabilization of the jamming avoidance response reveals its locally unstable and globally nonlinear dynamics. *J. Exp. Biol.* 216:4272–84
139. Fussmann GF, Ellner SP, Shertzer KW, Hairston NG Jr. 2000. Crossing the Hopf bifurcation in a live predator-prey system. *Science* 290:1358–60
140. O'Connor SM, Donelan JM. 2012. Fast visual prediction and slow optimization of preferred walking speed. *J. Neurophysiol.* 107:2549–59
141. Reiser MB, Dickinson MH. 2008. A modular display system for insect behavioral neuroscience. *J. Neurosci. Methods* 167:127–39
142. Maimon G, Straw AD, Dickinson MH. 2008. A simple vision-based algorithm for decision making in flying *Drosophila*. *Curr. Biol.* 18:464–70
143. Reiser MB, Dickinson MH. 2010. *Drosophila* fly straight by fixating objects in the face of expanding optic flow. *J. Exp. Biol.* 213:1771–81
144. Beatus T, Guckenheimer JM, Cohen I. 2015. Controlling roll perturbations in fruit flies. *J. R. Soc. Interface* 12:20150075
145. Seelig JD, Chiappe ME, Lott GK, Dutta A, Osborne JE, et al. 2010. Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nat. Methods* 7:535–40
146. Seelig JD, Jayaraman V. 2015. Neural dynamics for landmark orientation and angular path integration. *Nature* 521:186–91
147. Green J, Adachi A, Shah KK, Hirokawa JD, Magani PS, Maimon G. 2017. A neural circuit architecture for angular integration in *Drosophila*. *Nature* 546:101–6
148. Heiligenberg W. 1991. The jamming avoidance response of the electric fish, *Eigenmannia*: computational rules and their neuronal implementation. *Semin. Neurosci.* 3:3–18
149. Zimmermann JB, Jackson A. 2014. Closed-loop control of spinal cord stimulation to restore hand function after paralysis. *Front. Neurosci.* 8:87
150. Srinivasan SS, Maimon BE, Diaz M, Song H, Herr HM. 2018. Closed-loop functional optogenetic stimulation. *Nat. Commun.* 9:5303
151. Rutishauser U, Kotowicz A, Laurent G. 2013. A method for closed-loop presentation of sensory stimuli conditional on the internal brain-state of awake animals. *J. Neurosci. Methods* 215:139–55
152. Schiff SJ. 2012. *Neural Control Engineering: The Emerging Intersection Between Control Theory and Neuroscience*. Cambridge, MA: MIT Press
153. Ljung L. 1978. Convergence analysis of parametric identification methods. *IEEE Trans. Autom. Control* 23:770–83
154. Marmarelis PZ, Marmarelis VZ. 1978. The white-noise method in system identification. In *Analysis of Physiological Systems*, pp. 131–80. Boston: Springer
155. Nickl RW, Ankarali MM, Cowan NJ. 2019. Complementary spatial and timing control in rhythmic arm movements. *J. Neurophysiol.* 121:1543–60

156. Ljung L. 2006. Frequency domain versus time domain methods in system identification—revisited. In *Control of Uncertain Systems: Modelling, Approximation, and Design*, ed. BA Francis, MC Smith, JC Willems, pp. 277–91. Berlin: Springer
157. Aström K, Eykhoff P. 1971. System identification—a survey. *Automatica* 7:123–62
158. Aguilar J, Zhang T, Qian F, Kingsbury M, McInroe B, et al. 2016. A review on locomotion robotics: the study of movement at the intersection of robotics, soft matter and dynamical systems. *Rep. Prog. Phys.* 79:110001
159. Demir A, Samson EW, Cowan NJ. 2010. A tunable physical model of arthropod antennae. In *2010 IEEE International Conference on Robotics and Automation*, pp. 3793–98. Piscataway, NJ: IEEE
160. Holmes P, Full RJ, Koditschek D, Guckenheimer J. 2006. The dynamics of legged locomotion: models, analyses, and challenges. *SLAM Rev.* 48:207–304
161. Pellis SM, Bell HC. 2011. Closing the circle between perceptions and behavior: a cybernetic view of behavior and its consequences for studying motivation and development. *Dev. Cogn. Neurosci.* 1:404–13
162. Marken RS, Mansell W, Khatib Z. 2013. Motor control as the control of perception. *Percept. Motor Skills* 117:236–47



Contents

Robotic Self-Replication <i>Matthew S. Moses and Gregory S. Chirikjian</i>	1
Robots That Use Language <i>Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek</i>	25
Magnetic Methods in Robotics <i>Jake J. Abbott, Eric Diller, and Andrew J. Petruska</i>	57
Mobile Sensor Networks and Control: Adaptive Sampling of Spatiotemporal Processes <i>Derek A. Paley and Artur Wolek</i>	91
Network Effects on the Robustness of Dynamic Systems <i>Ketan Savla, Jeff S. Shamma, and Munther A. Dableh</i>	115
Routing on Traffic Networks Incorporating Past Memory up to Real-Time Information on the Network State <i>Alexander Keimer and Alexandre Bayen</i>	151
Amphibious and Sprawling Locomotion: From Biology to Robotics and Back <i>Auke J. Ijspeert</i>	173
Stochastic Dynamical Modeling of Turbulent Flows <i>A. Zare, T.T. Georgiou, and M.R. Jovanović</i>	195
Robotics In Vivo: A Perspective on Human–Robot Interaction in Surgical Robotics <i>Alaa Eldin Abdelaal, Prateek Mathur, and Septimiu E. Salcudean</i>	221
The Synergy Between Neuroscience and Control Theory: The Nervous System as Inspiration for Hard Control Challenges <i>Manu S. Madhav and Noah J. Cowan</i>	243
Learning-Based Model Predictive Control: Toward Safe Learning in Control <i>Lukas Hewing, Kim P. Wabersich, Marcel Menner, and Melanie N. Zeilinger</i>	269

Recent Advances in Robot Learning from Demonstration <i>Harish Ravichandar, Athanasios S. Polydoros, Sonia Chernova, and Aude Billard</i>	297
Recent Scalability Improvements for Semidefinite Programming with Applications in Machine Learning, Control, and Robotics <i>Anirudha Majumdar, Georgina Hall, and Amir Ali Abmadi</i>	331
The Inerter: A Retrospective <i>Malcolm C. Smith</i>	361
Port-Hamiltonian Modeling for Control <i>Arjan van der Schaft</i>	393
Automated Planning for Robotics <i>Erez Karpas and Daniele Magazzeni</i>	417
Scientific and Technological Challenges in RoboCup <i>Minoru Asada and Oskar von Stryk</i>	441

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>