

Learning Outside the Brain: Integrating Cognitive Science and Systems Biology

This article uses concepts ranging from cognitive science to control theory to discuss how learning in biological systems is possible outside of, or even without, a brain.

By JEREMY GUNAWARDENA 

ABSTRACT | Learning is commonplace in organisms such as ourselves and even in organisms as far distant as the bee and the octopus. Such learning is implemented by brains, or neuronal networks, and has been extensively studied within ethology, psychology, cognitive science, and neuroscience. Whether learning also takes place in nonneuronal settings has remained a matter of sustained controversy, too often dominated by ideological views. In this survey, I will explain how learning can be rigorously interpreted as a form of information processing and then explore the evidence for whether learning also takes place in organismal contexts outside the brain, such as physiology, development, and individual cells. I will try to explain why it is important to build bridges in this way between cognitive science and systems biology, why concepts and methods from various branches of engineering may be helpful in this task, and what the eventual impact may be on how we think about the organism.

KEYWORDS | Active inference; Developmental Origins of Health and Disease (DOHaD); homeostasis; internal models; mind as model; mutual information.

I. INTRODUCTION

Learning may be informally defined, for now, as a persistent change in behavior in response to the same stimulus [1]. Humans excel in learning new skills, such as languages and sports, while other animals, such as dogs, birds, bees, and octopuses, have their own remarkable capabilities [2]. There are many forms of learning, some

of which have been characterized as the basis for scientific studies. Elementary forms of nonassociative learning in response to a single stimulus include habituation [3] or its counterpart, sensitization [see Fig. 1(a) and (b)]. More complex forms of conditioning involve the learning of associations between multiple stimuli. Ivan Pavlov (1849–1936) discovered early in the 20th century how a dog could learn to salivate [the conditioned response (CR)] to the ringing of a bell [the conditioned stimulus (CS)], factors normally unrelated to each other, by repeatedly associating the bell with the presentation of food [the unconditioned stimulus (US)], which prompts salivation [the unconditioned response (UR)] as part of the digestive process Fig. 1(c.1). This kind of classical conditioning often deals with involuntary reflexes, such as salivation, which are not under the organism’s intentional control. Somewhat later, building on the work of Edward Thorndike (1874–1949) and Clark Hull (1884–1952), Burrhus Frederick Skinner (1904–1990) elaborated the carrot and stick approach that generations of parents had figured out for themselves into instrumental, or operant, conditioning in which intentional, goal-directed behaviors are modified by positive or negative reinforcement [see Fig. 1(d)]. A YouTube video of Skinner’s pigeons playing ping pong¹ gives a vivid sense of the generality and effectiveness of instrumental conditioning.

These learning concepts originate in behaviorist psychology [4]. Learning has been reconsidered within the cognitive science that replaced behaviorism (see Section II), and I will suggest in Section VI a definition of learning that may be more appropriate for systems biology.

Learning is a central topic in ethology, psychology, cognitive science, and neuroscience. In these disciplines, it is assumed to be implemented by minds, brains, or

Manuscript received August 16, 2021; revised January 21, 2022; accepted March 17, 2022. Date of publication April 28, 2022; date of current version May 19, 2022.

The author is with the Department of Systems Biology, Harvard Medical School, Boston, MA 02115 USA (e-mail: jeremy@hms.harvard.edu).

Digital Object Identifier 10.1109/JPROC.2022.3162791

¹<https://www.youtube.com/watch?v=vGazyH6fQQ4>

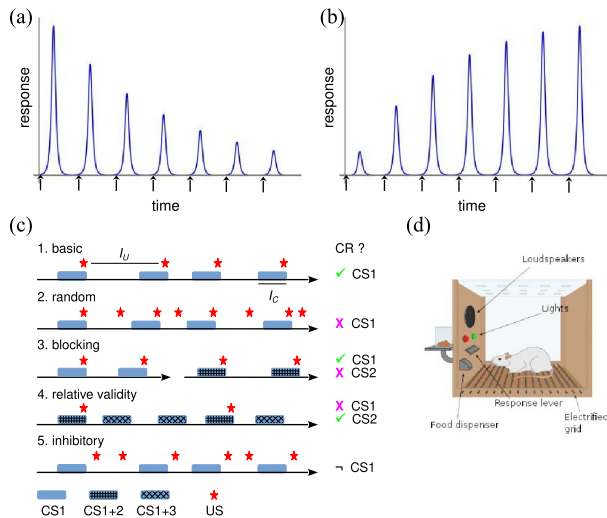


Fig. 1. Forms of learning. (a) *Habituation. Hypothetical illustration, showing diminishing responses to repeated stimuli (arrows). For additional properties, see [26].* (b) *Sensitization. The opposite effect to habituation, in which the response increases.* (c) *Classical conditioning, adapted from [27, Box 1]. The right-hand column shows whether or not a CR develops to the indicated CS. Protocol 1 is the traditional Pavlovian one in which the CS (CS1, blue bar), such as a tone or a light, is paired with the US (red star), shown here at a fixed delay from CS onset. I_U and I_C denote the US-US interval and the CS duration, respectively. In protocol 2, the US are presented at random with equal probabilities of occurring with or without the CS, and a CR does not develop. Protocols 3 and 4 are cue competition schemes involving multiple CSs. Protocol 3 has two phases: in the first phase, a CR is established to CS1; in the second phase, the compound stimulus CS1+2 is presented with the US, but a CR to CS2 does not develop. In protocol 4, CS1 is always presented as a compound with either CS2 or CS3, and the US occurs only for CS1+2. A CR develops to CS2 presented in isolation but not to CS1. In protocol 5, the CS is never paired with the US, but an inhibitory association still arises (see text).* (d) *Instrumental conditioning, showing a Skinner box, in which an organism can be given pleasant (food) or unpleasant (shock) reinforcement, depending on its responses (lever) to various stimuli (light and tone); illustration taken from the Wikipedia entry for Operant Conditioning Chamber under the CC BY-SA 3.0 license.*

neuronal networks, depending on the setting. The question that I want to discuss here is whether learning is also relevant in other biological contexts, such as physiology or development, in which the brain is not the only organ that is present, or in individual cells, in which there are no neuronal aspects at all. Our focus will largely be on mammals like us, with the fascinations of invertebrates, plants, amoeba, and prokaryotes [5]–[8] largely falling outside the scope of this limited survey. The questions of how cognition evolved and how it might be embodied have also elicited much recent interest [2], [9]–[12], and the “cognitive lens” offers an inclusive perspective on many aspects of biology [13]. However, cognition is a more elusive concept than learning. The latter has both experimental and theoretical rigor, as we will see, and more than enough subtlety to keep us busy.

But why should this question of learning be of interest to readers of the IEEE Proceedings and to systems biologists like me, whose day job is studying molecular information processing?

First, and perhaps most significantly, biological learning, as described above, can now be seen as a form of information processing, in which the representation within the organism of the information that has been acquired forms a model of the environment, from which predictions are made about future events (see Section II). The idea of mind as model can be traced back to the philosophy of Kant and to Helmholtz’s work on perception [14], and it has acquired broad foundations in modern cognitive science (see Section III-E). This information processing formulation abstracts learning from any specific biological implementation and suggests how it could be relevant to many different biological contexts or, indeed, to machines. The interplay between learning by organisms and learning by machines is long-standing. The deep learning in neural networks that have been so much in the news of late has its roots in neuroscience [15]. Nowadays, we have come full circle; computational learning, information theory, and control theory have become important tools for analyzing biological learning.

Second, learning provides a unifying idea, which links disparate areas and scales of biology, which are not directly concerned with minds, brains, or neurons. It offers a conceptual framework and an experimental paradigm for interpreting biological phenomena and understanding the organism in a different way. At least, that is one of the messages that I hope you will take away.

Unifying ideas are a feature of systems biology although opinions differ as to what this subject is about. It is often equated with “-omics” and somewhat less often with “modeling.” However, these are methods. What is the question? I believe that the question is asking how we get from dead molecules to living organisms. In the wake of the genome projects, we think that we know most of the molecules, so the important problem is no longer characterizing the molecules—which was molecular biology—but asking how the molecules collectively give rise to the phenomena of life. This requires thinking differently and asking different kinds of questions, rather than just using your favorite hammer. Whether learning takes place elsewhere than brains is the kind of different question, which may help us better perceive how living systems emerge out of molecules. The machine metaphor has been the key tactic in trying to bridge this divide [16], [17], but the classical machine that transduces energy or matter is not an agent that changes itself by constructing representations of its environment. Learning equips a machine with agency in the form of goals and purposes and thereby captures an essential feature of life.

Finally, there is a third reason why learning is so interesting. It offers a way to study nature, or the capability to learn, and nurture, or the environment that is learned, not in opposition to each other, as is all too often the case, but

as an integrated process through which organisms function in the world. We can thereby leave behind us one of the least productive arguments in biology.

My own route to thinking about learning goes back, in part, to a lecture by Dennis Bray, in which he described the extraordinary avoidance behaviors reported by Herbert Spencer Jennings (1868–1947) in a single-cell protist, the ciliate *Stentor roeseli* [18], [19]. Upon repeated irritation, *Stentor* works through a hierarchy of increasingly vigorous avoidance reactions before eventually detaching from its holdfast and swimming away. This may be seen as a form of nonassociative learning in which behavior changes in response to a single stimulus but in a more complex way than habituation or sensitization: the organism does something different rather than the same thing in a different way. Dennis and I were not the only ones to be captivated by such complex learning in a single cell [20], but I was hugely disappointed to be told by the experts that Jennings' experiments could not be reproduced [21]. A closer reading, however, suggested that it was the attempt at reproduction that was experimentally flawed, not Jennings' original work. That was my first exposure to the sociology of science and the particular problems surrounding learning [22, Chapter 1]. In the case of Jennings' work, it seemed that the scientific community did not want to believe that individual cells were capable of complex learning and preferred to trust a shoddy experiment that said the right thing to many good experiments that said the wrong thing. Ideology trumped evidence. This was troubling, to put it mildly—science is not supposed to work like politics—and it sparked an unofficial, decade-long, skunk-works project, which eventually confirmed that Jennings had been right all along [23]. This drew me into discussions with cognitive scientists about the broader aspects of learning [24], and these conversations have provided much of the background for this article.

If the close encounter with sociology was not warning enough, there are other minefields in the way of chasing unifying ideas across different disciplines. We perpetuate the myth, to the nonscientific public, that science is a unified approach to understanding the world. The reality, as we all know, is rather different. Even within biology, different disciplines occupy different conceptual landscapes, use different experimental paradigms, talk different languages, often with the same words, and impose stringent border controls on foreigners. Nobody, least of all me, is an expert in all the fields across which I want to trespass. It is particularly challenging for the outsider to adjudicate arguments within disciplines and the topic of learning seems to encourage some especially ferocious disputes. The author Philip Ziegler, in his study of the Black Death, spoke of “*the spectacle of rival historians, each established in his fortress of specialized knowledge, waiting to destroy the unwary trespasser*” [25]. Such a terrifying prospect is usually sufficient to keep us in our comfort zones. However, perhaps, once in a while, it is important to take the risk and do as Ziegler did, if only to see what happens and have

some fun. I am sure that this will invite bombardment from various fortresses, but I hope that it will also encourage others to correct my mistakes and press forward with the exploration. I must thank Mustafa Khamash and Jörg Stelling for providing this opportunity. What follows is really not their fault.

II. LEARNING AS INFORMATION PROCESSING

A. Learning Is Not Association

The relationship between learning and the forming of associations is long-standing and apparently well established. It seems intuitively plausible that it is the association between a CS and a US, or between an action and its reinforcement, arising from their close pairing in time [see Fig. 1(c.1)], which underlies learning during conditioning. This intuition about the cognitive aspect of learning has uncanny parallels with what neuroscientists have discovered about learning-dependent changes in neuronal networks. The synapses that link neurons are strengthened by long-term potentiation when the firing of an upstream neuron occurs slightly before the firing of a downstream neuron. As Donald Hebb (1904–1985) presciently suggested in the learning rule named after him, “*neurons that fire together, wire together*” [28], [29]. (The modern understanding of spike-timing-dependent plasticity involves both strengthening and weakening of a synapse depending on which neuron fires first [30].) For neuroscientists, learning, and the memories that accompany it, are implemented by changes in the strength of synapses [31]. The cognitive and neuroscientific accounts of learning both focus on associations, between stimuli and between neurons, respectively, and have strongly mutually reinforced each other.

Experiments carried out in the late 1960s showed, however, that what is learned during conditioning is different from an association. Rescorla [32] showed that the extent of conditioning depends on the relative frequencies, or probabilities, of the US occurring in the presence or absence of the CS. When these frequencies are equal, as in Rescorla's truly random protocol [see Fig. 1(c.2)], no learning takes place, despite the continuing association between CS and US. Cue competition studies with multiple stimuli revealed other unexpected properties. Kamin [33] showed with the blocking protocol [see Fig. 1(c.3)] that an association may not form even though the CS and US are repeatedly presented together. Wagner *et al.* [34] showed that organisms select the stimulus-response to learn depending on which stimulus gives a better prediction of the US [see Fig. 1(c.4)]. Finally, in the inhibitory protocol [see Fig. 1(c.5)], Rescorla [35] considered the extreme of relative frequencies in which the US occurs only in the absence of the CS and found that it then takes longer than normal to acquire the CR to that CS during subsequent training. An inhibitory association is formed despite the absence of pairing between CS and US.

In summary, association in time is neither necessary nor sufficient for learning, which appears to be guided instead by the information, which is most predictive of the US. Learning occurs when the organism is “surprised” by the failure of its predictions [33]. The idea of prediction error will ramify through the discussions that follow.

Further sophistication arises in timing, on which the presumed association between CS and US is based. It was generally supposed that, if the time from CS onset to US onset was short enough, then the CS was anticipating the US, and learning would take place. However, how short is short? If trials are conducted following Fig. 1(c.1) with the US-US interval, I_U , and the CS duration, I_C , being varied, then data from many studies show that the number of trials to the acquisition of a CR, which is a measure of the slowness of learning, varies approximately inversely with the ratio I_U/I_C on a log-log scale [36, Fig. 9]. In particular, there is no “short enough” CS-US interval for learning to occur: the interval can be long, provided that the interval between USs is increased in proportion.

This property of timescale invariance [36] is in marked contrast to the association between neurons in Hebbian learning. The latter exhibits an absolute threshold in timing: for spike timing-dependent plasticity, the firing of the presynaptic and postsynaptic neurons must occur within a window of some tens of milliseconds, depending on the neuronal context [30]. The relationship between the cognitive and neuroscientific accounts of learning based on the association between stimuli and between neurons, respectively, becomes much less clear in the light of timescale invariance [37]. Synaptic plasticity remains the conventional mechanism for explaining learning and memory in neuroscience [31], but its adequacy for the computational requirements of mental processing has been vigorously questioned [37], [38].

There is also more to timing than timescale invariance. It was observed early on, in both classical and instrumental paradigms, that the organism learns not just to do something but also when to do it [36]. This temporal aspect of learning is rarely emphasized, in contrast to the forming of associations [39], but many studies have confirmed the details. For example, during eyeblink conditioning in the rabbit, a puff of air to the eye (US) induces a blink (UR) with a tone or light as the CS. The rabbit does not just learn to blink (CR) but does so when the air puff would have been delivered [40]. Similar results have been found for fear conditioning in the rat [41]. These results show that the organism is learning the time interval I_C in Fig. 1(c.1), and it exhibits CRs in response to CSs at the time when it is expecting the US to be delivered.

These findings demonstrate that conditioning is more sophisticated than the mere formation of associations. Organisms appear to be extracting information about time series from the flow of events, assessing which series are predictive of others and when these predicted events will occur. They are undertaking computations to form an internal model, or “temporal map,” of the experienced

stimuli, which is being continually updated as new events are encountered [27]. The physiological nature of these internal models remains unclear, but abstract mathematical models have been put forward to account for the experimental findings. (Note the two senses in which “model” is being used here: one meaning a physiologically implemented internal representation of the external world, and the other meaning a mathematical construct developed by us. Regrettably, the semantic overloading of “model” will become even worse in Section II-C.) The mathematical model developed by Rescorla and Wagner [42] in the wake of their experimental findings was an influential starting point, but the later introduction of reinforcement learning theory by Sutton and Barto [43] marked a key turning point, which brought together mathematical psychology, optimization theory, and artificial intelligence.

B. Theories of Learning

We briefly describe the main ideas, largely following the treatment in [44], which provides more background. Consider a finite set of conditioned stimuli, S_i , presented to an agent, who tries to learn the value, or predictive strength, $V(S_i)$, of S_i , with respect to the single US. Rescorla and Wagner [42] introduced two central ideas. First, V is updated on each trial in proportion to the error between its current value and the ultimate, asymptotic value, λ , corresponding to the US. Second, the value of a compound stimulus is the sum of the component values. This leads to the following update rule applied synchronously to all stimuli:

$$V_{\text{new}}(S_i) = V_{\text{old}}(S_i) + \eta \left[\lambda - \sum_j V_{\text{old}}(S_j) \right]. \quad (1)$$

Here, $0 \leq \eta \leq 1$ is the learning rate, which may depend on S_i , and the sum on the right-hand side is taken over those stimuli presented in the compound on that trial. Equation (1) explains many of the cue competition experiments described above [42], although not the timescale invariance, to which we return below.

Sutton and Barto [43], [45] and Niv [44] expanded on Rescorla and Wagner’s ideas in several ways, within the setting of Markov decision processes. For a more detailed comparison with Rescorla–Wagner, see [46]. First, Sutton and Barto formalized the statistical relationships between stimuli as arising from a discrete-time Markov chain, X , specified by a conditional probability distribution over the stimuli, $\Pr(X_{n+1} = S_i | X_n = S_j)$. Second, they assumed that the agent is stochastically presented with a reward, r , in each state, through a conditional probability distribution, $\Pr(r | X)$, and that the value, $V^*(S_i)$, which the agent attempts to learn, is the total expected reward discounted at a rate $0 < \gamma < 1$

$$V^*(S_i) = \mathbf{E} [r(X_0) + \gamma r(X_1) + \gamma^2 r(X_2) + \dots | X_0 = S_i] \quad (2)$$

where $r(X_n)$ is the reward received after n steps of the Markov chain and the expectation is taken over both chain and reward distributions. The reward takes the place of the US and its accompanying UR and makes it easier to build a connection between Pavlovian and instrumental conditioning. Discounting ensures that $V^*(S_i)$ is finite. Animals, including humans, do assume that rewards diminish the later that they are received although, interestingly, they discount more nearly hyperbolically than exponentially [47]. Crucially, (2) places this approach in the setting of optimal control and dynamic programming [48]. Third, Sutton and Barto introduced the temporal difference prediction error

$$\delta_n(V) = r(X_n) + \gamma V(S(X_n)) - V(S(X_{n-1})) \quad (3)$$

where $S(X_n)$ is the stimulus received after n steps of the Markov chain and V is the current value function for stimuli. Note that $\delta_n(V)$ in (3) depends only on what is presented to the agent and requires no knowledge of the probability distributions for stimuli or rewards, which the agent is not expected to know. The update rule for the value function then follows the same form as (1)

$$V_{\text{new}}(S(X_n)) = V_{\text{old}}(S(X_n)) + \eta \delta_n(V_{\text{old}}). \quad (4)$$

One can appeal to results from dynamic programming to show that, under appropriate conditions, (4) converges to the true value function V^* defined by (2). Hence, temporal difference learning leads to the expected total reward, despite the stochastic dynamics of stimuli and rewards being unknown.

A further step may be taken by introducing actions that the agent can undertake in any state, which may change the probability distributions of stimuli and rewards [43], [44]. The temporal difference prediction error in (3) can be used to incrementally learn both the value function and a policy, which specifies the probability with which action should be taken in a given state. This formulation not only accommodates instrumental conditioning but makes more explicit the connections between learning and control, to which we will return in Section III-B. The agent seeks an optimal reward by iteratively learning how to control, through its own actions, the environment that stimulates it. The key problem, as expressed in (3), lies in balancing short-term exploitation of reward, or immediate control, with long-term exploration, or system identification.

The difference between the Rescorla–Wagner model and the reinforcement learning model is more profound than one of mathematical elaboration. Rescorla and Wagner proposed a phenomenological model to account for empirical data; Sutton and Barto proposed a normative model, based on rational principles, which claims that learning is a form of reward optimization. (A normative description of Rescorla–Wagner may be given under

suitable assumptions [46].) At first, the success of reinforcement learning lay in solving challenging sequential decision problems, such as backgammon [49], but then its significance for biological learning was revealed in spectacular fashion [50]. Studies of learning had progressed from the kinds of behavioral experiments described in Fig. 1 to recordings from electrodes implanted in particular brain areas of awake animals during learning tasks [51]. In experiments with monkeys, phasic bursts of activity by dopamine-releasing neurons in the midbrain ventral tegmental area were found to follow the temporal difference prediction error (3) for that task [50], [52]. The error signal shifts in a characteristic way over time from the first receipt of a reward to the stimulus that increasingly predicts the reward, only to return with the opposite sign if the reward does not appear as predicted [44, Fig. 2]. This pattern of activity is exactly what was observed in dopaminergic neurons [50, Fig. 1]. Similar results have been found for more complex, higher order learning tasks in other animals, including humans [53]. It is hard to imagine a more striking example of a mathematical theory explaining complex biological data and, thereby, prompting a paradigm shift in biological thinking. The reward prediction error hypothesis for dopamine signaling has become a centerpiece of modern neuroscience [54]–[56].

Reinforcement learning has also had another spectacular success, in making the leap from learning backgammon to teaching humans how to play Go [57], a game long thought to require uniquely human capabilities at its highest levels. This may suggest that it has swept all before it. However, the Markov decision process setting described above uses discrete time, in terms of which order and sequence can be specified but not time intervals. While it accounts well for many complex aspects of conditioning, it is not manifestly timescale invariant. Gallistel and colleagues have developed an alternative, largely normative account of Pavlovian and instrumental conditioning, based on rates of occurrence and information theory [27], [58], [59]. They consider stimuli as stochastic processes impinging on an agent and ascribe the predictive strength, or degree of contingency, of two processes, to the mutual information between them, suitably normalized. If the agent is to estimate such quantities, the problem of temporal discretization still has to be confronted and is addressed empirically using Weber’s principle that the minimal discernible difference in a measure, including a measure of a time interval, is proportional to its value [59]. The resulting theory is intuitively attractive and readily seen to be timescale invariant. It also accounts for many complex aspects of conditioning. It lacks, as yet, the neurophysiological basis of temporal difference error prediction, but the focus on mutual information will be important to us later (see Section VI).

The subtleties of disentangling association from temporal contingency, the impact of the Rescorla–Wagner model, and the continuing misunderstandings around it are carefully explored in Gallistel’s tribute to his erstwhile

colleague, Robert Rescorla, in a special issue devoted to the latter's memory [60].

The information-theoretic approach draws attention to the probability distributions of occurrence times of stimuli but focuses on quantitative measures of these distributions, such as their entropy. The Bayesian revolution in cognitive science [61]–[63] has taken the radical step of presuming that an agent learns the actual probability distributions of values, or suitable approximations to them, as representations of the agent's beliefs—the agent tracks its own uncertainty—and these distributions are updated by Bayes rule upon receipt of new stimuli [64]. Further discussion of the Bayesian brain hypothesis would take us too far afield, but Gershman [46] has proposed a unification of the Bayesian and reinforcement learning approaches, and we will encounter the Bayesian viewpoint again in Section III-E.

C. Reflexive Versus Reflective Models

The approaches described above share the common feature that, whatever internal model of the environment is adopted by an agent, whether based on predictive values, entropies, or distributions, it is used to immediately determine the agent's actions. An additional possibility, first articulated by Edward Chase Tolman (1886–1959) [65], [66], is that an internal model is used to explore, or simulate, the world in advance of action and, thereby, plan better actions for the future. Tolman [65] noted that rats left to explore a maze without reinforcement from food undertake latent learning, which can be revealed by significantly faster explicit learning under reinforcement, compared to rats encountering the maze for the first time. In Tolman's view, the free-ranging rats were latently constructing a cognitive map of the maze and using this internal model to make better predictions of what routes to take when there was food waiting for them. At choice points in the maze, rats are seen to hesitate and look back and forth as if mentally working through their options, a behavior interpreted as vicarious learning [66]. Tolman's prescient suggestion of a cognitive map underlying spatial learning was later given a neurophysiological basis through the remarkable discoveries of place cells and grid cells in the hippocampal region of the brain [67], [68]. Such cells do, indeed, appear to transiently map out the potential forward paths when rats pause at a choice point [69].

The idea that agents use two types of models appears repeatedly in cognitive science from different perspectives with different terminologies and interpretations [66]: retrospective versus prospective, reflexive versus reflective, habitual versus goal-directed, and model-free versus model-based. (This last usage, which comes from reinforcement learning, uses model in yet another sense to mean the model of the environment given by the probability distributions for stimuli and rewards of the Markov decision process. Temporal difference learning based on (3) is called model-free because it does not use

this environmental model. For us, the continually updated stimulus values still provide a model of the environment. As noted above, the word “model” has become dangerously overloaded. We will do our best to specify the intended meaning when the context does not make it clear.) An informal illustration of the dichotomy comes from Daniel Kahneman [70], who speaks of fast versus slow thinking. A bat and a ball together cost \$1.10; the bat costs \$1.00 more than the ball; and how much does the ball cost? The faster answer—given by most people, including me—says one thing; the slower answer based on actually solving the linear equations says another.

The different terminologies introduced above illustrate different aspects of the implied tradeoffs: one type of model relies on faster, ingrained habits for immediate purposes, requires fewer computational resources but can make more errors; the other type of model is slower, consumes more computational resources to reflect about potential future scenarios, plans accordingly, and may be more accurate. The latter model implies greater agency on the part of the organism, which simulates the world in advance of dealing with it, rather than relying on pre-existing, perhaps even genetically acquired, habits. Note that retrospective, habitual, or reflexive models may still be predictive—after learning, the CS in Pavlovian conditioning is predicting the US—but there is no prospective planning involved. Animals appear to use a balance of models, rather than a dichotomy. It is a commonplace observation that repeated training gradually shifts the balance from reflective to reflexive behavior; it would be hard to learn to play the piano or tennis otherwise. Nevertheless, the distinction will be helpful to keep in mind, and we will adopt the reflexive-reflective terminology as being the least misleading for our purposes.

In computational terms, both types of models can be formulated from a reinforcement learning perspective. In place of the temporal difference error in (3) being used to guide the immediate choice of action in a reflexive manner, the agent learns some approximation to the decision tree of actions and rewards and uses that to prospectively plan what action to take in a reflective manner. Multistage Markov decision processes can distinguish between the two kinds of models and experiments with human subjects, and fMRI data show both models being used and grounded in different brain areas [66], [71].

D. Summary

This limited introduction can hardly do justice to the topic of learning, but the citations should offer some further orientation. It has shown, I hope, that learning is not just the formation of associations; it is a form of information processing that agents, both living and artificial, can implement; it involves, to varying degrees, the construction of representations, or internal models, of an agent's environment; such models may be reflective or reflexive, depending on whether or not they are used to

simulate the future; a rich experimental repertoire exists to analyze different forms of learning; this repertoire is accompanied by sophisticated theoretical frameworks that draw on mathematics, engineering, and computer science; and in respect of learning by brains or neuronal networks, learning is studied in a strongly interdisciplinary manner across ethology, psychology, cognitive science, and neuroscience. Let us turn now to learning outside the immediate context of the brain, in physiology (see Section III), development (see Section IV), and individual cells (see Section V). Having done that, we will offer a definition of learning that seems more appropriate for systems biology (see Section VI).

III. PHYSIOLOGY, VISCERAL LEARNING, AND ALLOSTASIS

A. From Psychology to Physiology

The separation of the study of the body and the study of the mind was perhaps inevitable in the baleful light of Cartesian dualism, but they came together for Ivan Pavlov, whose pioneering work in digestive physiology led to his Nobel Prize and prompted the studies of conditioned reflexes for which he is so much better known [72]. Subsequent attempts to bring the two disciplines closer have had mixed success, but the argument for doing so, and the potential biological impact, has much to do with learning.

Physiology takes us inward, from the central nervous system (CNS) and the focus on stimuli external to the organism to the peripheral nervous system and internal stimuli. The body's tissues are immersed in a fluid medium, the internal milieu, which rather resembles the sea water in which the first cells lived. It was Claude Bernard (1813–1878) who first articulated the idea that the constancy of the internal milieu is one of the conditions of animal life [73]. He pointed out that this constancy arose not because the organism was impervious to change, but, on the contrary, because it sensed departures from constancy and actively compensated for them. Walter Cannon (1871–1945) clarified what we now call negative feedback as the key feature in Bernard's dynamical vision of stability and coined the word homeostasis to describe the latter [74], [75]. He thereby codified the central concept of physiology and the first systems concept in biology. Norbert Wiener (1894–1964) and Arturo Rosenblueth (1900–1970), inspired partly by Cannon and partly by World War II demands for automated artillery, drew the analogy between feedback control in animals and machines. These ideas became one of the founding themes of cybernetics [76]. The inspiring ambitions for this new discipline, ranging across the physical and social sciences [77], proved too grand to sustain but bioengineering, in the guise of control theory, domesticated the main ideas, and provided post-WWII physiology with a solid quantitative foundation [78].

Homeostasis covers the maintenance of the fluid internal milieu—pH, salt balance, O₂, CO₂, glucose, and so on—and also other quantities relevant to the functioning of the visceral organs, such as heart rate, breathing, blood pressure, and temperature. These are usually regulated below the level of consciousness by the autonomic part of the peripheral nervous system (the other part being the voluntary nervous system that controls the skeletal muscles) together with the system of endocrine hormones released from glands. The electrical and chemical systems of communication are closely linked: the hypothalamus regulates the pituitary gland at the base of the brain, which, in turn, chemically regulates multiple hormonal axes, such as the hypothalamus-pituitary-adrenal (HPA) stress axis, while the autonomic nervous system (ANS) innervates all glands.

B. Internal Models

The connection between control and learning was mentioned in Section II-B. At one level, it seems natural: to control a system, one must know something about it, and that knowledge must be represented within the controller. We can see this in the simplest biological control systems. Many of these exhibit zero steady-state error in response to a step perturbation [79]–[81], the hallmark of integral control under linear assumptions. The control variable, y , follows the equation:

$$\frac{dy(t)}{dt} = k(x(t) - x_{set})$$

which ensures that, at steady state, the controlled variable, x , must be at its set point, x_{set} . Hence, y is keeping track of—acquiring information about and remembering—an aggregated measure of the system's deviation from its set point

$$y(t) = k \int_{u=0}^{u=t} (x(u) - x_{set}) du. \quad (5)$$

The internal representation of the controlled system's history provided by y may be limited, but it is all that is needed.

Integral control is an instance of the internal models principle [82], [83], or, with some latitude in interpretation, of the much vaguer principle of requisite variety from cybernetics [84], [85]: the controller contains a model of the perturbations to which the system is exposed. In this case, the perturbations are step changes, the Laplace transform of a step is $1/s$, and so the controller contains the same element, which is also an integrator. This internal model of the perturbations is not the same as that provided by the controlled variable, y , of the system's history, as above, but they seem evidently related to each other.

Internal models and control theory have been central to understanding movement systems [86]–[88]. Humans

and many other animals exhibit remarkably complex, fast, and accurate motor activities—playing the piano or hitting a backhand in tennis—and their control cannot be based solely on visual feedback because it is too slow. Instead, internal models are used to plan the complex multidimensional trajectory of limb movements and issue the motor commands to achieve the desired trajectory. In the language of Section II-C, this is highly reflective, not merely reflexive like the integral controller. There is also a subtle difference between the kinds of internal models being discussed here, in terms of what they represent. Models arising from the internal models' principle represent the perturbations to which the system is exposed; the models found in movement control systems usually represent the systems themselves [89]. Biology may exploit both types of the internal model, depending on the predictability of the perturbations [90].

Particularly striking use of reflective internal models is found in the involuntary vestibular–ocular reflex (VOR). When you run over broken ground, the background is bouncing around in a capricious way, yet you perceive a stable image, unlike a video camera following the same course. The accelerometers in the vestibular apparatus of the inner ear send signals to the brainstem and cerebellum, which compute the effect of head movements on the visual field and issue motor commands to the muscles around the eye, which compensates for the changes [79], [91]. The sophistication of the VOR is startling, but evolution has had a long time to improve it, from at least as far back as our fish ancestors [92]. Indeed, the necessity to disentangle sensory effects caused by one's own movements from those caused externally must have been an evolutionary priority for any organism that actively moves. This will be something to keep in mind when we discuss learning in single cells (see Section V-E).

This discussion of internal models has strayed in two respects from the avowed topic of this section. We have been discussing movement, not homeostasis, and internal models themselves, rather than learning. The purpose was to show that internal models do exist, and indeed, that control theory tells us that they must exist as a consequence of homeostasis. Internal models of movement must, in any case, be learned—think again of playing the piano or tennis—and also continuously adjusted in the face of novel demands, fatigue, injury, and aging [93]. Even involuntary internal models are learned. Classic experiments by Erismann and Kohler [94] showed that human subjects equipped with prism glasses, which reverse the optical field, are highly visually disabled to begin with but recover their visual abilities and their VOR recalibrates to the reversed flow of the optical background over a period of weeks [95]. How such learning takes place remains unclear, but adaptive control theory may suggest potential theoretical models [96].

Two kinds of internal models have appeared in our discussion: those coming from control theory and those coming from learning theory. It seems implausible that

these are different, but there is no rigorous justification for why they are the same. The relationship between them remains to be clarified.

C. Visceral Learning

Let us turn now from movement back to the internal milieu. In doing so, we move from what is observable and controllable, where the CNS plays the dominant role, to what is inaccessible and elusive, and the province of the ANS. The CNS is evidently capable of complex learning. Following Pavlov's work on salivation, it seemed plausible that the ANS was capable of classical conditioning (there was substantial Russian work in this direction that was poorly understood in the West [97, Chapter 5]), but instrumental conditioning, which is both voluntary and broader in scope, was generally considered beyond the capabilities of the ANS. This consensus remained unshaken despite the discovery of instrumental conditioning in invertebrates [98].

Aside from this general prejudice against the ANS, it is not straightforward to design experiments on visceral learning. Homeostatic control is hierarchical and distributed across multiple tissues and organ systems, and it is challenging to control for the many potential confounding factors. Difficulties of reproducibility and interpretation have been conspicuous. The renowned psychologist Neil Miller (1909–2002) received much attention when he reported experiments showing instrumental conditioning of heart rate, intestinal contractions, and blood pressure in rats [99], but his own Ph.D. student, Barry Dworkin, was later unable to replicate these findings [100]. This crisis had two effects. Miller [101] continued more successfully with studies of instrumental learning in humans, becoming one of the founders of the field of biofeedback. As he liked to point out, the sphincters that control waste release are under autonomic control, and there seems no difficulty in young humans learning to control them through reinforcement. Whether toilet training was seen as compelling evidence by fellow scientists is not so clear.

As for Dworkin, he devoted himself to unambiguously demonstrating visceral learning, developing paralyzed animal preparations, which could be reliably studied over weeks and focusing on the baroreflex that homeostatically regulates blood pressure. He showed classical conditioning to both exteroceptive (auditory tone) and interoceptive (nerve stimulation) cues [102]. He found that activation of the baroreceptors reduced pain sensitivity, which could potentially lead to instrumental learning of high-blood pressure in human subjects [103], [104]. The implications of learning for clinical conditions, especially drug addiction, have become an important theme [105], [106]. Repeated administration of a drug often results in habituation [see Fig. 1(a)] or tolerance. However, this does not arise because the drug is increasingly ignored by the organism but, rather, because compensatory physiological mechanisms are triggered, which increasingly reduce the drug's

effects. Such compensatory responses can elicit Pavlovian conditioning from cues that accompany drug administration, such as the place or manner of administration. These effects can be revealed by administration of a placebo under the same conditions: instead of the usual response to the drug, an opposite compensatory effect occurs [97, Chapter 6]. Such withdrawal symptoms are often associated with learned cues [105]. The responses can be very subtle, as found with the conditioning of glucose levels in response to insulin administration, where both hypoglycemic and hyperglycemic responses are observed, depending on the details of the context [107].

Dworkin [97, Chapter 4] also developed the first formal model that integrated Rescorla–Wagner theory with feedback control. In his phenomenological and *ad hoc* formulation, feedback control was interpreted recursively to allow integration with (1), and the effect of conditioning was to modulate the gain of the feedback loop. Accordingly, conditioning was seen as a means of improving feedback regulation, rather than as a means of doing away with it (see Section III-E and Fig. 2). The internal models (see Section III-B) are not part of the story and remain unidentified.

D. From Homeostasis to Allostasis

The role of behavior in homeostasis was clear to Bernard and especially to Cannon. It is particularly evident in the homeostasis of energy and water, which requires interaction with the environment. Hunger and thirst compel an organism to seek food and drink. Indeed, organisms often eat or drink before they become hungry or thirsty, indicating anticipatory behavior [108]. Pavlov himself identified the cephalic phase of insulin secretion, which anticipates before food is consumed the subsequent rise of glucose levels. If these behavioral aspects have long been evident, they have not always been remembered in explaining homeostasis. Set points, negative feedbacks, and control theory provide a compelling abstraction, whose very success may have prompted complacency.

Concerns about the adequacy of the mathematical formulations came from several directions. Curt Richter (1894–1988) introduced the behavioral dimensions in his pioneering studies of appetite regulation [109]. The concept of a set point was already troubling because it is a model parameter whose biological implementation is often obscure. Where does the set point come from? Observations suggest that it is an emergent property subject to modulation by context, rather than one determined by genetics. Some set points, such as that for weight, were argued not to be set at all but were, instead, “settling” points arising from dynamical balances, which could vary with context [110, Sec. 2.1.2]. A dynamic rather than static picture of internal quantities was revealed by real-time measurements: in place of regulation around a set point, a temporal trajectory was being regulated within limits [111, Fig. 2]. There was *stability through change* [112]. (Ironically, this was exactly how Cannon thought [74, Fig. 1];

set points came only with cybernetics.) The significance of anticipatory responses in advance of feedback errors, and the learning that accompanies them, became more widely acknowledged [111], [113], along with the need to situate homeostasis in the broader context of life history and evolutionary fitness [112]. When pushed too far, organisms were found to exhibit stress responses that could reorganize visceral control systems, suggesting how pathological conditions could emerge from sustained overload [112], [114]. Bruce McEwen (1938–2020) made the important discovery of receptors in the hippocampus for stress-related, steroid hormones from the HPA axis, showing that the endocrine system was signaling not only to the visceral tissues but also to the brain [115]. The hypothalamus was talking to the hippocampus! The subtext here was the growing realization that cerebral functions are not just localized in specific brain regions but also require networks of interactions between the regions [110, Sec. 5]. In this case, the network connections go via the body. Brain regions involved in learning and memory were, thereby, seen to be implicated in homeostasis.

Of the many attempts to reformulate homeostasis to accommodate such developments, the one that has acquired the most traction centers on the concept of allostasis, a term coined by Sterling and Eyer [111] but brought to prominence, especially by McEwen [112], [114]. To quote from [111], “*homeostasis’ is flawed: the goal of regulation is not to preserve constancy of the internal milieu ... ‘allostasis’ proposes that efficient regulation requires anticipating needs and preparing to satisfy them before they arise.*” There has been much debate about what this means [116], whether it is really useful [117] and whether allostasis is genuinely different from homeostasis [78]. For us, the most significant aspect of this debate is the focus on anticipation [111], [118]. In the language of Section II-C, this is suggestive of reflective, rather than merely reflexive, internal models. Control theory has no difficulty with reflective models, as we saw with movement control (see Section III-B), but you would be hard-pressed to find them in studies of homeostasis. However, the proponents of allostasis have not yet identified or characterized the models for which such strong claims are being made. The perspective that learning theory brings may offer a way to clarify the concepts and move the debate forward.

E. Return of the Body?

The attention of cognitive scientists, who have long been preoccupied with the CNS, has recently begun to turn inward to interoception and homeostasis [120]. This may be related to a growing interest in emotion and its role in cognitive processes [121]. Keramati and Gutkin [119] have put forward a model of homeostatic reinforcement learning (HRL) (see Fig. 2), which is quite different to the merger attempted by Dworkin (see Section III-C). In their




Fig. 2. Homeostatic reinforcement learning based on [119, Fig. 1]. Graph of the drive function in (6) for a 1-D internal milieu, showing the set point at the minimum value. The reinforcement learning system on the right issues actions as displacements in the milieu (7) for which reductions in drive provide rewards (8).

formulation, the internal milieu is represented by a state vector, $H = (h_1, \dots, h_N)$, whose individual components are the typical regulated quantities, such as temperature or glucose level. A set point, H^{set} , is assumed, and the discrepancy, $D(H)$, between the current value of H and the set point is measured in some appropriate norm, such as the Euclidean distance

$$D(H) = \sqrt{\sum_{i=1}^n (h_i - h_i^{set})^2}. \tag{6}$$

An arbitrary reinforcement learning system can be associated with this minimalist internal milieu, with the actions of the Markov decision process, $K = (k_1, \dots, k_N)$, yielding additive changes to the state

$$H_{new} = H_{old} + K. \tag{7}$$

What closes the loop between the two systems is that the homeostatic discrepancy is interpreted as a motivational drive whose reduction defines the reward, $r(H, K)$

$$r(H_{old}, K) = D(H_{old}) - D(H_{new}). \tag{8}$$

The notion of drive has a long history in psychology [110, Sec. 2.2]. It has been thought of as an intervening variable between stimuli and responses. The idea that drive reduction is coupled to survival goes back to Clark Hull, and this behavioristic interpretation is offered as justification for the drive in the HRL model [122].

It is not difficult to see that, in this formulation, as long as rewards are discounted, reward maximization corresponds to minimizing the distance to the set point. The HRL model has several other attractive features and accounts for experimental data on drug-induced tolerance (see Section III-C) and on differential learning of oral and gastric ingestion of water. It has also offered insights into observed features of cocaine addiction [123].

The HRL model is notable for its minimalist treatment of the internal milieu. The complexity of distributed tissues

and autonomic and endocrine feedback is replaced by the norm computed in (6). The set point is imposed rather than emergent. It is perhaps not surprising that what is accounted for by the HRL model has more to do with learning than with homeostasis. An initial physiological question would be to ask whether a reinforcement learning explanation can be found for the emergence of integral control (see Section III-B). The physiologist would be disappointed by the lack of attention to such issues.

Homeostasis has also been swept up into Karl Friston’s theory of active inference, also known as the “free-energy” principle. Active inference brings the idea of mind as model (see Section I) to a Bayesian culmination. This theory has its roots in the concept of predictive coding in the brain [124]–[126] but approaches it from a variational Bayesian perspective. Gershman [127] has provided a succinct analysis of the relationship between active inference and the Bayesian brain hypotheses. Briefly, predictive coding stipulates a feedback loop between higher (more abstract) and lower (more concrete) brain areas, in which a higher area sends down feedforward predictions about lower area activities based on an encoded internal model of the latter, while the lower area sends up feedback errors between the predictions and its actual activities. The internal model encodes assumptions about the hidden causes underlying lower area activities. Inference of activities from models flows downward, while learning of models from activities flows upward. In the Bayesian view, predictions are probability distributions over model states so that uncertainties are tracked and used to guide choices, and prediction errors are used to update the model according to Bayes’ rule. Computing the resulting posterior distribution remains one of the main challenges with Bayesian approaches. The variational method seeks the Bayesian posterior within some specified family of probability distributions as the minimum distribution under Kullback–Leibler divergence. By choosing the family appropriately, a computationally tractable approximation to the true posterior may sometimes be found. As a variational Bayesian method, active inference is characterized by minimizing not the Kullback–Leibler divergence itself but a related functional called the “free energy” [128], from which the alternative name arises.

Friston [88] has argued that this free-energy principle encompasses, with appropriate assumptions, not only exteroceptive perception of the world but also motor actions in the world (hence, “active inference”), as well as interoceptive perception and homeostatic regulation of the visceral body [129]. In this grand unified theory of the brain, the feedback loops are arranged in a hierarchy between exteroceptive and interoceptive sensorimotor processes in the periphery and cognitive abstractions in higher cortical areas [130].

The integrative scope of the free-energy principle is appealing, but the technical details must be treated with caution. Some of the mathematical claims made for the theory do not hold up to rigorous scrutiny [131].

Friston [132] also claims that “free energy” can be grounded in nonequilibrium physics and can thereby explain how life emerges from matter. This is the greatest of all unsolved problems in physics, so it is not hard to see why Wired magazine referred to Friston as “*The Man who Explained Everything*” [133]. However, free energy in physics is a concept of equilibrium thermodynamics, and it is known that there is no function of state, whether energy or entropy, which can account in general for systems away from thermodynamic equilibrium [134]. The physical interpretation of Fristonian free energy is also not rigorously grounded [131]. Where firm ground gives way to speculation remains contentious; Raviv [133] issues the following trigger warning, “*The free-energy principle is maddeningly difficult to understand. So difficult, in fact, that entire rooms of very, very smart people have tried and failed to grasp it.*” Here, we regard the free-energy principle as an attractive Bayesian hypothesis about the integrated functioning of organisms, whose details need to be clarified in any given context.

In contrast to the HRL model discussed previously, no data are presented in support of active inference in homeostasis [129]. Indeed, the interoceptive aspect of the theory focuses on the ANS to the exclusion of the endocrine system [129, Fig. 2]. The hormonally related implications of active inference would be that, for example, the beta cells in the pancreatic Islets of Langerhans, which regulate glucose levels through insulin secretion, are conveying prediction error signals up the active inference hierarchy. Nothing like this is known or has even been looked for; there has been nothing like the experimental observation of reinforcement error signals in dopaminergic neurons (see Section II-B). Moreover, it is not clear how to reconcile the active inference interpretation with the multiple tissues—pancreas, liver, and muscle—which are implicated in overall glucose regulation. The interplay between insulin sensitivity in muscle and beta-cell function in the pancreas has been a central concern among those studying glucose dysregulation in diabetes [135], and it is implausible that these factors are irrelevant to the internal models. Once again, the physiologist would be disappointed by the lack of attention to physiology.

The inclusion of interoceptive homeostasis in active inference solves a different problem for theories based on the minimization of prediction errors. There is a very easy solution to the minimization problem: the organism stays in a dark room and does nothing [136]. Happily, with interoception now included, the organism’s glucose levels will steadily drift from their set points, increasing the prediction errors and forcing the organism to go to the supermarket [129]. Interoception saves active inference from philosophical embarrassment.

Despite its speculative character, there is something immensely appealing about the unity and scope of the theory of active inference. It offers a universal normative account of how sensation and action are integrated: the organism maximizes the evidence for its models of the

internal and external worlds. It is the kind of foundational theory that one would hope to be correct, and it frames some of the questions that we should be asking about how organisms function. We will come back to it when considering single cells (see Section V-E).

F. Summary

The engagement of cognitive science with the body is a welcome development, but the chasm between physiology and psychology remains clearly visible. For physiologists, learning has offered a means to explore visceral regulation (see Section III-C), and the distinction between reflexive and reflective models may yet help to resolve the confusions surrounding allostasis (see Section III-D). For cognitive scientists, despite aspirations toward universal theories, such as active inference, physiology has been useful largely in service to other concerns (see Section III-E). While both parties rely on internal models (see Section III-B), the locus of these models remains unclear. Cognitive scientists locate them where they have always been, entirely in the brain, but the physiological evidence suggests the involvement of other tissues and the endocrine system, whose collective role in the internal models has not been clarified.

There has been a striking lack of attention to such physiological issues in recent cognitive approaches. The great successes of cognitive science may have encouraged an imperialistic mindset toward colonial domination of other fields, with the regrettable consequences for the indigenous participants of exclusion, if not extinction. This presents a particular difficulty for physiology, which has suffered a crisis of confidence following the molecular revolution [137]. Concurrently, the kinds of invasive animal experiments undertaken by Dworkin and others (see Section III-C) would now be hard to justify under modern ethical standards. It is not clear who would be willing and able to undertake the kinds of experiments required to inspire liberation physiology. Perhaps, bioengineering and synthetic biology offer a new route to analyzing learning in the internal milieu and resisting colonialization [138].

IV. DEVELOPMENTAL ORIGINS OF LEARNING

A. Internal Models in Development

The discussion in Section III was focused on the adult organism, as is often the case in physiological studies. While convenient, this ignores the facts of ontogeny: the animal organism develops from a single fertilized egg through one of the most complex and fascinating processes in all of biology. It has also been one of the most hidden and inaccessible since, in the case of placental mammals, it takes place largely within the body of adult females. Humans are unusual in their neoteny, compared to their primate cousins, and undergo an extensive period of infancy, childhood, and adolescence during which developmental processes continue [139]. Development is where

our familiar machine metaphors break down entirely; we have no machines that construct themselves. Nevertheless, the principal metaphor for understanding development, programming, is borrowed from computation: the organism's genome contains a developmental program that orchestrates ontogeny. It has always been evident that this genomic program must be sensitive to external information and the conditions under which the fetus, embryo, infant, or child is growing. However, the extent and significance of that sensitivity have been the subject of much controversy [140], and the program metaphor of development has largely focused, instead, on how genomic information is processed [141]. In this section, we will survey the evidence that conditions during development determine adult physiology in a predictive manner that is suggestive of an internal model.

B. DOHaD and PARs

It is not surprising that certain adult diseases, such as chronic bronchitis, are correlated with infant mortality across geographical areas because both arise from poor social conditions, which affect infants in the post-neonatal period. The rates of both outcomes have declined over the last century, at least in countries such as the U.K. The case of coronary heart disease, however, is more complex. Its rate has increased over the last century. The conventional explanation for this has been the generally unhealthy nature of western lifestyles. It was paradoxical, therefore, when the epidemiologist David Barker uncovered a striking correlation between the incidence of neonatal mortality during the early 20th century (1910–1925) in England and Wales, and the incidence of coronary heart disease in the same geographical areas in the period 1968–1978 [142, Table 1 and Fig. 1]. Neonatal mortality in the period in question was largely associated with adverse growth of the embryo, rather than with postbirth conditions. Barker's findings led him to propose the hypothesis, which is often named after him, that the uterine environment determines the adult phenotype, and adverse prenatal conditions, such as poor nutrition, could increase the risk of ill health in adult life.

For those brought up in the physical and engineering sciences, statistical studies of this nature can seem elusive. It is important to be cautious in interpreting them because issues of confounding, nonreproducibility, and p-value hacking can be extremely serious [143], [144]. In this case, the data were good, the samples were large, and Barker's work stimulated several subsequent studies that confirmed the general findings [145]. The existence of well-documented cohorts in certain countries, such as those who experienced the Dutch Hunger Winter of 1944, deliberately caused by the Nazi occupying forces, has provided the opportunity for long-term studies that have deepened Barker's initial insights [146]. For example, the timing of embryonic or early childhood events can be important, with particular outcomes predominantly

occurring when events take place during sensitive periods, or critical windows, during development. The discipline that coalesced around these studies has become known as the Developmental Origins of Health and Disease (DOHaD) [147].

The basis for understanding Barker's hypothesis has been the phenomenon of developmental plasticity: the genome can give rise to different phenotypic outcomes depending on context [148]. Such plasticity begins to subvert the program metaphor of development, but it is well established across animals [149]. Insect metamorphosis is an extreme example. Among vertebrates, the sex of certain reptiles is not chromosomally hardwired but determined instead by the temperature at which the egg develops, thereby relying on the gradient within the nest to allocate offspring between the sexes. Some lizards appear to use both chromosomal and temperature-dependent systems, depending on context [150]. Among mammals, West-Eberhard [151] describes the case of Slijper's goat, which was born with congenital paralysis of its forelegs but, nevertheless, learned (the word seems particularly appropriate here) to walk and run on two legs. The many reports of two-legged quadrupeds on the Internet have made it harder to relegate this example to anecdotal status. Meadow voles present a less unusual but still instructive example. The body size, reproductive status, and fur depth of vole pups depend on the amount of daylight experienced by their mothers prior to conception, which correlates well with the seasonal environments, which the pups will subsequently encounter [152]. Among humans, developmental plasticity is manifest in language acquisition, which depends on the linguistic community in which we grow up.

In the setting of developmental plasticity, Barker's hypothesis led Bateson [153] and Gluckman and Hanson [154] to offer a further refinement, now called the predictive adaptive response (PAR) hypothesis [155]. This suggests that mothers provide a forecast of the environmental conditions that their offspring are likely to encounter in later life, and the offspring use this information predictively to regulate their development to match the expected subsequent environment. If those environmental conditions are not met, the mismatch can increase the risk of adult ill-health. Meadow voles provide an example of such predictive anticipation [152]. In humans, offspring born to mothers experiencing a lower nutritional environment are well adapted to those same conditions but poorly adapted to a richer nutritional environment, which can place them at risk for obesity, metabolic syndrome, and cardiovascular diseases [154].

Much of the debate around the PAR hypothesis has focused on its adaptive and evolutionary rationale. This is an important question, but it leads into minefields [156] that will delay us here. What is more relevant for us is the suggestion of learning and prediction during development. Developmental biologists typically view plasticity as arising from alternative paths triggered in response to environmental cues by switch or rheostat-like mechanisms [157].

Can that be seen as learning? We will return to this question in Section VI-C but, first, introduce here some important molecular mechanisms that are implicated in the processes we are discussing.

C. Epigenetic Mechanisms of Development and Learning

One of the primary obstacles to initial acceptance of Barker's hypothesis was the lack of a mechanism through which environmental effects on an embryo could manifest themselves as ill-health in the adult [147]. Epigenetics provided the solution [158]. (I use the word epigenetics here in the weak sense to mean only certain biochemical mechanisms associated with DNA without the stronger implication of creating states that are heritable through cell division [159].) Eukaryotic DNA is wrapped around nucleosomal protein cores to form chromatin [160], and both DNA and the attendant histone proteins can be chemically marked by cytosine methylation on DNA [161] and a veritable zoo of protein post-translational modifications (PTMs) on multiple sites on the histones [162]. Such PTMs are enzymatically regulated, reversible covalent additions of chemical moieties to amino acid residues [163]; examples include phosphorylation, methylation, and ubiquitination. Along with these chemical marks, noncoding RNA entities have emerged, which are intimately involved in epigenetic chromatin regulation [164]. Epigenetic mechanisms participate in regulating gene expression and implementing heritable cell memory during lineage development. They also provide the means whereby environmental information becomes encoded in the internal states of the organism. For example, in a well-controlled, cross-fostering study, Michael Meaney's group analyzed variation in licking, grooming, and arched-back nursing by rat mothers [165]. Those mothers who scored highly for these nurturing instincts raised offspring who had lower stress levels and were less fearful in open-field exploration. The female offspring also scored highly when raising their own pups, showing that the maternal behaviors were transmitted trans-generationally. (This is not Lamarckian inheritance but Lamarck's ghost hovers nearby [166].) Subsequent analysis found correlated epigenetic changes at the stress-regulating glucocorticoid receptor gene in the hippocampus, which emerged during nursing, persisted into adulthood and were necessary for the behavioral responses [167]. Similar forms of epigenetic embryonic programming have now been found in other contexts [168], [169].

The findings of Meaney and colleagues can be interpreted as a PAR, in which the mother's fearlessness and lack of stress, which results in strong nurturing behavior, provides information to the offspring, which predicts that their future environment may also be relatively stress-free. It seems plausible that the accompanying epigenetic changes encode in some way the information that is transferred between the offspring and the mother. In

these studies, the brain emerges as the potential source of the internal models underlying the learning, but related studies have shown similar environmentally dependent epigenetic effects in tissues like the primate liver [170]. This suggests that the internal models underlying developmental learning may be broadly distributed in the body, as also noted in discussing physiology (see Section III-F).

Epigenetic changes in non-neuronal tissues may seem far removed from the internal models of cognitive science, for which neurons provide potential computational resources. However, a striking connection between these different contexts arises at the molecular level in considering the memories that accompany learning. A central problem in neuroscience has been the identity of the "engram" or the physical implementation of a memory [171], [172]. How can memories be biochemically encoded over a lifetime? The developmental context provides another memory, which can also extend over a lifetime, in the identity of terminally differentiated cells. A cell of a particular type typically remembers what it is until it is lost to damage, aging, or death (exceptions being malignant transformation or transdifferentiation during wound healing). Epigenetic mechanisms participate in establishing and maintaining cellular identity during lineage specification in development [173], as they do in encoding environmental memories. In the brain, memories are related to changes in synapses and synaptic strengths during learning, even if the exact nature of the relationship remains contentious (see Section II-A). An important neuroscientific finding in recent years has been that accompanying these changes at synapses are changes of gene expression within brain cells, along with essentially the same kinds of epigenetic changes that are observed during development [174]–[176]. In other words, the two kinds of memories, the one cognitive and the other developmental and environmental, are not so different in terms of their molecular implementation. Perhaps they are also not so different in terms of their role in learning.

D. Summary

The process by which a multicellular organism constructs itself remains astonishing to contemplate, and we struggle to understand it with limited metaphors, such as programming. The evidence from human epidemiological and animal experimental studies is that what is being programmed includes capabilities to anticipate the world that the adult organism will subsequently occupy, based on generously provided maternal information, and to reprogram the developmental process accordingly. Epigenetic mechanisms are implicated both in such developmental programming and also in cognitive learning, a molecular convergence that may reflect functional similarities. Since these mechanisms operate within single cells, we must now turn to those fundamental agents that undertake the construction of the multicellular organism.

V. LEARNING IN SINGLE CELLS

A. The Cell as the Unit of Life

It is easy to forget, in our reductionist enthusiasm, that the cell is the unit of life. Physiological, developmental, and cognitive phenomena depend ultimately on the cells that make up the relevant parts of the organism. The foundation for understanding how learning works in the contexts that we have considered above lies, ultimately, in the capabilities of individual cells. Neuroscience has traditionally focused on networks of cells as the source of cognitive functions, but modern studies are finding surprising capabilities in individual neurons (see Section V-C). Moreover, multicellular organisms evolved from unicellular organisms and the continuities in that transition may be as important as understanding the novelties. From what ancestral features did multicellular learning evolve? Finally, cells are an experimentally tractable system whose capabilities can be studied now in contrast to the challenges involved in the other contexts that we have discussed (see Section V-E).

The great increase in our understanding of the molecular circuitry within cells has prompted analogies with cognition and learning, which reflects the search for a more functional perspective [177], [178]. However, these studies have not yet exploited the advances in cognitive science described above (see Section II). We will use those advances to suggest a fresh perspective on how cells function (see Section V-E). Learning, and especially habituation, in single-cell organisms has been reviewed in [179] and [180].

B. Single-Cell Organisms

Single-cell organisms face the same challenges as all organisms, of navigating through their environment if they are motile, seeking food, avoiding predators, finding mates, and competing and cooperating with other organisms. Debates about their capabilities date back to at least the late 19th century [181], [182]. One school of thought, represented by Jacques Loeb (1859–1924), asserted that cellular behavior could be reduced to chains of overt responses toward external cues. The other school, represented by Herbert Spencer Jennings (see Section I), asserted that cells were capable of agency, decision-making, and learning. Experimental studies were challenging and dogged by questions of reproducibility and interpretation [183]. With the advent of behaviorism, the consensus shifted decisively to Loeb's corner. Single cells were acknowledged to be capable of habituation and sensitization but not of exhibiting more complex behaviors, such as associative learning [183], [184]. Despite the modern resurgence of interest in cognitive capabilities [2], [9], [12], including those of single cells [179], [180], the consensus against complex learning in single-cell organisms has hardly shifted.

Jennings' own experiments on avoidance behavior in *Stentor roeseli* [18] were judged not to be reproducible.

As explained in Section I, the disavowal of Jennings' work had an ideological flavor, and it took until very recently to show that he was right all along [23].

The situation with regard to classical conditioning is more complex. Several attempts have been made to condition various protist species with such mixed results that one participant gave up in exasperation [183]. I will single out three positive studies that seem particularly noteworthy. In the 1950s, Gelber [185] reported conditioning in *Paramecium aurelia* using a needle as the CS with a coating of bacterial food as the US; after repeated training, the *Paramecia* were drawn to the needle even when the food was absent. Her findings were contested on various grounds, which carry less weight from a modern standpoint [24], and her work was marginalized to the point of being forgotten. In the 1970s, Hennessey *et al.* [186] reported conditioning in *Paramecium caudatum* using an AC electrical shock as the US and vibration as the CS. They observed individual organisms and adopted Rescorla's truly random protocol [see Fig. 1(c.2)] in a well-controlled series of experiments. Recently, De La Fuente and colleagues [187], [188] reported conditioning in three species of amoeboid protists, *Amoeba proteus*, *Metamoeba leningradensis*, and *Amoeba borokensis*, using a chemical attractant as the US and galvanotaxis in response to an electric field as the CS. Like Gelber, their experiments were undertaken on populations, but, unlike Hennessey *et al.*, they were not accompanied by controls based on the learning theory described in Section II-A [see Fig. 1(c)].

All three studies give the distinct impression that some form of learning is taking place. That of Hennessey *et al.* remains especially compelling. However, none of the studies has been replicated. Such lack of attention testifies to the prevailing consensus against complex learning in single-cell organisms, despite the evidence just described to the contrary.

C. Single Cells in Multicellular Organisms

In contrast to single-cell organisms, single cells in multicellular organisms have relinquished their autonomy as independent agents for the benefit of the collective. Learning on their part would seem unnecessary, in keeping with the prevailing metaphor of development as the unfolding of a genomic program. The striking similarity of identical twins testifies to the efficiency of their identical genomes in guiding very similar development (it is as well to note here that identical twins start from the same cell, not just the same DNA). However, a different kind of twin offers a counterpoint to the programming narrative. Abby and Brittany Hensel are conjoined twins born in 1990. They have one body, two legs, two arms, and two heads. By all accounts (the Web is the only source of information), they have active, successful lives as teachers. During their development, their cells had to repeatedly confront situations that were decidedly not in the program. Instead of giving up or producing nonsense, they did an astonishingly

good job of making a viable organism. The success speaks to profound capabilities for agency at the cellular level.

Even if we consider normal development, whatever that might be, any given cell will find itself in contexts about which it can have little knowledge as to what neighboring cells it will encounter, what state they are in along their developmental trajectories, how fast they are changing, what signals will be encountered, and what accidents may have occurred, either to itself or to its neighbors, from stochasticity or defects. The situation may be more constrained than the world of a single-cell organism, but it makes up for it in the relentless complexity of possibilities. If we were designing a cell to cope with such an environment, we would be well advised to program it to learn from that environment and work out what is best to do, whatever context it encounters. The genome may specify what species to make, but it is the myriad decisions made by individual cells during ontogeny that determines what individual emerges. Learning may be just as necessary for non-autonomous single cells if only to construct their individual organism in the first place.

In the 1990s, Daniel Koshland (1920–2007) undertook a series of experiments in rat PC12 cells on the habituation of noradrenaline secretion to a variety of repetitive stimuli [189], [190]. What is interesting about these studies is that Koshland interpreted them as analyzing learning and memory and explicitly placed them in the context of contemporary neuroscience. (The best insight into this aspect of Koshland’s thinking comes from his survey on bacterial chemotaxis and neuroscience [8].) For example, he carefully tested [191] for the characteristics of habituation put forward in the learning literature by Thompson and Spencer [192]. What is also interesting, for a different reason, is the paucity of attention paid to Koshland’s habituation work subsequently. According to PubMed, the initial paper [189] has 12 citations, none at all in the decade 1996–2006, while the final paper [190] has none. Despite extensive study of the molecular machinery underlying secretion, there has been no attempt to interpret it from a learning perspective, and when learning has resurfaced in the cell biology literature, it has lost touch with Koshland’s work [177], [178].

PC12 cells are derived from a tumor of the chromaffin cells of the adrenal gland and are of ANS lineage, so Koshland did not stray far from neurons. It is from them, however, that the most compelling evidence has emerged to date about cellular capabilities. The dendritic arbors of pyramidal neurons are capable of computations, such as XOR, which were previously thought to require neuronal networks [193]. Studies of eyeblink conditioning in the cerebellum of ferrets found that individual Purkinje cells can be conditioned to learn a timed pattern of sequential responses [194]. Olfactory sensory neurons exhibit unique patterns of gene expression that appear to encode their individual olfactory perceptions [195]. It seems that neurons, at least, are more sophisticated than previously imagined. They are, indeed, highly specialized, but, as Koshland

pointed out when comparing bacteria to neurons [8], their molecular machinery is shared by all cells. If other kinds of cells are not thought to exhibit such behavior, this may be due to a failure of imagination on our part, rather than incompetence on the part of the cells.

D. Multisite PTMs as Potential Mechanisms

The epigenetic mechanisms discussed above (see Section IV-C) are an obvious candidate for implementing cellular learning, but they typically operate on a longer timescale than is observed, for instance, in ciliate avoidance [23] or eyeblink conditioning [194]. A faster alternative may be protein PTM, which occurs much more broadly than in epigenetics. An instructive example is that of *E. coli* chemotaxis, which exhibits zero steady-state error and has been modeled as an integral control mechanism [196]. We noted previously that there is a close, but still unclarified, relationship between control and learning, with both giving rise to internal models (see Section III-B). In the chemotaxis example, the integral controller variable that is a memory of the system’s history (5) is implemented by the total level of methylation arising from four sites on the appropriate receptor [196]. Four methylation sites on the same bacterial protein appear unusual. Although empirical data are not definitive [197], it seems that, in marked contrast to eukaryotes, bacterial PTMs take place on at most a couple of sites on a given protein. The four methylation sites in *E. coli* cannot be reduced without degrading the chemotactic response [198].

Another example of PTMs encoding information comes from a remarkable interval timer found in *Drosophila* neurons involved in regulating the duration of mating [199]. This timer appears to be implemented by autophosphorylation of CaMKII, a dodecameric, calcium-sensitive protein, each component of which has two phosphorylation sites, giving 24 sites in all on each oligomer. Eukaryotic cells have profound capabilities for such multisite PTMs: the hub protein p53 has over 100 sites of modification [200], giving each molecule of p53 the capacity to exhibit some 10^{30} “modforms,” or global patterns of modification [201]. Only a tiny fraction of these modforms can be present at any one time, but the potential for information encoding is vast. If learning gives rise to internal models that represent information acquired from the environment, the repertoire of protein modforms may be a good place to look for them.

E. Active Inference for Single Cells

We discussed Friston’s active inference hypothesis in the context of homeostasis where it suggests how predictive internal models can integrate both the organism’s response to the external world and the need to maintain its internal state (see Section III-E). Because active inference is formulated abstractly, it can as well be applied to individual cells [202]–[204]. Single cells have a variety of sensory modalities, including receptors for chemical,

mechanical, and electrical stimuli. We typically do not list their senses, in the way we do our own, but perhaps we should to see things from the cell's perspective. Cells have an internal state, given by their patterns of gene expression, nutrient and energy levels, time on their circadian clock, cell cycle period, and so on. Cells have capabilities for action, such as secretion, phagocytosis, and movement. The active inference hypothesis suggests that cells have an internal model of the world, which is continually updated by the discrepancy between the cell's expectations based on the model and the information received from its external and internal senses. From this perspective, the goal-directed behavior of cells, whether to survive and reproduce in the world, for cells that are organisms, or to also construct the organism and maintain it homeostatically, for cells that are parts of larger wholes, is an emergent property of minimizing prediction errors.

Circumstantial evidence suggests that cells have internal models. The example of the VOR (see Section III-B) reveals a fundamental problem that must be universal for all organisms: how do they disentangle the effects of their actions on their own senses from the effects of the environment? If you cannot tell the difference between bumping into a rock and a predator bumping into you, your chances of survival are likely to be slim. The evolutionary pressure to do better must have been fierce. We know that learning based on a reflective internal model has been the solution to this problem for the camera eye of humans, so it is not unreasonable to suspect that something similar, if not as sophisticated, may have evolved ancestrally in eukaryotic cells.

The implications of active inference for single-cell organisms are similar to those for any organism, but this hypothesis also offers a fresh perspective on development. Here, the internal model of a cell would encode its current assumptions about the individual being constructed by the cells that communicate with it, whether chemically, mechanically, or electrically. Lineage development may reflect successive updates of this internal model and attendant changes of internal state, which are inherited by successor cells that adopt transient identities. For example, during hematopoiesis, a pluripotent stem cell can successively become a precursor hematopoietic cell, a myeloid progenitor, a myeloblast, and a monocyte, until, finally, a terminally differentiated macrophage emerges. The necessity for the internal model to be inherited implies that is encoded epigenetically in the strong sense (see Section IV-C). Along the way, a cell may find itself in contexts where something has gone wrong and its prediction error has diverged, but it still “knows” what to do, which is whatever optimally reduces the prediction error. The individuality of the organism (see Section V-C) emerges from the unique trajectories taken by clones of cells, each cell autonomously attempting to minimize its own prediction errors in the collective context of all the other cells that are doing the same. What is programmed is

not the answer, in the eventual individual that is formed—no program could achieve that—but the flexibility to figure out who that individual should be. Moreover, no cell needs to have a global view of the whole organism—a cell may not know that it is part of a conjoined twin—to do its local job as best it can in the circumstances in which it finds itself.

Levin and colleagues [203] outline a simulation of how active inference at the level of individual cells could accomplish the morphogenesis of a schematic multicellular organism from a single cell; more details are provided in [204] with accompanying commentaries in the same issue.

Learning is an explicit part of active inference, so, if the latter is to be tested, it becomes important to determine the learning capabilities of single cells. This has now become more feasible than in studying the physiology or development of a multicellular organism. Technologies such as microfluidics allow live cells to be exposed to complex environments [205], potentially including other cells, whose information content can be varied systematically. Overt responses, such as changes in shape, movement, or secretion, could be tracked by video microscopy and potentially categorized by automated methods [206]. It remains challenging to monitor the cell's internal state, for which only a few measurements of protein state can be obtained concurrently from live cells. The extensive understanding of molecular circuitry may suggest appropriate candidate proteins. An alternative approach to acquiring the necessary data would be to use fixed (dead) cells and exploit spatially resolved measurements of RNA in single cells [207]. Such methods may potentially offer enough signals to detect an increase in mutual information between an individual cell and its environment. As we suggest in the following, such an increase may be considered a necessary condition for learning (9). As to identifying what is learned, and detecting prediction errors within an active-inference hierarchy, that will require substantial effort and ingenuity.

F. Summary

The question of whether single-cell organisms can learn is tantalizingly poised. Evidence for it exists; interest in the question has resurfaced, but definitive experiments have not yet been undertaken. In part, this may stem from the lack of a consensus framework for assessing whether learning has taken place, a problem that we address in the following (see Section VI). Friston's theory of active inference offers a normative suggestion for how learning would enable individual cells to integrate sensing and action and flexibly achieve their goals, whether in the world or within the organism. In the latter context, development may come to be seen not as the programmed construction of the organism but as the process of collectively learning which individual organism to construct.

VI. DEFINING LEARNING FOR SYSTEMS BIOLOGY

A. Learning as the Acquisition of Useful Information

Sections III–V have revealed biological contexts in which forms of learning appear to be taking place, which do not all readily fit into the classical paradigms of Fig. 1. These paradigms were formulated for organisms with brains and overt behavioral responses, which reflects their historical origins in behaviorist psychology. If we strip that baggage away and keep the information processing, what is left are two interconnected features

- 1) increase of mutual information between environmental states and system states
- 2) in which the internal representation of external information can influence subsequent behavior.

Recall that the mutual information, $0 \leq I(X, Y) \leq 1$, between two random variables, X and Y , having the joint space of discrete values $\mathcal{X} \times \mathcal{Y}$, is the Kullback–Leibler divergence between the joint probability distribution, $\Pr_{(X, Y)}$, and the product of the two marginal distributions, \Pr_X, \Pr_Y

$$I(X, Y) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Pr_{(X, Y)}(x, y) \log \left(\frac{\Pr_{(X, Y)}(x, y)}{\Pr_X(x) \Pr_Y(y)} \right). \quad (10)$$

I suggest that (9) constitutes a definition of learning that is more appropriate for studying contexts outside the brain, while accommodating what was discussed in Section II. The focus on mutual information is hardly new and borrows, in particular, from the approach of Gallistel and colleagues (see Section II-B).

Equation (9) is broadly consistent with the “umbrella” definition put forward in [1] but focuses on what is happening inside the system rather than on the overt responses or actions that the system exhibits to its environment. In that sense, it is anti-behaviorist. It accordingly requires knowledge of the system’s internal states, but this has become more readily available at the cellular and molecular levels.

Equation (9.2) is needed in addition to (9.1) to distinguish learning from other forms of information transfer. A system that is driven by its environment could conceivably show an increase in mutual information, as required by (9.1), but this may not correspond to what we think of as learning. For example, when you change time zones, you are exposing yourself to driving by a perturbation to your day–night cycle, and you can literally feel the increase in mutual information as you adjust to jetlag. The recovery from jetlag reprograms your circadian clock in a way that is subsequently used to guide behavior. That recovery process leads to (9.2) and confirms that your body is not just being driven but has actually learned from the experience.

The two requirements of (9) have rather different characters. Equation (9.1) is the fundamental necessary condition for learning to be taking place; (9.2) draws attention to the representation of external information in the internal model, whose existence is then sufficient to confirm learning. Equation (9.1) reflects an observer’s extrinsic perspective, which encompasses both the system and its environment; (9.2) reflects the system’s intrinsic perspective. The mutual information whose increase is detected by an observer may not itself be represented within the system, whose intrinsic representation may encode only what is computationally relevant. Equation (9.1) could conceivably be assessed by methods for estimating (10), which have been developed, for example, for cellular responses to chemical signals [208]–[211]. In contrast, (9.2) may be significantly harder to establish than (9.1). This difficulty is the downside of focusing inward. As previously noted, the physical nature of the cognitive engram remains one of the central problems of neuroscience (see Section IV-C). Although the nature of the representations outside the brain is equally perplexing, these two problems may also be more closely related to each other than previously appreciated (see Section IV-C). It may well be necessary to understand how information is represented by individual neurons in order to understand how cognitive memories are represented by brains. The question of how learned information is represented, encoded, and used is the central challenge to be confronted in adopting a learning-centric perspective, and (9.2) makes that challenge explicit and unavoidable.

B. Examples of Learning That Satisfy (9)

We have noted the close relationship between control and learning (see Section III-B) and discussed bacterial chemotaxis, which has been interpreted as an integral control mechanism (see Section V-D). Learning is not usually mentioned in respect of chemotaxis, but it seems plausible that learning may be necessary to efficiently navigate toward an attractant or away from a repellent. The integral control interpretation makes this more precise through (5), in which the controller variable maintains a memory of the system’s history. During chemotaxis, information in the environment, in the form of concentration gradients of attractants or repellents, becomes coupled to the motions of the bacterium and hence to changes in the controller variable. We would, therefore, expect (9.1) to be satisfied, although an interesting problem is to formally derive the increase in mutual information. Equation (9.2) is more straightforward since the controller variable manifestly influences subsequent behavior. Hence, in the light of (9), learning is implicit in bacterial chemotaxis. We may wonder if the same can be said of chemotaxis in other organisms.

Bacteria also offer a more surprising context in which the two requirements of (9) are clearly satisfied. CRISPR (clustered regularly interspaced short palindromic repeats) is widely known, in the guise of CRISPR–Cas, as a method

of genome engineering [212], but that is a human hijacking of a remarkable adaptive immune system found in many bacteria and most archaea [213]. Microbes with a CRISPR system excise segments of DNA or RNA from invading viruses and integrate these segments between interspaced palindromic repeats in a noncoding CRISPR memory locus within their genomes [213, Fig. 2]. On subsequent reinfection, this DNA locus is transcribed to provide RNA guide segments, which can recognize by base pairing the viral DNA from which they originated, allowing that viral DNA to be destroyed by the microbe. This is immunological learning within an individual cell. We can literally see the increase in mutual information (9.1) in the growth of the CRISPR array, while the nature of the representation and the way it guides behavior (9.2) are evident and based on the classical encoding of information by DNA. Substantially more information is learned here than in chemotaxis. We see that, despite the arguments over whether or not single cells can be conditioned (see Section V-B), single cells—and microbial single cells, to boot—are certainly capable of sophisticated learning, as defined by (9). Moreover, unlike the immune system in mammals [214], learning by CRISPR takes place within individual cells rather than populations of cells.

CRISPR is believed to be specific to microbes, but there is something puzzling about this. Evolution is always repurposing mechanisms for new contexts and it seems odd that it would completely give up on such a clever learning strategy in eukaryotes. One cannot help but wonder if there are more surprises lying in store for us in what is commonly referred to as “junk” DNA.

C. But Why Does It Matter?

CRISPR may exemplify (9), but it also awakens the Devil’s Advocate. What is the benefit of the learning perspective advocated here beyond being a trendy metaphor? After all, CRISPR was not found by studying learning in microbes but by conventional microbiology and molecular biology. Looking forward, what could learning bring to understanding CRISPR or allostasis or PARs? Facing up to such questions helps us understand what is at stake in taking (9) seriously.

At least four benefits accrue from (9): breadth of application; quantification; a focus on how memory is represented and encoded; and a richer metaphor for living systems.

In view of its abstract formulation, (9) is broadly applicable to many contexts that fall outside the scope of Fig. 1. Chemotaxis and CRISPR are good examples. Another is provided by contexts in which observed responses may be long delayed or even absent, as in PARs that manifest as an increased risk of ill health in adult life. Vaccination has a similar character. Like CRISPR, vaccination tangibly meets the two requirements of (9): information about a foreign pathogen is transferred from the world to the system and is represented to guide subsequent behavior. However, this learning may not manifest as

an overt immune response if the corresponding pathogen is not subsequently encountered.

The breadth of application of (9) encourages us to characterize learning within the ecological and evolutionary context in which it is found. If learning arises through one of the classical paradigms of Fig. 1, as in Koshland’s studies of habituation in mammalian cells, the additional controls and experiments developed in the psychology literature [26] offer a starting point for characterization. If learning does not fall into these classical settings but arises directly from (9), as with chemotaxis or CRISPR, it will have its own characteristics that need to be delineated. For example, viral pathogens are subject to mutation, raising the question of how much variation CRISPR can tolerate before it updates its memory. Do continued viral challenges simply accumulate indefinitely in the CRISPR memory, which implies a substantial burden on a microbial genome, or does CRISPR actively forget? The learning perspective suggests new kinds of questions to ask. We may then start to see that there are many kinds of learning in biology, each with the properties appropriate to its natural context.

Equation (9.1) is rigorously quantitative (10), which, unsurprisingly, brings many advantages. First, it offers a way to measure what is learned by specifying the relevant environmental and system states, and estimating their mutual information [208]–[211]. For example, judging from the careful study of meadow voles (see Section IV-B), it seems plausible that experiments could be designed, at least with nonhuman animals, in which the increase in mutual information could be estimated and a rigorous assessment made as to the extent of learning during a PAR. Human studies would be more challenging but not inconceivable [155]. In chemotaxis, quantifying the amount of information that is encoded by receptor methylation may help to explain the efficiency of chemotaxis. In CRISPR, quantifying the information learned may allow a rigorous cost-benefit analysis, which may, in turn, explain why only 50% of bacteria use this form of immunity. More generally, the amount by which mutual information is increased during learning may offer hints about the underlying representation: the internal model may be linked to those system states on which the increase in mutual information is concentrated, and this may help in establishing (9.2). In this way, quantifying the information that is learned can lead to insights into how the information is being used.

Second, quantification introduces a learning spectrum along which many biological processes can be accommodated. We have discussed the example of chemotaxis (see Section VI-B), and in a similar vein, homeostatic processes involved in adapting to environmental changes may also have their own components of learning. Another example brings us back to the question raised at the end of Section IV-B. The development of a multicellular organism is typically viewed as arising from a series of developmental switches within cells, in which choices are made by genetic regulatory networks in response to cues from neighboring cells [141]. Once again, this is not usually

interpreted as learning but some information is certainly transferred from the environment of a cell to its internal states. It is often suggested that the amount of information is limited to whether or not a cue is above the threshold and, therefore, amounts to only a few bits, depending on the patterns of cues that are involved. However, some cues exhibit multiple thresholds [215] and the allegedly low levels of information may reflect a broader bias arising from experimental designs that deliberately mitigate the complexity found *in vivo*. Irrespective of the actual amount, the information is represented within a cell in the states of the molecular circuit that implements the developmental switch, and these states subsequently guide the cell's behavior. We see that development also involves learning in the sense of (9). Such developmental mechanisms may then further underlie the learning of individuality suggested in Section V-C.

More broadly, there has been a long-standing dichotomy encountered at many scales in biology between so-called “instructive” and “selective” (“exploratory” and “permissive”) processes [140], [216]. For example, in contrast to a developmental switch, in which the pattern of external cues provides instruction as to the choice, certain multipotent stem cells may wander around in dynamical state-space attractors, and an external cue merely selects those cells in a particular region of the attractor landscape where they are primed to be responsive to this cue [217]. Both kinds of processes require engagement with the environment, but the selective process is seen as already containing the future possibilities, which the environment only selects. From the perspective of (9), the question is, rather, by how much is the mutual information increased? In selective processes, it may be only one bit, depending on the cue involved; in instructive processes, it may be substantially more. As Bateson [140] suggests, the sharp dichotomy between instruction and selection may be better seen as part of a learning spectrum, for which (9.1) provides the quantification.

We tend to use names other than “learning” to describe biological processes in which only a small number of bits are transferred between environment and system. One of the benefits of (9) is to alert us to the learning that may be concealed within such processes and encourage us to actually measure it.

As noted above with respect to the neuronal engram (see Section IV-C), the problem of how learned information is represented and encoded remains one of the major problems of contemporary science. Equation (9) draws attention to the centrality and broad scope of this problem beyond neuroscience.

Finally, to amplify one of the points made in Section I, perhaps, the most significant benefit of (9) is not in the specific scientific insights or problems that it throws up but in the different metaphors that it offers with which to think about living systems. They may be machines, but they are not classical machines that merely transduce energy and matter; they are autonomous agents that acquire

information from their environment from which they construct models to act in the world. If such a change of thinking takes root, it could have profound consequences for how we conceive of living systems in the light of science.

VII. CONCLUSION

We have covered a lot of ground but I am acutely conscious of what we have not discussed. Immunology is the feature of organisms which is most obviously a form of learning (see Section VI-B) and does not directly involve the brain (although that demarcation may be less clear than we like to think [218]). There has been a rich interplay between immunology and computer science [219], and computational immunology is thriving [220], but immunologists have not felt the urge, so far, to draw on the resources of cognitive science [221]. Immunological learning in mammals occurs in cell populations through processes of positive and negative selection akin to evolution [214]. Similar forms of evolutionary learning and predictive anticipation have been found in studies of bacterial populations [222], [223]. The spectre of evolution has lurked behind many of our discussions, and I should at least point to the interesting relationships that have emerged between evolution and learning [224], [225]. Several theoretical studies, such as [226]–[231], have suggested that single cells have the computational capabilities for complex learning.

I mentioned in Section I that ciliate behavior was, in part, responsible for my interest in learning. The other part was evolution, although from a different angle to that above. I was struggling to understand the evolution of biological complexity. The standard account in the textbooks struck me as entirely inadequate, a view reinforced by the work of Kirschner and Gerhart [216], [232] on the concept of “weak linkage.” Their book *The Plausibility of Life* is, in my opinion, one of the most significant contributions to the contentious subject of *evo-devo* [233], despite the problems that they identify being glossed over in the conventional account [234]. Learning within the organism offers a resolution of these problems, as hinted at in Section V-E, but crossing the treacherous minefields of *evo-devo* must be left to another time.

Systems biology emerged as a discipline to capitalize on the successes of biochemistry, molecular biology, and the genome projects. It has provided rich perspectives of the cell as an ensemble of elaborate machines that transduce matter and energy—polymerases, ribosomes, signaling pathways, gene regulatory and metabolic networks, the cytoskeleton, the cell cycle oscillator, the circadian clock, and so on—and design principles on which these machines are based [235]. I cannot be the only one to find this inspiring but also lacking some fundamental ingredient of living systems. Although biologists study life, they are often reluctant to grapple with its meaning [236]. It can have an ineffable, if not downright mystical, quality that repels scientific analysis. Learning offers an alternative

perspective, which is perfectly scientific, as we have seen (see Section II), but gives back to living systems some of the autonomy and agency that they have lacked as mere, unlearning machines. The rich understanding of learning within cognitive science (see Section II) provides invaluable resources for rebooting the machine metaphor. I have drawn on these resources to suggest an information-theoretic definition of learning that may be more appropriate for systems biology (9). In this view, living systems, at all scales of the biological hierarchy, not only the cognitive, are machines that acquire information from their environments to build internal models of the world that guide their behavior.

The disciplines of physiology (see Section III) and development (see Section IV) offer considerable evidence for the existence of internal models and the significance of learning, but it is at the cellular level (see Section V) that we may most readily hope to test these hypotheses. The example of CRISPR (see Section VI-B) confirms that individual microbial cells undertake substantial learning according to the definition in (9), but the learning capabilities of individual eukaryotic cells and their corresponding molecular mechanisms remain to be unraveled. Doing so may clarify, in turn, the mechanisms underlying the

long-standing conundrum of the cognitive engram (see Section IV-C).

Learning through an internal model offers a scientific way, shorn of mysticism, to articulate a concept of mind, bringing to fruition the ideas of Kant and Helmholtz of mind as model (see Section I). When I was a novice systems biologist, I once struggled to explain the frustrations of doing experiments on living cells, and a wise and experienced biologist smiled and said “*You mean the cells are in charge of the experiment?*” Indeed! They have minds of their own. ■

Acknowledgment

The author is most grateful to three anonymous reviewers, whose perceptive and constructive comments have clarified several sections of this article, and to Hessam Akhlaghpour, Randy Gallistel, Sam Gershman, Mike Levin, and Rosa Martinez-Corral for their helpful comments on the draft. He thanks the participants of the discussion group on cell learning for illuminating the broader context for the ideas described here. Notwithstanding the help received, he remains wholly responsible for the opinions expressed here and any remaining confusions and errors, for which he apologizes.

REFERENCES

- [1] A. B. Barron, E. A. Hebets, T. A. Cleland, C. L. Fitzpatrick, M. E. Hauber, and J. R. Stevens, “Embracing multiple definitions of learning,” *Trends Neurosci.*, vol. 38, pp. 405–407, Jul. 2015.
- [2] R. Powell, I. Mikhalevich, C. Logan, and N. S. Clayton, “Convergent minds: The evolution of cognitive complexity in nature,” *Interface Focus*, vol. 7, no. 3, Jun. 2017, Art. no. 20170029.
- [3] R. F. Thompson, “Habituation: A history,” *Neurobiol. Learn. Memory*, vol. 92, pp. 127–134, Sep. 2009.
- [4] J. Bruner, “A short history of psychological theories of learning,” *Daedalus*, vol. 133, no. 1, pp. 13–20, Jan. 2004.
- [5] A. Trewavas, “Aspects of plant intelligence,” *Ann. Botany*, vol. 92, no. 1, pp. 1–20, May 2003.
- [6] E. B. Rasmussen, “Zombies, invertebrates, and plants, oh my! Introduction to the special section on ‘learning: No brain required,’” *Perspec. Behav. Sci.*, vol. 41, no. 2, pp. 337–341, 2018.
- [7] R. P. Boisseau, D. Vogel, and A. Dussutour, “Habituation in non-neural organisms: Evidence from slime moulds,” *Proc. Roy. Soc. B, Biol. Sci.*, vol. 283, no. 1829, Apr. 2016, Art. no. 20160446.
- [8] D. E. Koshland, “Bacterial chemotaxis in relation to neurobiology,” *Annu. Rev. Neurosci.*, vol. 3, no. 1, pp. 43–75, Mar. 1980.
- [9] F. Baluska and M. Levin, “On having no head: Cognition throughout biological systems,” *Frontiers Psychol.*, vol. 7, p. 902, Jun. 2016.
- [10] F. Adams, “Cognition wars,” *Stud. Hist. Philosophy Sci. A*, vol. 68, pp. 20–30, Apr. 2018.
- [11] R. Solé, M. Moses, and S. Forrest, “Liquid brains, solid brains,” *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 374, no. 1774, Jun. 2019, Art. no. 20190040.
- [12] P. Lyon, F. Keijzer, D. Arendt, and M. Levin, “Reframing cognition: Getting down to biological basics,” *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 376, no. 1820, Mar. 2021, Art. no. 20190750.
- [13] S. Manicka and M. Levin, “The cognitive lens: A primer on conceptual tools for analysing information processing in developmental and regenerative morphogenesis,” *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 374, no. 1774, Jun. 2019, Art. no. 20180369.
- [14] L. R. Swanson, “The predictive processing paradigm has roots in Kant,” *Frontiers Syst. Neurosci.*, vol. 10, p. 79, Oct. 2016.
- [15] W. S. McCulloch and W. Pitts, “A logical calculus of the ideas immanent in nervous activity,” *Bull. Math. Biophys.*, vol. 5, pp. 33–115, Dec. 1943.
- [16] B. Alberts, “The cell as a collection of protein machines: Educating the next generation of molecular biologists,” *Cell*, vol. 92, no. 3, pp. 291–294, 1998.
- [17] L. H. Hartwell, J. J. Hopfield, S. Leibler, and A. W. Murray, “From molecular to modular cell biology,” *Nature*, vol. 402, no. S6761, pp. C47–C52, Dec. 1999.
- [18] H. S. Jennings, *Behavior of the Lower Organisms*. New York, NY, USA: Columbia Univ. Press, 1906.
- [19] D. Bray, *Wetware: A Computer in Every Living Cell*. New Haven, CT, USA: Yale Univ. Press, 2009.
- [20] O. Sacks, “The mental life of plants and worms, among others,” in *The New York Review of Books*, Apr. 2014.
- [21] J. H. Reynierse and G. L. Walsh, “Behavior modification in the protozoan stentor re-examined,” *Psychol. Rec.*, vol. 17, pp. 161–165, Apr. 1967.
- [22] H. Collins and T. Pinch, *The Golem. What You Should Know About Science*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [23] J. P. Dexter, S. Prabhakaran, and J. Gunawardena, “A complex hierarchy of avoidance behaviours in a single-cell eukaryote,” *Curr. Biol.*, vol. 9, pp. 4323–4329, Dec. 2019.
- [24] S. J. Gershman, P. E. Balbi, C. R. Gallistel, and J. Gunawardena, “Reconsidering the evidence for learning in single cells,” *eLife*, vol. 10, Jan. 2021, Art. no. e61907.
- [25] P. Ziegler, *The Black Death*. Cheltenham, U.K.: History Press, 2010.
- [26] C. H. Rankin et al., “Habituation revisited: An updated and revised description of the behavioral characteristics of habituation,” *Neurobiol. Learn. Memory*, vol. 92, pp. 135–138, Sep. 2009.
- [27] P. D. Balsam and C. R. Gallistel, “Temporal maps and informativeness in associative learning,” *Trends Neurosci.*, vol. 32, pp. 73–78, Feb. 2009.
- [28] T. J. Sejnowski, “The once and future Hebb synapse,” *Can. Psychol./Psychol. Canadienne*, vol. 44, no. 1, pp. 17–20, 2003.
- [29] R. E. Brown and P. M. Milner, “The legacy of Donald O. Hebb: More than the Hebb synapse,” *Nat. Rev. Neurosci.*, vol. 4, pp. 1013–1019, Dec. 2003.
- [30] N. Caporale and Y. Dan, “Spike timing-dependent plasticity: A Hebbian learning rule,” *Annu. Rev. Neurosci.*, vol. 31, pp. 25–46, Jul. 2008.
- [31] T. Takeuchi, A. J. Duszkievicz, and R. G. M. Morris, “The synaptic plasticity and memory hypothesis: Encoding, storage and persistence,” *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 369, no. 1633, Jan. 2014, Art. no. 20130288.
- [32] R. A. Rescorla, “Probability of shock in the presence and absence of CS in fear conditioning,” *J. Comparative Physiol. Psychol.*, vol. 66, no. 1, pp. 1–5, 1968.
- [33] L. J. Kamin, “Probability, surprise, attention and conditioning,” in *Punishment and Aversive Behavior*, B. A. Campbell and R. M. Church, Eds. New York, NY, USA: Appleton-Century-Crofts, 1969.
- [34] A. R. Wagner, F. A. Logan, and K. Haberlandt, “Stimulus selection in animal discrimination learning,” *J. Exp. Psychol.*, vol. 76, pp. 171–180, Feb. 1968.
- [35] R. A. Rescorla, “Conditioned inhibition of fear resulting from negative CS-US contingencies,” *J. Comp. Physiol. Psychol.*, vol. 67, pp. 504–509, Apr. 1969.
- [36] C. R. Gallistel and J. Gibbon, “Time, rate, and conditioning,” *Psychol. Rev.*, vol. 107, no. 2, pp. 289–344, 2000.
- [37] C. R. Gallistel and L. D. Matzel, “The neuroscience of learning: Beyond the Hebbian synapse,” *Annu. Rev. Psychol.*, vol. 64, no. 1, pp. 169–200, Jan. 2013.
- [38] C. R. Gallistel and A. P. King, *Memory and the Computational Brain. Why Cognitive Science Will Transform Neuroscience*. Chichester, U.K.: Wiley-Blackwell, 2010.
- [39] J. M. Pearce and M. E. Bouton, “Theories of associative learning in animals,” *Annu. Rev. Psychol.*, vol. 52, pp. 111–139, Feb. 2001.
- [40] N. E. White, E. J. Kehoe, J.-S. Choi, and J. W. Moore, “Coefficients of variation in timing of

- the classically conditioned eyeblink in rabbits," *Psychobiology*, vol. 28, pp. 520–524, Dec. 2000.
- [41] M. Davis, L. S. Schlesinger, and C. A. Sorenson, "Temporal specificity of fear conditioning: Effects of different conditioned stimulus-unconditioned stimulus intervals on the fear-potentiated startle effect," *J. Experim. Psychol., Animal Behav. Processes*, vol. 15, no. 4, pp. 295–310, 1989.
- [42] R. A. Rescorla and A. R. Wagner, "A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non reinforcement," in *Classical Conditioning II: Current Research and Theory*, A. H. Black and W. F. Prokasy, Eds. New York, NY, USA: Appleton-Century-Crofts, 1972.
- [43] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 1998.
- [44] Y. Niv, "Reinforcement learning in the brain," *J. Math. Psychol.*, vol. 53, pp. 139–154, Jun. 2009.
- [45] R. S. Sutton and A. G. Barto, "Towards a modern theory of adaptive networks: Expectation and prediction," *Psychol. Rev.*, vol. 88, pp. 135–170, 1981.
- [46] S. J. Gershman, "A unifying probabilistic view of associative learning," *PLOS Comput. Biol.*, vol. 11, no. 11, Nov. 2015, Art. no. e1004567.
- [47] M. W. Johnson and W. K. Bickel, "Within-subject comparison of real and hypothetical money rewards in delay discounting," *J. Exp. Anal. Behav.*, vol. 77, pp. 129–146, Mar. 2002.
- [48] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Belmont, MA, USA: Athena Scientific, 1996.
- [49] G. Tesauro, "Temporal difference learning and TD-Gammon," *Commun. ACM*, vol. 38, no. 3, pp. 58–68, Mar. 1995.
- [50] W. Schultz, P. Dayan, and P. R. Montague, "A neural substrate of prediction and reward," *Science*, vol. 275, pp. 1593–1599, Mar. 1997.
- [51] W. Schultz, P. Apicella, and T. Ljungberg, "Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task," *J. Neurosci.*, vol. 13, pp. 900–913, Mar. 1993.
- [52] P. R. Montague, P. Dayan, and T. J. Sejnowski, "A framework for mesencephalic dopamine systems based on predictive Hebbian learning," *J. Neurosci.*, vol. 16, pp. 1936–1947, Mar. 1996.
- [53] B. Seymour *et al.*, "Temporal difference models describe higher-order learning in humans," *Nature*, vol. 429, pp. 664–667, Jun. 2004.
- [54] W. Schultz, "Neuronal reward and decision signals: From theories to data," *Physiol. Rev.*, vol. 95, no. 3, pp. 853–951, 2015.
- [55] T. N. Lerner, A. L. Holloway, and J. L. Seiler, "Dopamine, updated: Reward prediction error and beyond," *Curr. Opin. Neurobiol.*, vol. 67, pp. 123–130, Apr. 2021.
- [56] M. D. Iordanova, J. O.-Y. Yau, M. A. McDannald, and L. H. Corbit, "Neural substrates of appetitive and aversive prediction error," *Neurosci. Biobehav. Rev.*, vol. 123, pp. 337–351, Apr. 2021.
- [57] D. Silver *et al.*, "Mastering the game of Go without human knowledge," *Nature*, vol. 550, pp. 354–359, Oct. 2017.
- [58] C. R. Gallistel, A. R. Craig, and T. A. Shahan, "Temporal contingency," *Behavioural Processes*, vol. 101, pp. 89–96, Jan. 2014.
- [59] C. R. Gallistel, A. R. Craig, and T. A. Shahan, "Contingency, contiguity, and causality in conditioning: Applying information theory and Weber's law to the assignment of credit problem," *Psychol. Rev.*, vol. 126, pp. 761–773, Oct. 2019.
- [60] C. R. Gallistel, "Robert Rescorla: Time, information and Contingency," *Revista de Historia de la Psicología*, vol. 41, no. 1, pp. 7–21, 2021.
- [61] N. Chater, J. B. Tenenbaum, and A. Yuille, "Probabilistic models of cognition: Conceptual foundations," *Trends Cog. Sci.*, vol. 10, pp. 287–291, Jul. 2006.
- [62] A. Perfors, J. B. Tenenbaum, T. L. Griffiths, and E. Xu, "A tutorial introduction to Bayesian models of cognitive development," *Cognition*, vol. 120, pp. 302–321, Sep. 2011.
- [63] S. J. Gershman, E. J. Horvitz, and J. B. Tenenbaum, "Computational rationality: A converging paradigm for intelligence in brains, minds, and machines," *Science*, vol. 349, pp. 273–278, Jul. 2015.
- [64] R. J. Kruschke, "Bayesian approaches to associative learning: From passive to active learning," *Learn. Behav.*, vol. 36, pp. 210–226, Aug. 2008.
- [65] E. C. Tolman, "Cognitive maps in rats and men," *Psychol. Rev.*, vol. 55, pp. 180–208, Jul. 1948.
- [66] R. J. Dolan and P. Dayan, "Goals and habits in the brain," *Neuron*, vol. 80, pp. 312–325, Oct. 2013.
- [67] T. Hartley, C. Lever, N. Burgess, and J. O'Keefe, "Space in the brain: How the hippocampal formation supports spatial cognition," *Phil. Trans. Roy. Soc. B: Biol. Sci.*, vol. 369, no. 1635, Feb. 2014, Art. no. 20120510.
- [68] E. I. Moser, M.-B. Moser, and B. L. McNaughton, "Spatial representation in the hippocampal formation: A history," *Nat. Neurosci.*, vol. 20, pp. 1448–1464, Nov. 2017.
- [69] A. Johnson and A. D. Redish, "Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point," *J. Neurosci.*, vol. 27, pp. 12176–12189, Nov. 2007.
- [70] D. Kahneman, *Thinking, Fast and Slow*. New York, NY, USA: Farrar, Strauss and Giroux, 2011.
- [71] N. D. Daw, S. J. Gershman, B. Seymour, P. Dayan, and R. J. Dolan, "Model-based influences on humans' choices and striatal prediction errors," *Neuron*, vol. 69, pp. 1204–1215, Mar. 2011.
- [72] D. P. Todes, *Ivan Pavlov. A Russian Life in Science*. Oxford, U.K.: Oxford Univ. Press, 2014.
- [73] C. Bernard, *Lectures on the Phenomena of Life Common to Animals and Plants*. Springfield, IL, USA: Charles C. Thomas Publications, 1974.
- [74] W. B. Cannon, "Organization for physiological homeostasis," *Physiol. Rev.*, vol. 9, no. 3, pp. 399–431, Jul. 1929.
- [75] W. B. Cannon, *The Wisdom of the Body*. New York, NY, USA: W. W. Norton and Co., Inc., 1932.
- [76] N. Wiener, *Cybernetics: Or Control and Communication in the Animal and the Machine*, 2nd ed. Cambridge, MA, USA: MIT Press, 1948.
- [77] N. Wiener, *The Cybernetics Group*, 2nd ed. Cambridge, MA, USA: MIT Press, 1948.
- [78] R. H. S. Carpenter, "Homeostasis: A plea for a unified approach," *Adv. Physiol. Educ.*, vol. 28, no. 4, pp. 180–187, Dec. 2004.
- [79] D. A. Robinson, "Integrating with neurons," *Annu. Rev. Neurosci.*, vol. 12, pp. 33–45, Mar. 1989.
- [80] H. El-Samad, J. P. Goff, and M. Khammash, "Calcium homeostasis and parieturent hypocalcemia: An integral feedback perspective," *J. Theor. Biol.*, vol. 214, no. 1, pp. 17–29, Jan. 2002.
- [81] J. H. Koeslag, P. T. Saunders, and E. Terblanche, "A reappraisal of the blood glucose homeostat which comprehensively explains the type 2 diabetes mellitus-syndrome X complex," *J. Physiol.*, vol. 549, pp. 333–346, Jun. 2003.
- [82] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, pp. 457–465, Sep. 1976.
- [83] E. Sontag, "Adaptation and regulation with signal detection implies internal model," *Syst. Control. Lett.*, vol. 50, pp. 119–126, Oct. 2003.
- [84] W. R. Ashby, "Requisite variety and its implications for the control of complex systems," *Cybernetica*, vol. 1, pp. 83–99, Jan. 1958.
- [85] R. C. Conant and W. R. Ashby, "Every good regulator of a system must be a model of that system," *Int. J. Syst. Sci.*, vol. 1, pp. 89–97, Oct. 1970.
- [86] M. Kawato, "Internal models for motor control and trajectory planning," *Curr. Opin. Neurobiol.*, vol. 9, no. 6, pp. 718–727, 1999.
- [87] E. Todorov, "Optimality principles in sensorimotor control," *Nat. Neurosci.*, vol. 7, pp. 907–915, Sep. 2005.
- [88] K. Friston, "What is optimal about motor control?" *Neuron*, vol. 72, pp. 488–498, Nov. 2011.
- [89] J. Huang, A. Isidori, L. Marconi, M. Mischiati, E. Sontag, and W. M. Wonham, "Internal models in control, biology and neuroscience," in *Proc. IEEE Conf. Decis. Control (CDC)*, Dec. 2018, pp. 5370–5390.
- [90] E. Roth, K. Zhuang, S. A. Stamper, E. S. Fortune, and N. J. Cowan, "Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*," *J. Exp. Biol.*, vol. 214, pp. 1170–1180, 2011.
- [91] M. A. Frens, "Forward models and state estimation in compensatory eye movements," *Frontiers Cellular Neurosci.*, vol. 3, p. 13, Nov. 2009.
- [92] M. Joshua and S. G. Lisberger, "A tale of two species: Neural integration in zebrafish and monkeys," *Neuroscience*, vol. 296, pp. 80–91, Jun. 2015.
- [93] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, "Principles of sensorimotor learning," *Nature Rev. Neurosci.*, vol. 12, pp. 739–751, Dec. 2011.
- [94] P. Sachse, U. Beeremann, M. Martini, T. Maran, M. Domeier, and M. R. Furtner, "The world is upside down"—The Innsbruck Goggle experiments of Theodor Erisman (1883–1961) and Ivo Kohler (1915–1985)," *Cortex*, vol. 92, pp. 222–232, Jul. 2017.
- [95] G. M. Jones, "Plasticity in the adult vestibulo-ocular reflex," *Philos. Trans. Roy. Soc. London. B, Biol. Sci.*, vol. 278, pp. 319–334, Apr. 1977.
- [96] C. Tin and C.-S. Poon, "Internal models in sensorimotor integration: Perspectives from adaptive control theory," *J. Neural Eng.*, vol. 2, no. 3, pp. S147–S163, Sep. 2005.
- [97] B. R. Dworkin, *Learning and Physiological Regulation*. Chicago, IL, USA: Univ. Chicago Press, 1993.
- [98] B. Brembs, "Operant conditioning in invertebrates," *Curr. Opin. Neurobiol.*, vol. 13, pp. 710–717, Dec. 2003.
- [99] N. E. Miller, "Learning of visceral and glandular responses," *Science*, vol. 163, pp. 434–445, Jan. 1969.
- [100] B. R. Dworkin and N. E. Miller, "Failure to replicate visceral learning in the acute curarized rat preparation," *Behav. Neurosci.*, vol. 100, no. 3, pp. 299–314, 1986.
- [101] E. E. Coons, *Neal E. Miller: 1909–2002* (Biographical Memoirs). Washington, DC, USA: National Academy of Sciences, 2014.
- [102] B. R. Dworkin and S. Dworkin, "Heterotopic and homotopic classical conditioning of the baroreflex," *Integr. Physiol. Behav. Sci.*, vol. 33, pp. 158–174, Jul. 1999.
- [103] B. R. Dworkin, T. Elbert, and H. Rau, "Blood pressure elevation as a coping response," in *Stress, Coping and Cardiovascular Disease*, P. M. McCabe, N. Schneiderman, T. Field, and A. R. Wellens, Eds. Mahwah, NJ, USA: Lawrence Erlbaum Associates Inc., 2000.
- [104] J. R. Jennings and A. F. Heim, "From brain to behaviour: Hypertension's modulation of cognition and affect," *Int. J. Hypertens.*, vol. 2012, Jan. 2012, Art. no. 701385.
- [105] S. Siegel, "Drug anticipation and drug addiction. The 1998 H. David Archibald lecture," *Addiction*, vol. 94, pp. 1113–1124, Aug. 1999.
- [106] T. E. Robinson and K. B. Berridge, "The incentive sensitisation theory of addiction: Some current issues," *Philos. Trans. Roy. Soc. B, Biol. Sci.*, vol. 363, pp. 3137–3146, Oct. 2008.
- [107] C. F. Flaherty, P. S. Grigson, and A. Brady, "Relative novelty of conditioning context influences directionality of glycemic conditioning," *J. Exp. Psychol., Animal Behav. Processes*, vol. 13, pp. 144–149, Apr. 1987.
- [108] M. L. Andermann and B. B. Lowell, "Toward a wiring diagram understanding of appetite control," *Neuron*, vol. 95, pp. 759–778, Aug. 2017.
- [109] S. C. Woods and D. S. Ramsay, "Homeostasis: Beyond Curt Richter," *Appetite*, vol. 49, pp. 388–398, Sep. 2007.
- [110] K. C. Berridge, "Motivation concepts in behavioral neuroscience," *Physiol. Behav.*, vol. 81, no. 2, pp. 179–209, Apr. 2004.
- [111] P. Sterling, "Allostasis: A model of predictive regulation," *Physiol. Behav.*, vol. 106, no. 1, pp. 5–15, Apr. 2012.

- [112] B. S. McEwen and J. C. Wingfield, "The concept of allostasis in biology and biomedicine," *Hormones Behav.*, vol. 43, no. 1, pp. 2–15, Jan. 2003.
- [113] D. S. Ramsay and S. C. Woods, "Physiological regulation: How it really works," *Cell Metab.*, vol. 13, pp. 361–364, Sep. 2016.
- [114] B. S. McEwen and E. Stellar, "Stress and the individual," *Arch. Internal Med.*, vol. 153, pp. 2093–2101, Sep. 1993.
- [115] B. S. McEwen, "Interview with Bruce S. McEwen," *Trends Neurosci.*, vol. 36, pp. 207–208, Apr. 2013.
- [116] B. S. McEwen and J. C. Wingfield, "What is in a name? Integrating homeostasis, allostasis and stress," *Horm. Behav.*, vol. 57, pp. 105–111, Feb. 2010.
- [117] T. A. Day, "Defining stress as a prelude to mapping its neurocircuitry: No help from allostasis," *Prog. Neuropsychopharmacol. Biol. Psychiatry*, vol. 29, pp. 1195–1200, Dec. 2005.
- [118] J. Schulkin and P. Sterling, "Allostasis: A brain-centered, predictive mode of physiological regulation," *Trends Neurosci.*, vol. 42, pp. 740–752, Oct. 2019.
- [119] M. Keramati and B. Gutkin, "Homeostatic reinforcement learning for integrating reward collection and physiological stability," *eLife*, vol. 3, Dec. 2014, Art. no. e04811.
- [120] E. H. Petzschner, S. N. Garfinkel, M. P. Paulus, C. Koch, and S. S. Khalsa, "Computational models of interoception and body regulation," *Trends Neurosci.*, vol. 44, no. 1, pp. 63–76, Jan. 2021.
- [121] A. K. Seth and K. J. Friston, "Active interoceptive inference and the emotional brain," *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 371, no. 1708, Nov. 2016, Art. no. 20160007.
- [122] O. J. Hulme, T. Morville, and B. Gutkin, "Neurocomputational theories of homeostatic control," *Phys. Life Rev.*, vol. 31, pp. 214–232, Dec. 2019.
- [123] M. Keramati, A. Durand, P. Girardeau, B. Gutkin, and S. H. Ahmed, "Cocaine addiction as a homeostatic reinforcement learning disorder," *Psychol. Rev.*, vol. 124, pp. 130–153, Mar. 2017.
- [124] M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: A fresh view of inhibition in the retina," *Proc. Roy. Soc. London. Ser. B, Biol. Sci.*, vol. 216, pp. 427–459, Nov. 1982.
- [125] D. Mumford, "On the computational architecture of the neocortex II. The role of cortico-cortical loops," *Biol. Cybern.*, vol. 66, pp. 241–251, Jan. 1992.
- [126] R. P. Rao and D. H. Ballard, "Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects," *Nature Neurosci.*, vol. 2, no. 1, pp. 79–87, 1999.
- [127] S. J. Gershman, "What does the free energy principle tell us about the brain?" *NBDT*, vol. 2, no. 3, 2019.
- [128] R. Bogacz, "A tutorial on the free-energy framework for modelling perception and learning," *J. Math. Psychol.*, vol. 76, pp. 198–211, Feb. 2017.
- [129] G. Pezzulo, F. Rigoli, and K. Friston, "Active inference, homeostatic regulation and adaptive behavioural control," *Prog. Neurobiol.*, vol. 134, pp. 17–35, Nov. 2015.
- [130] K. Friston, "The free energy principle: A unified brain theory?" *Nat. Rev. Neurosci.*, vol. 11, pp. 127–138, Feb. 2010.
- [131] M. Biehl, F. A. Pollock, and R. Kanai, "A technical critique of some parts of the free energy principle," 2020, [arXiv:2001.06408](https://arxiv.org/abs/2001.06408).
- [132] K. Friston, "Life as we know it," *J. R. Soc. Interface*, vol. 10, Sep. 2013, Art. no. 20130475.
- [133] S. Raviv, "The man who explained everything," *Wired*, vol. 26, p. 96, Dec. 2018.
- [134] R. Landauer, "Inadequacy of entropy and entropy derivatives in characterizing the steady state," *Phys. Rev. A, Gen. Phys.*, vol. 12, pp. 636–638, Aug. 1975.
- [135] J. J. Noland and K. Faerch, "Estimating insulin sensitivity and beta cell function: Perspectives from the modern pandemics of obesity and type 2 diabetes," *Diabetologia*, vol. 55, pp. 2863–2867, Nov. 2012.
- [136] Z. Sun and C. Firestone, "The dark room problem," *Trends Cogn. Sci.*, vol. 24, pp. 346–348, May 2020.
- [137] D. Noble, "Neo-Darwinism, the modern synthesis and selfish genes: Are they of use in physiology?" *J. Physiol.*, vol. 589, pp. 1007–1015, Mar. 2011.
- [138] N. Olsman, "Cybernetics, systems biology and the phenomenological gap," *IEEE Contr. Syst. Mag.*, vol. 41, no. 3, pp. 92–98, May 2021.
- [139] D. Lieberman, *The Story of the Human Body: Evolution, Health and Disease*. New York, NY, USA: Pantheon, 2013.
- [140] P. Bateson, "Biological approaches to the study of behavioural development," *Int. J. Behav. Develop.*, vol. 10, no. 1, pp. 1–22, Mar. 1987.
- [141] E. H. Davidson, *The Regulatory Genome: Gene Regulatory Networks in Development and Evolution*. Burlington, MA, USA: Academic, 2006.
- [142] D. J. P. Barker, "The origins of the developmental origins theory," *J. Intern. Med.*, vol. 261, pp. 412–417, Mar. 2007.
- [143] J. P. A. Ioannidis, "Why most published research findings are false," *PLoS Med.*, vol. 2, no. 8, p. e124, Aug. 2005.
- [144] L. D. Nelson, J. Simmons, and U. Simonsohn, "Psychology's renaissance," *Annu. Rev. Psychol.*, vol. 69, pp. 511–534, Jan. 2018.
- [145] I. C. Mcmillen and J. S. Robinson, "Developmental origins of the metabolic syndrome: Prediction, plasticity, and programming," *Physiol. Rev.*, vol. 85, no. 2, pp. 571–633, Apr. 2005.
- [146] L. C. Schulz, "The Dutch Hunger Winter and the developmental origins of health and disease," *Proc. Nat. Acad. Sci. USA*, vol. 107, pp. 16757–16758, Sep. 2010.
- [147] P. D. Wadhwa, C. Buss, S. Entringer, and J. M. Swanson, "Developmental origins of health and disease: Brief history of the approach and current focus on epigenetic mechanisms," *Seminars Reproductive Med.*, vol. 27, no. 5, pp. 358–368, 2009.
- [148] P. Bateson et al., "Developmental plasticity and human health," *Nature*, vol. 430, pp. 419–421, Jul. 2004.
- [149] M. J. West-Eberhard, *Developmental Plasticity and Evolution*. Oxford, U.K.: Oxford Univ. Press, 2003.
- [150] R. S. Radder, A. E. Quinn, A. Georges, S. D. Sarre, and R. Shine, "Genetic evidence for co-occurrence of chromosomal and thermal sex-determining systems in a lizard," *Biol. Lett.*, vol. 4, pp. 176–178, Apr. 2008.
- [151] M.-J. West-Eberhard, "Phenotypic accommodation: Adaptive innovation due to developmental plasticity," *J. Exp. Zool. B, Mol. Developmental Evol.*, vol. 304, pp. 610–618, Nov. 2005.
- [152] T. M. Lee and I. Zucker, "Vole infant development is influenced perinatally by maternal photoperiodic history," *Amer. J. Physiol.-Regulatory, Integrative Comparative Physiol.*, vol. 255, no. 5, pp. R831–R838, Nov. 1988.
- [153] P. Bateson, "Fetal experience and good adult design," *Int. J. Epidemiol.*, vol. 30, pp. 928–934, Oct. 2001.
- [154] P. D. Gluckman and M. A. Hanson, "The developmental origins of the metabolic syndrome," *Trends Endocrinol. Metabol.*, vol. 15, pp. 183–187, May 2004.
- [155] P. Bateson, P. Gluckman, and M. Hanson, "The biology of developmental plasticity and the predictive adaptive response hypothesis," *J. Physiol.*, vol. 592, pp. 2357–2368, Jun. 2014.
- [156] K. Laland et al., "Does evolutionary theory need a rethink?" *Nature*, vol. 514, pp. 161–164, Oct. 2014.
- [157] E. Lafuente and P. Beldade, "Genomics of developmental plasticity in animals," *Frontiers Genet.*, vol. 10, p. 720, Aug. 2019.
- [158] R. Bonasio, S. Tu, and D. Reinberg, "Molecular signals of epigenetic states," *Science*, vol. 330, pp. 612–616, Oct. 2010.
- [159] M. Ptashne, "On the use of the word 'epigenetic,'" *Current Biol.*, vol. 17, no. 7, pp. R233–R236, Apr. 2007.
- [160] S. Venkatesh and J. L. Workman, "Histone exchange, chromatin structure and the regulation of transcription," *Nat. Rev. Mol. Cell Biol.*, vol. 16, pp. 178–189, Mar. 2015.
- [161] D. Schübeler, "Function and information content of DNA methylation," *Nature*, vol. 517, pp. 321–326, Jan. 2015.
- [162] M. A. J. Morgan and A. Shilatifard, "Reevaluating the roles of histone-modifying enzymes and their associated chromatin modifications in transcriptional regulation," *Nature Genet.*, vol. 52, pp. 1271–1281, Dec. 2020.
- [163] C. T. Walsh, *Posttranslational Modification Proteins*. Englewood, CO, USA: Roberts and Company, 2006.
- [164] D. Holoch and D. Moazed, "RNA-mediated epigenetic regulation of gene expression," *Nature Rev. Genet.*, vol. 16, no. 2, pp. 71–84, Feb. 2015.
- [165] D. Francis, J. Diorio, D. Liu, and M. J. Meaney, "Nongenomic transmission across generations of maternal behavior and stress responses in the rat," *Science*, vol. 286, pp. 1155–1158, Dec. 1999.
- [166] M. Szyf, "Lamarck revisited: Epigenetic inheritance of ancestral odor fear conditioning," *Nature Neurosci.*, vol. 17, no. 1, pp. 2–4, Jan. 2014.
- [167] I. C. G. Weaver et al., "Epigenetic programming by maternal behavior," *Nature Neurosci.*, vol. 7, pp. 847–854, Aug. 2004.
- [168] M. C. Vogt et al., "Neonatal insulin action impairs hypothalamic neurocircuit formation in response to maternal high-fat feeding," *Cell*, vol. 156, no. 3, pp. 495–509, Jan. 2014.
- [169] S. Das, S. Min, and V. Prahlaad, "Gene bookmarking by the heat shock transcription factor programs the insulin-like signaling pathway," *Mol. Cell*, vol. 81, pp. 4843–4860, Dec. 2021.
- [170] K. M. Aagaard-Tillery et al., "Developmental origins of disease and determinants of chromatin structure: Maternal diet modifies the primate fetal epigenome," *J. Mol. Endocrinol.*, vol. 41, no. 2, pp. 91–102, Aug. 2008.
- [171] M.-M. Poo et al., "What is memory? The present state of the engram," *BMC Biol.*, vol. 14, no. 1, p. 40, Dec. 2016.
- [172] J. Tee and D. P. Taylor, "Where is memory information stored in the brain?" 2021, [arXiv:2112.05362](https://arxiv.org/abs/2112.05362).
- [173] P. A. Steffen and L. Ringrose, "What are memories made of? How Polycomb and Trithorax proteins mediate epigenetic memory," *Nat. Rev. Mol. Cell Biol.*, vol. 15, pp. 340–356, May 2014.
- [174] J. J. Day and J. D. Sweatt, "Epigenetic mechanisms in cognition," *Neuron*, vol. 70, pp. 813–829, Jun. 2011.
- [175] E.-L. Yap and M. E. Greenberg, "Activity-regulated transcription: Bridging the gap between neural activity and behavior," *Neuron*, vol. 100, pp. 330–348, Oct. 2018.
- [176] M. B. Chen, X. Jiang, S. R. Quake, and T. C. Südhof, "Persistent transcriptional programmes are associated with remote memory," *Nature*, vol. 587, pp. 42–437, Nov. 2020.
- [177] A. Koseska and P. I. H. Bastiaens, "Cell signaling as a cognitive process," *EMBO J.*, vol. 36, pp. 568–582, Mar. 2017.
- [178] P. Csermely et al., "Learning of signaling networks: Molecular mechanisms," *Trends Biochem. Sci.*, vol. 45, pp. 284–290, Apr. 2020.
- [179] S. K. Y. Tang and W. F. Marshall, "Cell learning," *Current Biol.*, vol. 28, Jun. 2018, Art. no. R1180.
- [180] A. Dussotour, "Learning in single cell organisms," *Biochem. Biophys. Res. Commun.*, vol. 564, pp. 92–102, Jul. 2021.
- [181] P. J. Pauly, "The Loeb–Jennings debate and the science of animal behavior," *J. Hist. Behav. Sci.*, vol. 17, pp. 504–515, Oct. 1981.
- [182] H. S. Jennings, "Diverse ideals and divergent conclusions in the study of behavior in lower organisms," *Amer. J. Psychol.*, vol. 21, pp. 349–370, Jul. 1910.
- [183] P. B. Applewhite, "Learning in protozoa," in *Biochemistry and Physiology of Protozoa*, M. Levandowsky and S. H. Hunter, Eds. New York, NY, USA: Academic Press, 1979, ch. 11,

pp. 55–341.

[184] E. M. Eisenstein, *Aneural Organisms in Neurobiology*. New York, NY, USA: Plenum Press, 1975.

[185] B. Gelber, “Investigations of the behavior of *Paramecium aurelia*: I. Modification of behavior after training with reinforcement,” *J. Comparative Physiol. Psychol.*, vol. 45, no. 1, pp. 58–65, 1952.

[186] T. M. Hennessey, W. B. Rucker, and C. G. McDiarmid, “Classical conditioning in *Paramecia*,” *Anim. Learn. Behav.*, vol. 7, pp. 417–423, Dec. 1979.

[187] I. M. De la Fuente *et al.*, “Evidence of conditioned behavior in amoebae,” *Nature Commun.*, vol. 10, no. 1, p. 3690, Dec. 2019.

[188] J. Carrasco-Pujante *et al.*, “Associative conditioning is a robust systemic behavior in unicellular organisms: An interspecies comparison,” *Frontiers Microbiol.*, vol. 12, Jul. 2021, Art. no. 707086.

[189] P. N. McFadden and D. E. Koshland, Jr., “Habituation in the single cell: Diminished secretion of norepinephrine with repetitive depolarization of PC12 cells,” *Proc. Nat. Acad. Sci. USA*, vol. 87, pp. 2031–2035, Mar. 1990.

[190] P. T. Martin and D. E. Koshland, Jr., “Neurosecretory habituation in PC12 cells: Modulation during parallel habituation,” *Proc. Natl. Acad. Sci. USA*, vol. 92, pp. 5052–5056, May 1995.

[191] L. Cheever and D. E. Koshland, Jr., “Retention of habituation in PC12 cells,” *Proc. Nat. Acad. Sci. USA*, vol. 89, Nov. 1992, Art. no. 10084.

[192] R. F. Thompson and W. A. Spencer, “Habituation: A model phenomenon for the study of neuronal substrates of behavior,” *Psychol. Rev.*, vol. 73, no. 1, pp. 16–43, 1966.

[193] A. Gidon *et al.*, “Dendritic action potentials and computation in human layer 2/3 cortical neurons,” *Science*, vol. 387, pp. 83–87, Jan. 2020.

[194] D.-A. Jirenhed, A. Rasmussen, E. Johansson, and G. Hesselow, “Learned response sequences in cerebellar Purkinje cells,” *Proc. Nat. Acad. Sci. USA*, vol. 114, pp. 6127–6132, Jun. 2017.

[195] T. Tsukahara, D. H. Brann, S. L. Pashkovski, G. Guitchounts, T. Bozz, and S. R. Datta, “A transcriptional rheostat couples past activity to future sensory responses,” *Cell*, vol. 184, pp. 6326–6343, Dec. 2021.

[196] T.-M. Yi, Y. Huang, M. I. Simon, and J. Doyle, “Robust perfect adaptation in bacterial chemotaxis through integral feedback control,” *Proc. Nat. Acad. Sci. USA*, vol. 97, pp. 4649–4653, Apr. 2000.

[197] M. Zhang, J.-Y. Xu, H. Hu, B.-C. Ye, and M. Tan, “Systematic proteomic analysis of protein methylation in prokaryotes and eukaryotes revealed distinct substrate specificity,” *Proteomics*, vol. 18, no. 1, Jan. 2018, Art. no. 1700300.

[198] A. Krembel, R. Colin, and V. Sourjik, “Importance of multiple methylation sites in *Escherichia coli* chemotaxis,” *PLoS ONE*, vol. 10, no. 12, Dec. 2015, Art. no. e0145582.

[199] S. C. Thornquist, K. Langer, S. X. Zhang, D. Rogulja, and M. A. Crickmore, “CaMKII measures the passage of time to coordinate behavior and motivational state,” *Neuron*, vol. 105, pp. 334–345, Jan. 2020.

[200] C. J. DeHart, J. S. Chahal, S. J. Flint, and D. H. Perlman, “Extensive post-translational modification of active and inactivated forms of endogenous p53,” *Mol. Cellular Proteomics*, vol. 13, no. 1, pp. 1–17, Jan. 2014.

[201] S. Prabakaran, G. Lippens, H. Steen, and J. Gunawardena, “Post-translational modification: Nature’s escape from from genetic imprintation and the basis for cellular information processing,” *Wiley Interdiscip. Rev. Syst. Biol. Med.*, vol. 4, pp. 565–583, Jan. 2012.

[202] G. Auletta, “Information and metabolism in bacterial chemotaxis,” *Entropy*, vol. 15, pp. 311–326, Jan. 2013.

[203] K. Friston, M. Levin, B. Sengupta, and G. Pezzulo, “Knowing one’s place: A free-energy approach to pattern regulation,” *J. Roy. Soc. Interface*, vol. 12, no. 105, Apr. 2015, Art. no. 20141383.

[204] F. Kuchling, K. Friston, G. Georgiev, and M. Levin, “Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems,” *Phys. Life Rev.*, vol. 33, pp. 88–108, Jul. 2020.

[205] A. M. Streets and Y. Huang, “Chip in a lab: Microfluidics for next generation life science research,” *Biomicrofluidics*, vol. 7, no. 1, Jan. 2013, Art. no. 011302.

[206] S. Han, E. Taralova, C. Dupre, and R. Yuste, “Comprehensive machine learning analysis of hydra behavior reveals a stable basal behavioral repertoire,” *eLife*, vol. 7, Mar. 2018, Art. no. e32605.

[207] N. Battich, T. Stoeger, and L. Pelkmans, “Image-based transcriptomics in thousands of single human cells at single-molecule resolution,” *Nature Methods*, vol. 10, pp. 1127–1133, Nov. 2013.

[208] R. Cheong, A. Rhee, C. J. Wang, I. Nemenman, and A. Levchenko, “Information transduction capacity of noisy biochemical signaling networks,” *Science*, vol. 334, pp. 354–358, Oct. 2011.

[209] J. Selimkhanov *et al.*, “Accurate information transmission through dynamic biochemical signaling networks,” *Science*, vol. 346, pp. 1370–1373, Dec. 2014.

[210] S. A. Cepeda-Humerez, J. Ruess, and G. Tkačik, “Estimating information in time-varying signals,” *PLOS Comput. Biol.*, vol. 15, no. 9, Sep. 2019, Art. no. e1007290.

[211] J. B. Lee, L. M. Caywood, J. Y. Lo, N. Levering, and A. J. Keung, “Mapping the dynamic transfer functions of eukaryotic gene regulation,” *Cell Syst.*, vol. 12, no. 11, pp. 1079–1093, 2021.

[212] J. A. Doudna and E. Charpentier, “The new frontier of genome engineering with CRISPR-Cas9,” *Science*, vol. 346, no. 6213, Nov. 2014.

[213] P. Horvath and R. Barrangou, “CRISPR/Cas: The immune system of bacteria and archaea,” *Science*, vol. 327, pp. 70–167, Jan. 2010.

[214] K. A. Hogquist, T. A. Baldwin, and S. C. Jameson, “Central tolerance: Learning self-control in the thymus,” *Nature Rev. Immunol.*, vol. 5, pp. 772–778, Oct. 2005.

[215] H. Greenfield, J. Lin, and M. C. Mullins, “The BMP signaling gradient is interpreted through concentration thresholds in dorsal–ventral axial patterning,” *PLoS Biol.*, vol. 19, no. 1, Jan. 2021, Art. no. e3001059.

[216] M. Kirschner and J. Gerhart, “Evolvability,” *Proc. Nat. Acad. Sci. USA*, vol. 95, pp. 8420–8427, Jul. 1998.

[217] T. Kalmar *et al.*, “Regulated fluctuations in Nanog expression mediate cell fate decisions in embryonic stem cells,” *PLoS Biol.*, vol. 7, no. 7, Jul. 2009, Art. no. e1000149.

[218] T. Koren *et al.*, “Insular cortex neurons encode and retrieve specific immune responses,” *Cell*, vol. 184, pp. 1–14, Nov. 2021.

[219] S. Forrest, S. A. Hofmeyr, and A. Somayaji, “Computer immunology,” *Commun. ACM*, vol. 40, no. 10, pp. 88–96, 1997.

[220] A. K. Chakraborty, “A perspective on the role of computational models in immunology,” *Annu. Rev. Immunol.*, vol. 35, pp. 403–439, Apr. 2017.

[221] D. L. Farber, M. G. Netea, A. Radbruch, K. Rajewsky, and R. M. Zinkernagel, “Immunological memory: Lessons from the past and a look to the future,” *Nature Rev. Immunol.*, vol. 16, pp. 124–128, Feb. 2016.

[222] I. Tagkopoulou, Y.-C. Liu, and S. Tavazoie, “Predictive behavior within microbial genetic networks,” *Science*, vol. 320, pp. 1313–1317, Jun. 2008.

[223] A. Mitchell *et al.*, “Adaptive prediction of environmental changes by microorganisms,” *Nature*, vol. 460, pp. 220–224, Jul. 2009.

[224] G. E. Hinton and S. J. Nowlan, “How learning can guide evolution,” *Complex Syst.*, vol. 1, pp. 495–502, Jun. 1987.

[225] R. A. Watson and E. Szathmáry, “How can evolution learn?” *Trends. Ecol. Evol.*, vol. 31, pp. 57–147, Feb. 2016.

[226] C. T. Fernando *et al.*, “Molecular circuits for associative learning in single-celled organisms,” *J. Roy. Soc. Interface*, vol. 6, pp. 463–469, May 2009.

[227] B. Bryant, “Chromatin computation,” *PLoS ONE*, vol. 7, no. 5, May 2012, Art. no. e35703.

[228] J. Macia, B. Vidiella, and R. V. Solé, “Synthetic associative learning in engineered multicellular consortia,” *J. Roy. Soc. Interface*, vol. 14, no. 129, Apr. 2017, Art. no. 20170158.

[229] M. Gabalda-Sagarra, L. B. Carey, and J. Garcia-Ojalvo, “Recurrence-based information processing in gene regulatory networks,” *Chaos, Interdiscipl. J. Nonlinear Sci.*, vol. 28, no. 10, Oct. 2018, Art. no. 106313.

[230] H. Akhlaghpour, “An RNA-based theory of natural universal computation,” *J. Theor. Biol.*, vol. 537, p. 110984, Mar. 2022.

[231] S. Biswas, S. Manicka, E. Hoel, and M. Levin, “Gene regulatory networks exhibit several kinds of memory: Quantification of memory in biological and random transcriptional networks,” *iScience*, vol. 24, no. 3, Mar. 2021, Art. no. 102131.

[232] M. Kirschner and J. Gerhart, “The theory of facilitated variation,” *Proc. Nat. Acad. Sci. USA*, vol. 104, pp. 8582–8589, May 2007.

[233] M. W. Kirschner and J. C. Gerhart, *The Plausibility Life*. New Haven, CT, USA: Yale Univ. Press, 2005.

[234] B. Charlesworth, “On the origins of novelty and variation,” *Science*, vol. 310, pp. 1619–1620, Dec. 2005.

[235] U. Alon, *An Introduction to Systems Biology: Design Principles of Biological Circuits*. London, U.K.: Chapman & Hall, 2006.

[236] J. Monod, *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology*. New York, NY, USA: Alfred A. Knopf, 1971.

ABOUT THE AUTHOR

Jeremy Gunawardena is a pure mathematician by training. He spent some years in industrial research at Hewlett-Packard Research Labs in Bristol, U.K., and Palo Alto, CA, USA, before returning to academic life at Harvard Medical School in Boston, MA, USA. He is currently an Associate Professor in the Department of Systems Biology. His research group studies cellular information processing using a combination of experiments and theory.

