



Design Principles for Neurobotics

Jeffrey L. Krichmar^{1,2*†} and Tiffany J. Hwu^{3†}

¹ Department of Cognitive Sciences, University of California, Irvine, Irvine, CA, United States, ² Department of Computer Science, University of California, Irvine, Irvine, CA, United States, ³ Riot Games, Los Angeles, CA, United States

In their book “How the Body Shapes the Way We Think: A New View of Intelligence,” Pfeifer and Bongard put forth an embodied approach to cognition. Because of this position, many of their robot examples demonstrated “intelligent” behavior despite limited neural processing. It is our belief that neurobots should attempt to follow many of these principles. In this article, we discuss a number of principles to consider when designing neurobots and experiments using robots to test brain theories. These principles are strongly inspired by Pfeifer and Bongard, but build on their design principles by grounding them in neuroscience and by adding principles based on neuroscience research. Our design principles fall into three categories. First, organisms must react quickly and appropriately to events. Second, organisms must have the ability to learn and remember over their lifetimes. Third, organisms must weigh options that are crucial for survival. We believe that by following these design principles a robot’s behavior will be more naturalistic and more successful.

OPEN ACCESS

Keywords: adaptive behavior, embodiment, learning, memory, neuromodulation, value

Edited by:

Florian Röhrbein,
Technische Universität Chemnitz,
Germany

Reviewed by:

André Cyr,
University of Ottawa, Canada
Hendry Ferreira Chame,
Université de Lorraine, France

*Correspondence:

Jeffrey L. Krichmar
jkrichma@uci.edu

[†]These authors have contributed
equally to this work

Received: 23 February 2022

Accepted: 19 April 2022

Published: 25 May 2022

Citation:

Krichmar JL and Hwu TJ (2022)
Design Principles for Neurobotics.
Front. Neurobot. 16:882518.
doi: 10.3389/fnbot.2022.882518

1. INTRODUCTION

Neurobotics is a powerful tool for testing brain theories and increasing our understanding of neuroscience. The robot controller is modeled after some aspect of the nervous system. Unlike human or other animal studies, the neuroboticist has access to every aspect of this artificial brain during the lifetime of the agent. Therefore, the neuroboticist can analyze and perturb the nervous system in ways that a neuroscientist cannot with present recording technology. Not only can a neurobot be tested under laboratory conditions that are similar to those of an animal experiment in order to provide direct comparisons, but it can also be tested in more natural conditions to see how these brain functions might respond to real-world situations.

The actions of a neurobot are embedded in its environment. By choosing an appropriate morphology, simple mechanical designs can perform complex functions by taking advantage of environmental features, thus alleviating slow, power-hungry nervous systems from having to make these calculations. This is known as morphological computation (Pfeifer and Bongard, 2006). Neurobot designs can be degenerate, that is, they can contain multiple systems capable of performing the same functions (Edelman and Gally, 2001). In this way the agent can still survive in the environment should one system fail. Similar to many mobile operating systems, neurobot computation can follow the brain architecture by having multiple processes run in parallel in an event-driven manner, continuously responding to concurrent events. These ideas have roots in behavior-based robots (Brooks, 1991; Arkin, 1998) and the design of neuromorphic hardware (Merolla et al., 2014; Davies et al., 2018).

To adapt to a changing environment a neurobot must be able to learn, store and recall information. Memory systems in neurobotics are particularly applicable in spatial memory for navigation and contextual memory for learning representations of the environment (Milford et al., 2016; Gaussier et al., 2019; Hwu et al., 2020). Success within dynamic environments, such as the real-world, requires the processing of risk, reward, and uncertainty by some notion of value and the ability to adapt (Oudeyer and Kaplan, 2007; Krichmar, 2008; Merrick, 2017). Through such systems, neurobots are able to predict future events and adapt to changes in the environment.

The world is full of trade-offs and changing needs that require us to make choices. Incorporating behavioral trade-offs into neurobots such as reward vs. punishment, invigorated vs. withdrawn activity, expected vs. unexpected uncertainty for attention, exploration vs. exploitation of choices, foraging for food vs. defending one's territory, coping with stress vs. keeping calm, and social interaction vs. solitary restraint can lead to interesting behavior in neurobotics (Canamero et al., 2006; Hiolle et al., 2012; Krichmar, 2013; Lones et al., 2018). Many of these trade-offs are regulated by the neuromodulators and hormone levels in the brain.

In this article, we present a set of principles to take into consideration when designing neurobots. They fall into three categories: 1) Embodiment and reactions, 2) Adaptive behavior, learning and memory, and 3) Behavioral trade-offs. Following these design principles can make neurobots more naturalistic and more interesting. Many of the ideas put forth in this article are based on material from our forthcoming book (Hwu and Krichmar, 2022).

2. EMBODIMENT AND REACTIONS - RESPONDING TO THE HERE AND NOW

In this first set of neurobotic design principles, we focus on what (Pfeifer and Bongard, 2006) called the “here and now.” These design principles are grounded in neuroscience and are focused on processes that respond to events. Even without learning and memory, these processes lead to flexible, adaptive behavior.

2.1. Embodiment

Brains do not work in isolation; they are closely coupled with the body acting in its environment (Chiel and Beer, 1997). Biological organisms perform *morphological computation*; that is, certain processes are performed by the body that would otherwise be performed by the brain (Pfeifer and Bongard, 2006). Moment-to-moment action can be handled at the periphery by the body, sensors, actuators, and reflexes at the spinal cord level. This allows the central nervous system, which is slower and requires more processing than the body or peripheral nervous system, to predict, plan, and adapt by comparing its internal models with current information from the body (Shadmehr and Krakauer, 2008; Hickok et al., 2011).

In biology, the morphology and behavior must fit within the organism's ecological niche. Therefore, the layout of its sensors

and actuators, their resolution and range are tuned to meet the specific organism's needs (Ziemke, 2003). As an example, consider a toddler flailing his or her arms. The child's arms more easily move toward the front of the torso than the back. By chance the toddler's hand touches an object causing a reflexive grasping motion. This leads to the fingers, where most of the tactile sensors are located, touching the object. The child will then move this hand in its easiest direction, which tends to be toward the face, where a range of sensors for vision, olfaction, and taste receptors reside. Comparing that with the embodiment and design of an insect or a fish, it is clear that these design implementations are specific to the organism's niche. Attention to these environmental details can provide guidance for the design of neurobots. The form of the organism's body shapes its behavior and its brain function. This requires a brain and body that is engaged with the environment, which is what we should strive for in designing our neurobots.

For many tasks that we carry out with ease, our brains are too slow to sense, process and move. For example, skiing down a hill or catching a wave on a surfboard happens too fast for our central nervous system to position the body appropriately. But the form and compliance of the body can position itself properly and adjust itself to perturbations without brain control. This is morphological computation in action.

Trapping a soccer ball is another example of morphological computation. In a RoboCup tournament, the Segway soccer team from The Neurosciences Institute solved this difficult sensorimotor problem with a very cheap design (Fleischer et al., 2006). On a large playing field it was nearly impossible for the robot to catch a fast-moving soccer ball, given that the Segway was large and cumbersome and had a slow camera frame rate and slow IR sensor refresh rate. Soccer balls would bounce off the robot before it had a chance to respond. After much trial and error, the team used plastic tubing that was fastened around the robot's body like a hula-hoop at just the right height (see **Figure 1**). Any ball that was passed to the Segway robot was trapped by the tubing, giving the robot time to use its camera and IR proximity sensors to place the ball in its kicking apparatus. In a sense, this is what humans do when playing soccer. They use soft pliant materials angled appropriately to soften the impact of a ball coming toward them. Many actions like these take place without much thought (i.e., brain processing).

By putting more emphasis on designs that exploit the environment, we can offload some of the control from the cognitive robot's central nervous system onto the body itself. This should allow the robot to be more responsive to the environment and more fluid in its actions. In addition, it frees up the nervous system to put more emphasis on planning, prediction, and error correction rather than reflexive movements. Too often cognitive neuroscientists forget that the body and the peripheral nervous system are performing many vital, moment-to-moment behaviors and tasks without central control. Even functions that are thought to be purely mental have a basis in embodiment. For example, it has been argued that interactions in which a person needs to understand another is an embodied process rather than an internal simulation (Gallagher, 2001).



FIGURE 1 | Video sequence showing robot capturing a soccer ball using morphological computation. A pliant plastic hoop around the Segway robot allowed trapped the soccer ball, allowing the slower IR sensors and camera to confirm that a soccer ball was caught.

2.2. Efficiency Through Cheap Design

Cheap design means finding the simplest solution to the challenge the robot is facing. One way to do this is by exploiting the environment. For example, winged insects and fast swimming fish exploit their environment by creating vortices with their wing beats or fin movements, which causes additional thrust and more energy to come out than the animal put in. Most of our touch receptors are where we need them most, at our fingertips. It would be wasteful to have this fine level of resolution on the back of our hands or arms. This is what is meant by cheap.

For example, in legged locomotion, roboticists have put much time and effort in creating robust controllers for legged locomotion. The importance of the cheap design principle can be observed when comparing the biped locomotion of passive dynamic walking robots to sophisticated humanoid robots, such as Honda's Asimo or Aldebaran's NAO. Passive dynamic walking robots exploit gravity, friction, and the forces generated by their swinging limbs (Collins et al., 2005). As a result, they require very little energy or control to move (see **Table 1**). In contrast, robots such as Asimo need complex control systems and long-lasting batteries to achieve the same result. Although these passive walkers are not necessarily biologically inspired, once the engineers or artists implement a design that minimizes energy expenditure, the gait looks very natural.

However, it should be stressed that passive dynamics and morphological computation are not enough to support a complete range of natural behaviors. Rather it frees up the system from expending energy and computational resources on some functions, while allowing it to concentrate on other functionality. Saving energy is a recurring theme in biology since biological organisms are under tight metabolic constraints (Beyeler et al., 2019; Krichmar et al., 2019). However, there is a trade-off that comes with efficiency. For example, it is more efficient to walk on four legs, but then arms are not available for manual dexterity and gestures, which is important for bipedal organisms. In a complete system, passive control is closely coupled with spinal cord reflexes, which in turn are in close communication with motor cortex and other areas of the brain. These issues should be taken into consideration when designing neurobots.

It is not only the body that follows the principle of cheap design: brains do as well. Biological systems are

TABLE 1 | Energy consumed during legged locomotion (Collins et al., 2005). Unit weight per unit distance.

Agent	Energy consumption
Asimo	3.23
Cornell biped	0.20
Humans	0.20

under extreme metabolic constraints and need to represent information efficiently. Therefore, the nervous system must encode information as cheaply as possible. The brain operates on a mere 20 watts of power, approximately the same power required for a ceiling fan operating at low speed (Krichmar et al., 2019). Although being severely metabolically constrained is at one level a disadvantage, evolution has optimized brains in ways that lead to incredibly efficient representations of important environmental features that are distinctly different from those employed in current digital computers. The brain utilizes many means to reduce its functional metabolic energy use. Indeed, one can observe at every level of the nervous system strategies to maintain high performance and information transfer while minimizing energy expenditure.

At the neuronal coding level, the brain uses several strategies to reduce neural activity without sacrificing performance. Neural activity (i.e., the generation of an action potential, the return to resting state, and synaptic processing) is energetically very costly, and this drives the minimization of the number of spikes necessary to encode the neural representation of a new stimulus. Such sparse coding strategies appear to be ubiquitous throughout the brain (Olshausen and Field, 2004; Beyeler et al., 2019). Efficient coding reduces redundancies and rapidly adapts to changes in the environment. At a macroscopic scale, the brain saves energy by minimizing the wiring between neurons and brain regions (i.e., number of axons) and yet still communicates information at a high level of performance (Laughlin and Sejnowski, 2003). Information transfer between neurons and brain areas is preserved by a small-world network architecture, which reduces signal propagation (Sporns, 2010). These energy-saving ideas should be taken into consideration in constructing neural controllers for robots, which like biological organisms

have limited energy resources. Moreover, many of these strategies could inspire new methodologies for constructing power-efficient artificial intelligence.

2.3. Sensory-Motor Integration

In the brain, the sensory and motor systems are tightly coupled. An organism or robot may get new sensory information that causes an action. Each action then creates new sensory information. Neurorobots can take advantage of this tightly coupled loop. For example, figure-ground segregation is a difficult computer vision problem in which a scene of static objects needs to be recognized from the background (e.g., a small toy block sitting on a similarly colored table). However, segmentation can be facilitated by sensorimotor integration in a very natural way as was demonstrated in a robot experiment by Fitzpatrick and Metta (2003). In their study, the robot's hand moved until it happened to hit the toy block, triggering motion detector responses in its visual system. In this way the robot's motor system generated sensory information, both visual and tactile, which led to the unexpected recognition of an object on the table. In the nervous system, a copy of the action, called a motor efference copy, is fed back to the brain. It creates an expectation that can be used to error check the movement and the expected sensory experience. Because hitting the toy creates a violation of both tactile and visual sensory expectations, the toy block is easily differentiated from the table.

It is important to emphasize how much the sensory and motor nervous systems are intertwined. Too often neuroscientists study these systems separately, but they are highly interconnected and work in concert. Although there are brain areas specialized for sensing such as auditory cortex and visual cortex, and there are areas of the brain devoted to action such as the motor cortex, most of the cortex is associational and cannot be called exclusively sensory or motor systems. These associational cortical areas are highly interconnected and the delineation between perception and action becomes blurred (Fuster, 2004). The parietal cortex receives multimodal sensory inputs and is important for planning movements. By multimodal we mean that the brain area receives more than one sense: auditory, olfactory, taste, visual, touch, or vestibular. The frontal cortex also receives multimodal inputs and is important for decisions, control of actions, and action selection. These multimodal association areas have direct influence on what we perceive and how we move.

2.4. Degeneracy

Degeneracy is the ability of elements that are structurally different to perform the same function or yield the same output (Edelman and Gally, 2001). To be fault tolerant and flexible a system's architecture should be designed such that different subsystems have different functional processes and there is an overlap of functionality between subsystems. In this design, if any subsystem fails the overall system can still function. This is different from redundancy, in which an identical system copy is kept in case there is a system failure (e.g., redundant array of independent disk [RAID] computer memory systems). Degeneracy appears throughout biology, from low-level processes such as the genetic code and protein folding

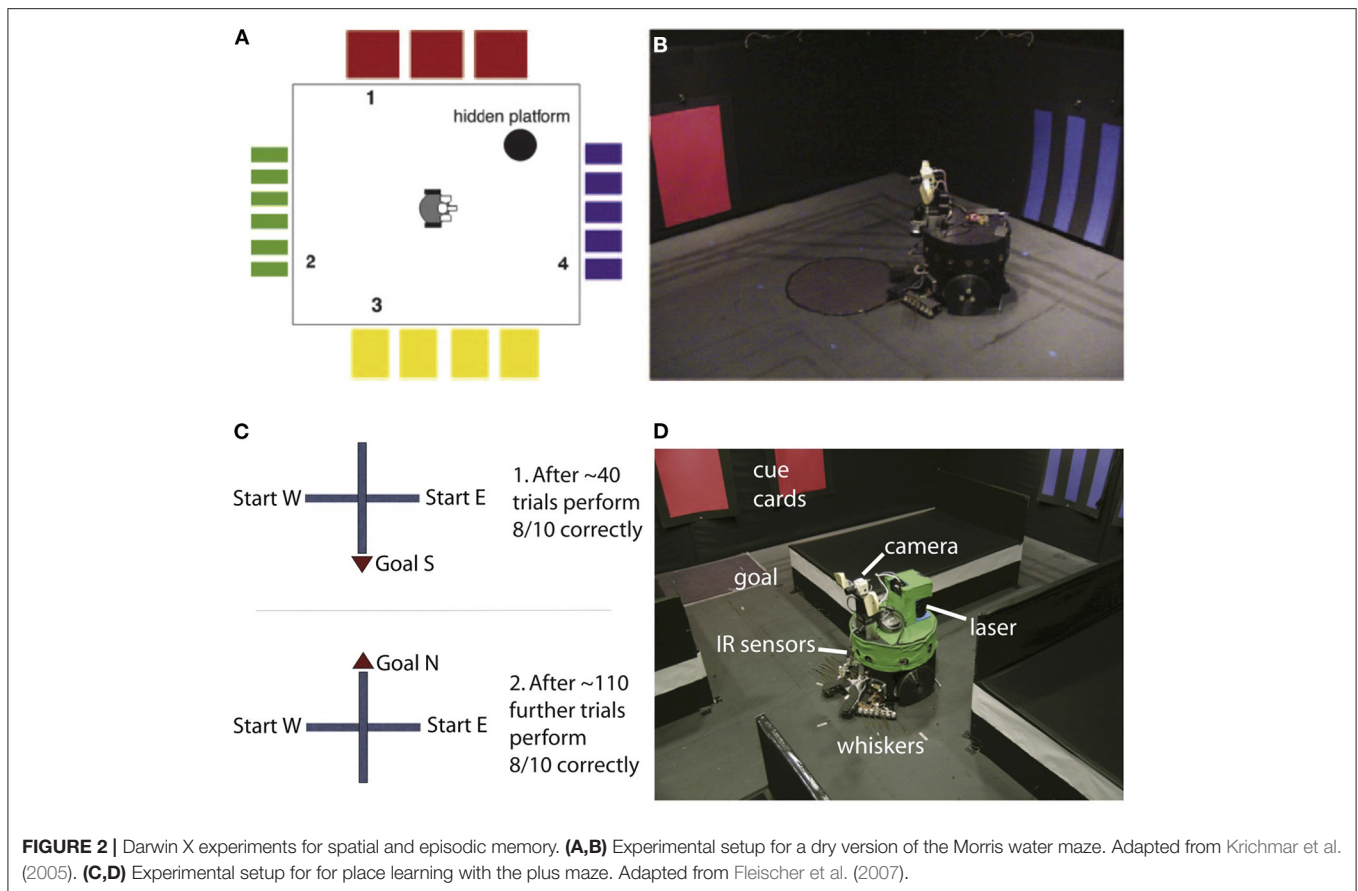
to system-level processes such as behavioral repertoires and language. For example, there are four nucleotide bases in DNA (thymine, cytosine, adenine, and guanine). It takes three bases to encode an amino acid, which is the building block of proteins. This means that there are potentially 4^3 or sixty-four possible combinations, but only twenty amino acids make up the proteins found in the human body. In many cases, different triplets encode the same amino acid. Therefore, the genetic code is considered degenerate. As a result, the genetic code is fault tolerant to mutations. The heterogeneity of neuron types within and between brain regions, as well as between organisms is another example of degeneracy. For instance, the nervous system has numerous cell types which can be distinguished by their anatomy, connectivity or firing behavior (Ascoli et al., 2007; Wheeler et al., 2015). Furthermore, organisms like the nematode *C. Elegans* have neurons that don't fire action potentials (Sarma et al., 2018). Despite this variability, these neuronal elements communicate with sensors, actuators, and other brain regions that have similar properties and often operate in the same environment. At the other end of the biological spectrum is communication. We have an almost infinite number of ways to communicate the same message. The same message could be communicated through voice, text, email, Morse code, gesture, or facial expressions. Degeneracy and variability is not only important to demonstrate for biological realism, it also leads to robustness and fault tolerance when operating in noisy, dynamic environments.

Degeneracy at multiple levels was nicely demonstrated by the neurobot Darwin X, which was used to demonstrate spatial and episodic memory (Krichmar et al., 2005; Fleischer et al., 2007). Darwin X solved a dry version of the Morris water maze and a place-learning version of a plus maze (see **Figure 2**). In the Morris water maze, a rat swims through murky water until it finds a platform hidden beneath the surface. After several days of exploration, the rat will swim directly to the platform from any starting location. If the hippocampus is damaged, the rat cannot learn the location of the platform. In the dry version, a reflective piece of paper was the proxy for the platform. It was the same color as the floor, so the robot could not see the platform but it could "feel" when it was on the platform by means of a downward-pointing light sensor.

Darwin X had an extensive model of the hippocampus formation and its surrounding cortical regions. Similar to a rodent, as the robot explored its environment, hippocampal place cells emerged. In the Darwin X experiments, combinations of place cells were used to plan routes to goals. The robot's behavior and neural activity were directly compared with rodent experiments. Like the rat, place cells could be used to predict not only the current location of the robot but also the location from which the robot came and the location to which it was heading. In the robot experiments, several levels of degeneracy emerged.

Degeneracy at the Neuronal Level

Because the neuroboticists were able to track every neuron in Darwin's simulated nervous system, they were able to trace the neuronal activity that led to hippocampal place activity. Although hippocampal place activity was similar on different



trials when the robot passed through the same location on the same heading, the neuronal activity leading to that neuron's place activity on a given trial differed dramatically. That is, different neural activation patterns led to the same hippocampal place cell outcome.

Degeneracy at the Systems Level

Darwin X received sensory input from its camera (vision), whiskers (somatosensory), compass (head direction), and laser range finder (depth/distance). Darwin X's spatial memory was multimodal and degenerate. Even when one or more of its sensory modalities were lesioned, Darwin X's behavior and place cell activity remained stable. Different sensory pathways led to the same outcome of knowing where Darwin was located.

Degeneracy at the Behavioral Level

In the Morris maze task, nine different Darwin X subjects, which consisted of the identical robots with slightly different nervous systems due to variations in initial synaptic connection probabilities, solved the same spatial navigation task in unique ways. Some subjects bounced off the "red" wall to the hidden platform, some bounced off the "blue" wall, and others went directly toward the platform location. The proficiency of each subject differed as well. Some were better learners than others. However, despite their idiosyncrasies they all shared the same outcome of solving this task.

2.5. Multitasking and Event-Driven Processing

Cognitive scientists tend to study the brain in a serialized fashion by focusing on one subsystem at a time, be it a type of memory or a specific perceptual effect. But intelligence emerges from many processes operating in parallel and driven by events. We (humans and other organisms) are multitaskers, and to multitask we must do things in parallel. The brain is the ultimate event-driven, parallel computer. There is no overarching clock as in computer architectures. Rather the brain responds to events when they happen. Therefore, neurorobots should have a multitask design that responds to multiple, asynchronous events in a timely manner.

Neurons throughout the brain are responding simultaneously to multiple events. Although the different parts of the brain may be acting somewhat independently, they are highly interactive. The sensory system is telling the motor system what it senses, and the motor system is telling the sensory system what its last action was. This brain analogy can be extended to the whole organism, in which control is parallel, asynchronous, and spatiotemporally matched with the real world.

A classic way of studying cognitive science was a serial process of "sense, think, and act," which guided many artificial intelligence robot designs. In response to this, Rodney Brooks and Ron Arkin developed behavior-based robots that responded

asynchronously to events (Brooks, 1991; Arkin, 1998). There are parallels between the subsumption architecture's layered hierarchical design and the nervous system. Neuroscientist Larry Swanson proposed a basic plan for the nervous system that somewhat follows this design; a four-component functional systems model with a motor system controlling behavior and visceral functions (i.e., internal organs and bodily functions), whose output is a function of activity in sensory, cognitive, and behavioral state systems, with feedback throughout (Swanson, 2007). It should be noted that in this view the central nervous system is only one component. Areas that regulate basic behaviors and internal monitoring are subcortical.

Correlates of the subsumption architecture can be observed in the brain's structure. Prescott et al. (1999) suggested that the hierarchical, layered subsumption architecture could describe the neural mechanisms of defense for the rat. For example, the lower levels were reactive and included withdrawal, startle, and freezing. The higher levels subsumed the lower levels by suppressing behavior or predicting outcomes through conditioning. Sensory input provides stimuli that can trigger behavioral responses. Another example involves self-monitoring systems in the nervous system (Chiba and Krichmar, 2020). **Figure 3** shows an overview of this architecture with brain systems on the left and the corresponding robot control systems on the right. There is low-level control for sensor processing and motor reflexes. The autonomic nervous system may subsume these lower levels to maintain homeostasis or adapt set points. Higher levels can set states or context that may shape responses.

Multitasking and event-driven processing is prevalent in current technology due in part to the ubiquity of real-time and embedded systems. Most modern computing devices, including smartphones, desktop computers, onboard automotive computers and entertainment systems, have parallel processes to handle asynchronous events. Neuromorphic computing architectures developed by researchers and major chip companies such as IBM and Intel are asynchronous and highly parallel, and they are composed of many computing units that act like neurons (Merolla et al., 2014; Davies et al., 2018). This architectural design also follows the *Efficiency through Cheap Design* principle described in Section 2.2. Neuromorphic hardware architectures use orders of magnitude less power than conventional computing by not relying on a synchronous clock and using spiking elements (Krichmar et al., 2019). A neuron uses most of its energy when it fires an action potential and when an action potential is processed at the synapse. Because neurons do not fire often (in the typical range of 10–100 Hz), the nervous system is in low-power mode between spikes. This idea was not lost on most neuromorphic hardware designers. Furthermore, communication bandwidth is reduced because information is sent only when there is a spike event.

3. ADAPTIVE BEHAVIOR - LEARNING AND MEMORY

Adaptation requires learning and remembering what was learned so that it can be applied in the future. Motivation is a key driver

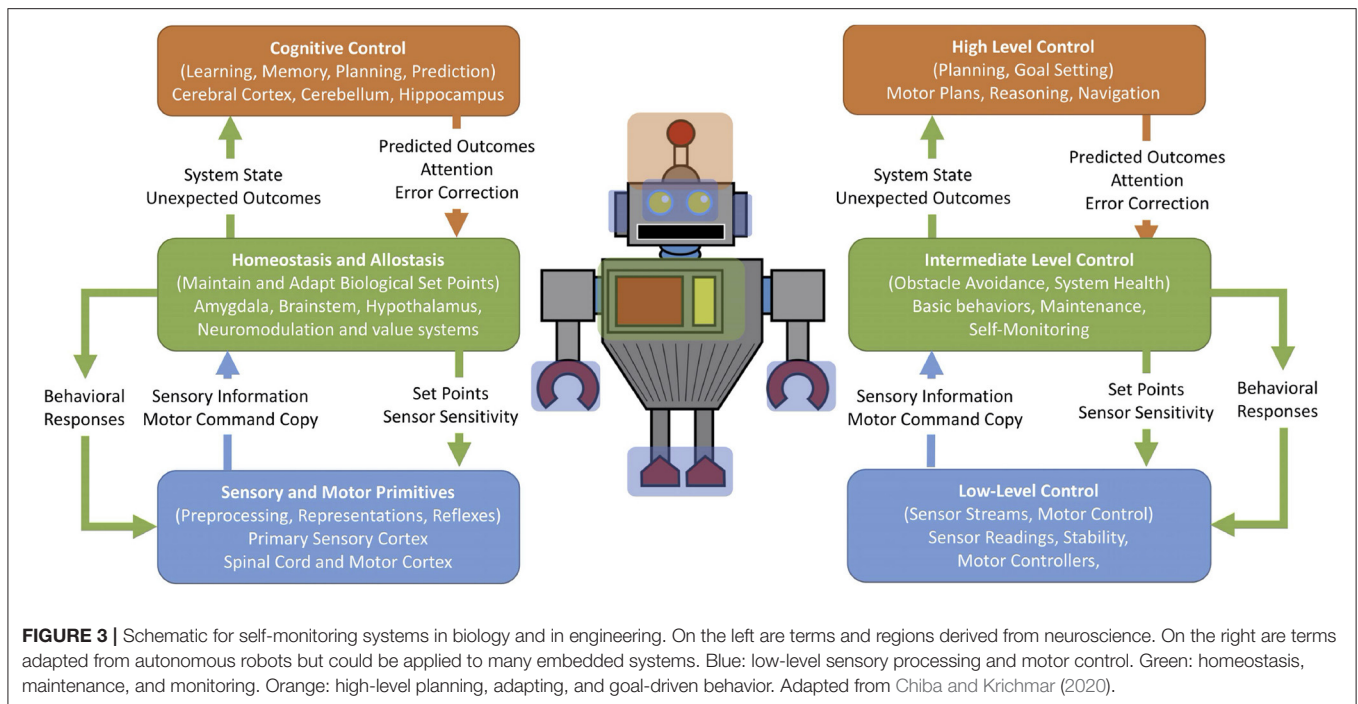
of learning. Motivators take many forms, which are called value systems. Another key aspect of adaptive behavior is the ability to predict future events. This requires building up a memory of expectations and the ability to adjust when expectations do not meet the current situation.

3.1. Learning and Memory

Unlike artificial systems, our brains allow us to learn quickly, incrementally, and continuously. With just a few presentations of something new, we can learn to recognize an object or situation or even learn a new skill. When we learn something new, we do not forget what we have learned previously. Moreover, we can take what we learn from one situation and apply it to another. On the other hand, artificial learning systems struggle under these situations, suffer from catastrophic forgetting of previous learning when something new is learned, and have difficulties generalizing learning from one task to another.

A brain region important for learning and memory is the hippocampus. The hippocampus is necessary to learn new memories and to consolidate those new experiences into long-term memories that can last a lifetime. The hippocampus can rapidly learn new autobiographical and semantic information, sometimes in the first experience (i.e., one-shot learning). Over time, this information becomes consolidated in the rest of the brain. Having a rapid learning system that can interact with a slower long-term storage area, which has been called complementary learning systems (McClelland et al., 1995; Kumaran et al., 2016), is thought to be the means by which our brains overcome catastrophic forgetting (i.e., forgetting previously learned information when learning new information). This aligns with another memory model, known as the hippocampal indexing theory (Teyler and DiScenna, 1986), which states that memories in the form of neocortical activation patterns are stored as indices in the hippocampus that are later used to aid recall. Although this may be an oversimplification, the notion that the hippocampus and medial temporal lobe integrates multimodal information from the neocortex makes sense and is backed by experimental evidence.

Our memories have context, and this contextual information can help us generalize when we encounter novel yet similar situations. In the literature this is called a schema, which is the memory of a set of items or actions bound together by a common context (van Kesteren et al., 2012). For example, if you are in a restaurant, you expect to see tables, chairs, a menu, waiters, and so forth. If you go to a new restaurant, that common context information can be used to rapidly consolidate the novel information into the restaurant schema. This requires mental representations that are flexible enough to learn tasks in new contexts and yet stable enough to retrieve and maintain tasks in old contexts (Carpenter and Grossberg, 1987). Tse et al. (2007) demonstrated this by training rats on different schemas, which were collections of associations between different foods and their locations in an enclosure. They found that the rats were able to learn new information quickly if it fit within a familiar schema. Additionally, the rats were able to learn new schemas without forgetting previous ones. The hippocampus was necessary for learning schemas and any new information matching a schema.



A subsequent study showed increased plasticity in the medial prefrontal cortex (mPFC) when information was consistent with a familiar schema (Tse et al., 2011). However, the hippocampus was not necessary to recall these memories, even after a short time period (e.g., 48 h). This challenged the idea of complementary learning systems because new information could rapidly be consolidated in cortical memory under these conditions.

The Tse et al. (2007) schema experiment was replicated with a neural network model based on interactions between the hippocampus and the mPFC (Hwu and Krichmar, 2020), and later tested on a robot required to create and utilize a schema (Hwu et al., 2020). A contextual pattern of objects and locations projected to the mPFC, in which each individual neuron encoded a different schema. The ventral hippocampus (vHPC) and dorsal hippocampus (dHPC) created triplet indices of schema, object, and location. The vHPC and dHPC drove activity that eventually activated a place cell through a winner-take-all process in the action network. This action neuron caused the robot to go to that location in search of the cued object. Contrastive Hebbian learning (CHL) was used to make the association between schema, object, and location. CHL utilizes oscillatory epochs, in which the duration depends on the familiarity and novelty of information, to learn associations. This assimilation of new information has similarities to adaptive resonance theory or ART (Grossberg, 2013). The model contained neuromodulators to encode novelty and familiarity. For example, if an object is novel and the context is unfamiliar, a new schema must be learned. However, if an object is novel and the context is familiar, the object can be added to an existing schema.

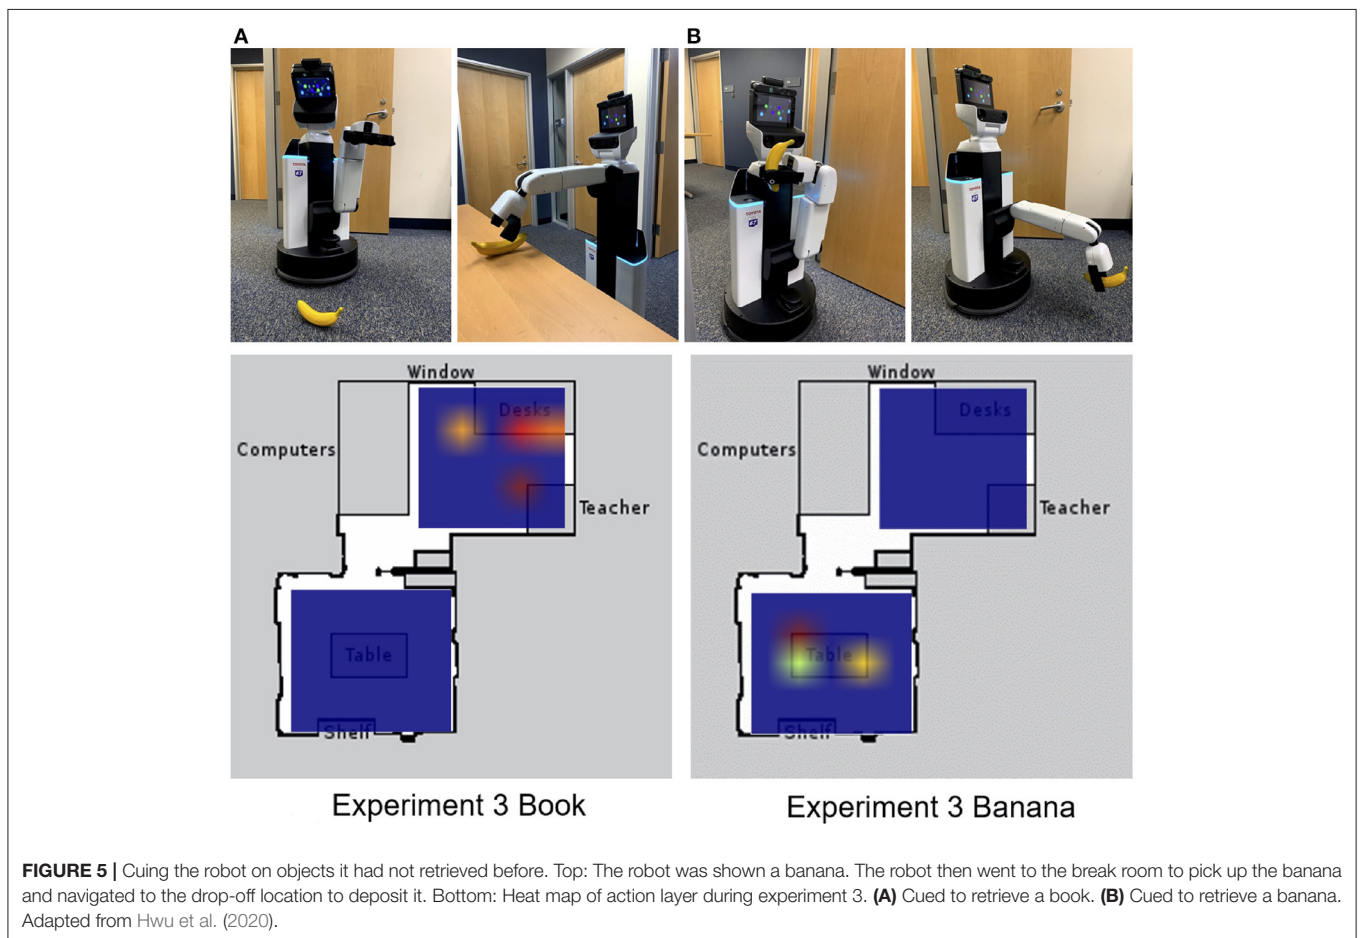
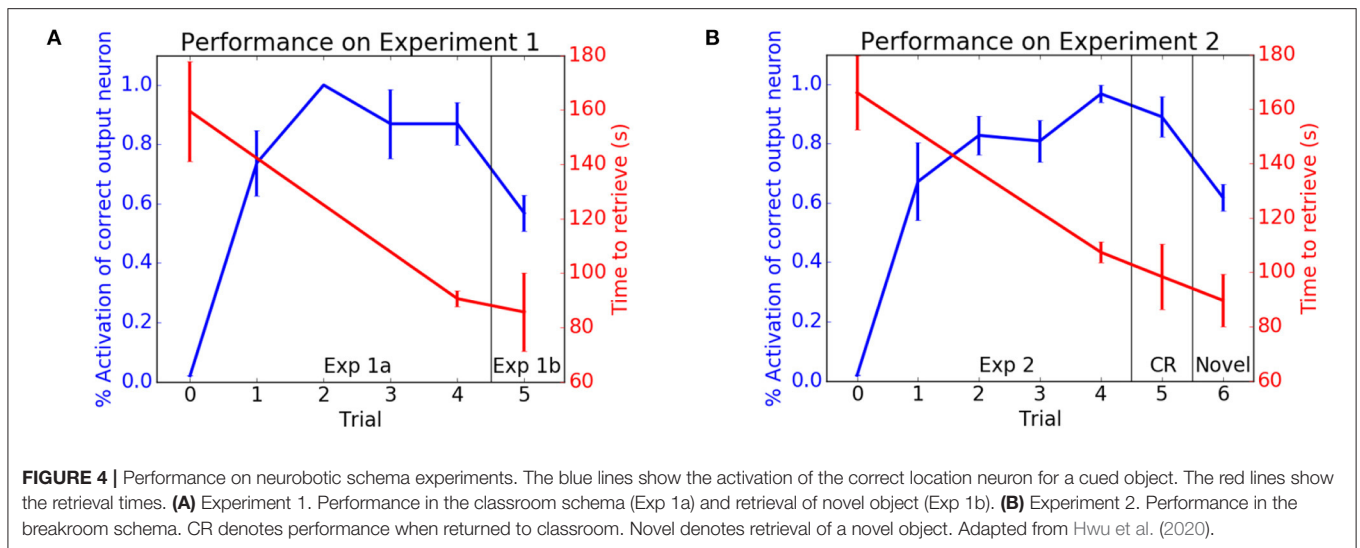
The schema model was embedded on the Human Support Robot (HSR) from Toyota (Yamamoto et al., 2018) and given

the task of finding and retrieving objects in a classroom and a break room. In a trial, the robot was prompted to retrieve an object, which required prior knowledge of the schema to which the object belonged and the location of the object. This caused the robot to navigate toward a location, recognize the object, grasp the object, and then return the object to its starting location.

In the first experiment, the HSR was placed in a room with typical classroom items (e.g., apple, bottle, computer mouse, book). **Figure 4A** shows the performance of the HSR retrieving classroom items. With each trial the number of correct places recalled increased and the time to retrieve an item decreased. After training and testing in the classroom, an original item was replaced by a novel item (Exp 1b in **Figure 4A**). Although the object was novel, the HSR knew it belonged in the classroom and was able to quickly consolidate this new information into the existing classroom schema.

In the second experiment, the HSR was placed in a room with typical break room items (e.g., apple, cup, banana, microwave oven). After training and testing, the classroom schema was tested again to see if the robot was able to maintain performance of prior tasks. As with the classroom, **Figure 4B** shows that the robot was able to learn this new break room schema without forgetting objects' locations in the classroom (CR in **Figure 4B**).

In the third experiment, the HSR was tested to see whether schemas could help with the retrieval of items that it was never explicitly trained to retrieve. If the HSR was cued with a book, it searched for the book on the desk in the classroom because books are likely to be found in a classroom schema (see **Figure 5A**). If the HSR was cued by showing it a banana, the HSR searched for the banana in the break room first rather than in the classroom (see **Figure 5B**).



The neurobotic schema experiments showed how context is tied to spatial representations via hippocampal interaction with the mPFC. Moreover, it demonstrates

how ideas from memory models in the brain may improve robotic applications and issues in artificial intelligence, such as catastrophic forgetting and lifelong learning. A robot

that has contextual memory could have applications for assistive technologies.

3.2. Value Systems

Robots should be equipped with a value system that constitutes a basic assumption of what is good and bad for its health and well-being. A value system facilitates the capacity of a biological brain to increase the likelihood of neural responses to an external phenomenon (Merrick, 2017). The combined effects of perception, experience, reasoning, and decision making contribute to the development of values in animals. Value can also be thought of as a measure of the effort one is willing to expend to obtain reward or to avoid punishment.

In addition to rewards and punishment, intrinsic motivation can be considered as a value system (Oudeyer and Kaplan, 2007). This can take the form of seeking novelty, fun, play or curiosity, that is, obtaining value for its own sake rather than to satisfy some need. For example, Oudeyer and colleagues created a robotic playground where a Sony AIBO dog with a motivation to explore and learn new things learned to manipulate objects (Oudeyer et al., 2007). The robot first spent time in situations that were easy to learn and then shifted its attention progressively to more difficult situations, avoiding situations in which nothing could be learned.

Neuromodulators are thought to act as the brain's value systems (Krichmar, 2008). Neuromodulators are chemicals that signal important environmental or internal events. They cause organisms to adapt their behavior through long-lasting signals to broad areas of the nervous system. Neuromodulators in the brain influence synaptic change (i.e., learning and memory) to satisfy global needs according to value.

To shape behavior, cognitive robots should have an innate value system to tell the robot that something was of value and trigger the appropriate reflexive behavior. From this experience the agent can learn which stimuli were predictive of that value and try to maximize the acquisition of good value while minimizing the acquisition of bad value. Many of these value-based robots employ models of the dopaminergic neuromodulatory system to shape behavior (Sporns and Alexander, 2002; Cox and Krichmar, 2009; Fiore et al., 2014; Chou et al., 2015).

Besides the dopaminergic reward system, there are multiple neuromodulators signal different value types (Doya, 2002; Krichmar, 2008). The serotonergic system is involved in harm aversion or impulsiveness (Miyazaki et al., 2018). The noradrenergic system signals oddball or unexpected events (Yu and Dayan, 2005). The cholinergic system is thought to increase attention to important features and at the same time to decrease the allocation of attention to distractors (Yu and Dayan, 2005). Acetylcholine and noradrenaline could be thought to signal intrinsic value by allocating attention and triggering learning (Avery and Krichmar, 2017). These systems interact with each other through direct and indirect pathways, and they all respond strongly to novelty by sending broad signals to large areas of the brain to cause a change in network dynamics resulting in decisive action.

Introducing saliency into the environment can lead to attentional signaling. For example, the robot CARL was designed to test a computational framework for applying neuromodulatory systems to the control of autonomous robots (Cox and Krichmar, 2009). The framework was based on the following premises (Krichmar, 2008): (1) the common effect of the neuromodulatory systems is to drive an organism to be decisive when environmental conditions call for such actions and allow the organism to be more exploratory when there are no pressing events; and (2) the main difference between neuromodulatory systems is the environmental stimuli that activate them. In the experiment, two out of four objects were salient, and CARL learned the appropriate action for each (see **Figure 6**). Unexpectedly, a strong attentional bias toward salient objects, along with ignoring the irrelevant objects, emerged through its experience in the real world. The selective attention could be observed both in CARL's behavior and in CARL's simulated brain. These neurobotic experiments showed how phasic neuromodulation could rapidly focus attention on important objects in the environment by increasing the signal-to-noise ratio (SNR) of neuronal responses. The model further suggested that phasic neuromodulation amplifies sensory input and increases competition in the neural network by gating inhibition.

The neuromodulatory system also regulates attention allocation and response to unexpected events. Using the Toyota Human Support Robot (Yamamoto et al., 2018), the influence of the cholinergic (ACh) system and noradrenergic (NE) systems on goal-directed perception was studied in an action-based attention task (Zou et al., 2020). In this experiment, a robot was required to attend to goal-related objects (the ACh system) and adjust to the change of goals in an uncertain domain (the NE system). Four different actions (i.e., eat, work-on-computer, read, and say-hi) were possible in the experiment and each of them was associated with different images of objects. For example, the goal action "eat" might result in attention to objects such as apple or banana, whereas the action "say-hi" should increase attention to a person. During the experiment, the goal action changed periodically and the robot needed to select the action and object that it thought the user wanted on the basis of prior experience. The ACh system tracked the expected uncertainty about which goal was valid, and the NE system signaled unexpected uncertainty when goals suddenly changed (Yu and Dayan, 2005). High ACh activity levels allocated attention to different goals. Phasic NE responses caused a rapid shift in attention and a resetting of prior goal beliefs. The model demonstrated how neuromodulatory systems can facilitate rapid adaptation to change in uncertain environments. The goal-directed perception was realized through the allocation of the robot's attention to the desired action/object pair. **Figure 7** shows the robot deciding which object to bring to the user. The bottom of **Figure 7** shows views from the robot's camera as it correctly guesses that the user's goal is to eat. Its top-down attention system finds an appropriate object, which is an apple in this case.

One problem that remains unsolved in neurobotics is that these artificial value systems are dissociated from the agent's

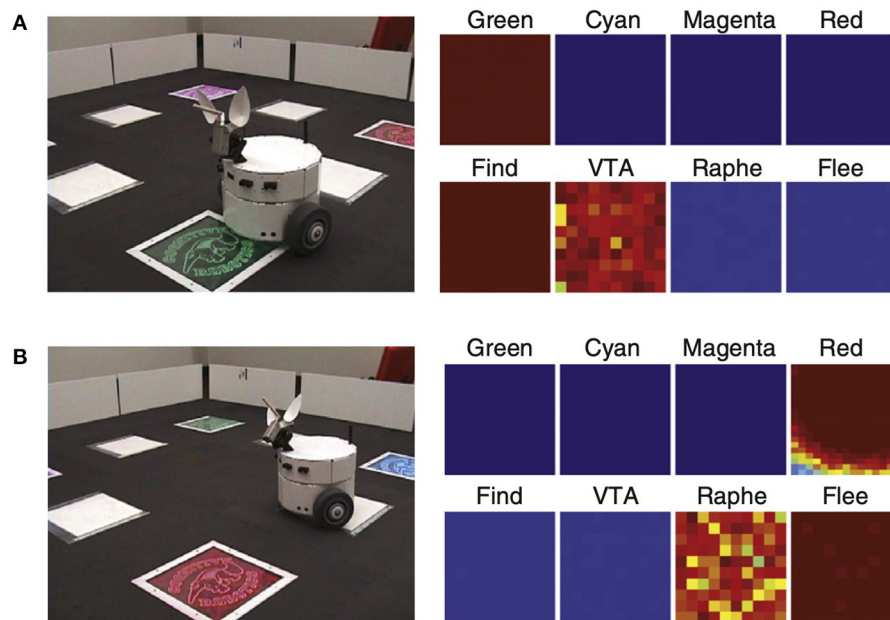


FIGURE 6 | CARL robot in colored panel task. The panels could flash any of six different colors. One color, green, signaled positive value. Another color, red, signaled negative value. The remaining colors were neutral. Positive and negative signals were transmitted from the panel to a receiver on the bottom of CARL. **(A)** CARL during an approach or find response. The panels on the right show strong neuronal activity in its simulated green visual area, the dopaminergic system (VTA), and the find motor neurons. **(B)** CARL during a withdrawal or flee response. The panels on the right show strong neuronal activity in its simulated red visual area, the serotonergic system (raphe), and the flee motor neurons. Adapted from Cox and Krichmar (2009).

body. Real pain, hunger, thirst, and fatigue drive a true value system. Without this connection to body-dependent drives, an artificial value system does not signal the immediacy of the agent's need and lacks to some degree the ability to shape the agent's behavior. It would be interesting to tie something like the robot's battery level to its hunger value. With faster-charging batteries or better solar cells this might be possible. An interesting step in this direction is the work on self-monitoring systems that can recognize drops in their performance, adapt their behavior, and recover. For example, Cully et al. (2015) developed a novel method for adapting gaits on a hexapod robot. In a sense, the robot controller had a memory of potential gaits. If one or more of the robot's legs were damaged, the robot would detect the damage, "imagine" different ways of moving, and then choose the new gait that it thought would work best under the new circumstances. In this way, the robot knew something was wrong and was able to adapt its behavior quickly without intervention.

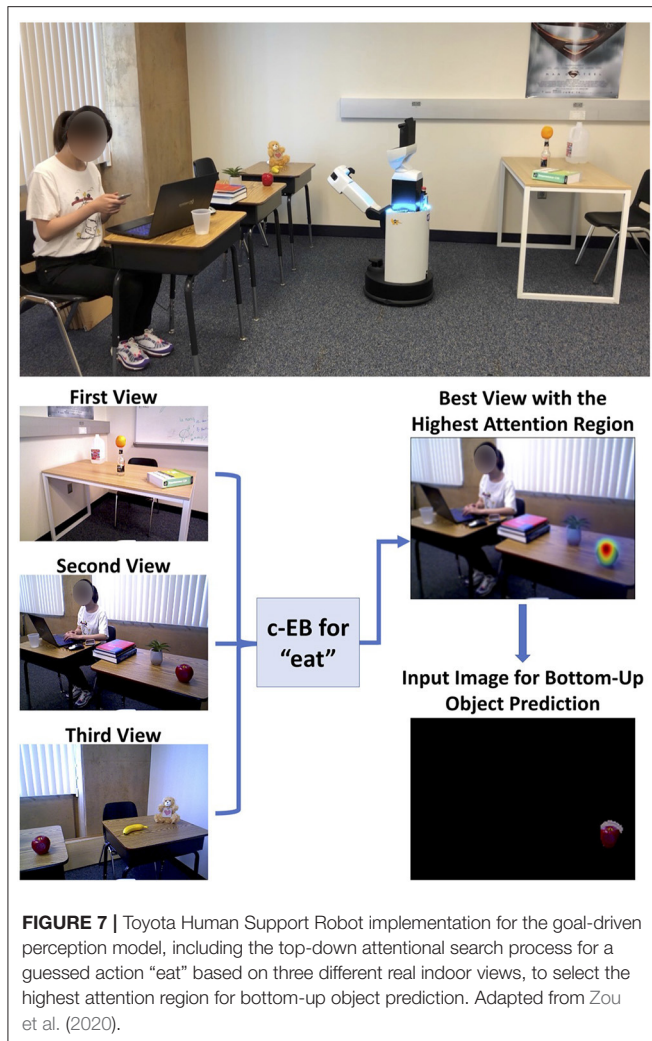
3.3. Prediction

Predicting outcomes and planning for the future is a hallmark of cognitive behavior. Much of the cortex is devoted to predicting what we will sense or the outcome of a movement or what series of actions will lead to big payoffs. Thus, a neurobot should strive to have these predictive capabilities. Through prediction and active inference, agents anticipate future sensory input based on prior experience. This minimizes free energy by predicting future outcomes so that they minimize the expenditures required to deal with unanticipated events (Friston, 2010). The idea

of minimizing free energy has close ties to many existing brain theories, such as Bayesian brain, Predictive Coding, cell assemblies, Infomax, and the theory of neuronal group selection (Edelman, 1993; Rao and Ballard, 1999; Friston, 2010). In the theory of neuronal group selection (Edelman, 1993), plasticity is modulated by value. Value systems control which neuronal groups are selected and which actions lead to evolutionary fitness; that is, they predict outcomes that lead to positive value and avoid negative value. In this sense predicting value is inversely proportional to surprise.

Prediction is crucial for fitness in a complex world and a fundamental computation in cortical systems (George and Hawkins, 2009; Clark, 2013; Richert et al., 2016). It requires the construction and maintenance of an internal model. In a similar fashion, model-based reinforcement learning builds an internal model made up of the likelihood and expected value for transitions between states (Solway and Botvinick, 2012).

Prediction has been useful in developing robot controllers. For example, in a humanoid robot experiment it was shown that having a predictive model helped the robot make appropriate reactive and proactive arm gestures (Murata et al., 2014). In the proactive mode the robot's actions were generated on the basis of top-down intentions to achieve intended goals. In the reactive mode the robot's actions were generated by bottom-up sensory inputs in unpredictable situations. In another robot experiment the combination of model-based and model-free reinforcement learning was used in a sorting task (Renaudo et al., 2015). The robot had to push cubes on a conveyor belt. The model-based



system improved performance by maintaining a plan from one decision to the next. However, the experiments suggested that the model-free system scales better under certain conditions and may be better in the face of uncertainty.

Jun Tani’s group has developed several predictive robot controllers using recurrent neural networks (Tani, 2016; Ahmadi and Tani, 2019; Chame et al., 2020). For example, they trained a hierarchy of continuous time recurrent neural networks (CTRNN) to learn different movements. Learning was achieved via backpropagation through time (BPTT), a means to apply an error signal to a sequence of neural activities. A teacher guided a humanoid Sony QRIO robot through different behavioral tasks. The CTRNNs received visual information and proprioceptive joint angles from the humanoid robot. Important to the learning were the different timescales of the CTRNNs. Slower higher-level CTRNNs sent predictions to the faster lower-level CTRNN units. Prediction errors from the lower levels were propagated to the higher levels for adjustments. Movements that appeared repeatedly were segmented into

behavioral primitives. These primitives were represented in fast context dynamics in a form that was generalized across object locations. On the other hand, the slow context units appeared to be more abstract in nature, representing sequences of primitives in a way that was independent of the object location. Tani (2016) speculated that this prediction multiple timescale hierarchy had similarities to the cortex. Fast responding motor primitives can be found in the primary motor cortex, and the slower prefrontal cortex sends top-down predictions to the primary motor cortex. Similarly, the primary visual cortex sends sensory information and prediction errors to the slower parietal cortex, which sends top-down predictions for the primary visual cortex. In this group’s recent work, they show the potential for prototyping robotics agents, modeled after active inference from the free energy principle theory (Friston, 2010), for human-robot interaction and socially assistive robotics (Chame et al., 2020).

4. BEHAVIORAL TRADE-OFFS - CONTEXTUAL DECISION-MAKING

Biological organisms need to consider many trade-offs to survive. These trade-offs regulate basic needs such as whether to forage for food, which might expose oneself to predators, or hide in one’s home, which is safer but does not provide sustenance. These trade-offs can be cognitive as in introverted or extroverted behavior. Interestingly, many of these trade-offs are regulated by chemicals in our brain and body, such as neuromodulators or hormones. These modulatory areas monitor and regulate environmental events. They send broad signals to the brain that can dramatically change behaviors, moods, decisions, etc. The brain can control these modulatory and hormonal systems by setting a context or making an adjustment when there are prediction errors (Chiba and Krichmar, 2020).

We discuss the neuroscience behind the trade-offs and neurobots that incorporate these trade-offs. We consider these behavioral trade-offs to be neurobotic design principles. By applying them to neurobots, we may realize behavior that is more interesting and more realistic.

4.1. Reward vs. Punishment

Dopaminergic neurons have phasic responses that match quite well with a reward prediction error signal used to shape behavior (Schultz et al., 1997). What about punishment? One model suggested that tonic serotonin tracked the average punishment rate and that tonic dopamine tracked the average reward rate (Daw et al., 2002). They speculated that a phasic serotonin signal might report an ongoing prediction error for future punishment. It has been suggested that the serotonergic and dopaminergic systems activate in opposition for goal-directed actions (Boureau and Dayan, 2011). This trade-off between reward and punishment can be quite nuanced when invigoration of activity can lead to rewards and punishment can lead to inaction.

4.2. Invigorated vs. Withdrawn

The dopamine and serotonin systems also regulate a trade-off between invigorated novelty seeking and withdrawn risk-averse behavior. It has been suggested that serotonin modulates the desire to withdraw from risk, which can take place in social interactions or in foraging for food (Tops et al., 2009). One can imagine that too much withdrawal from society could lead to symptoms of depression.

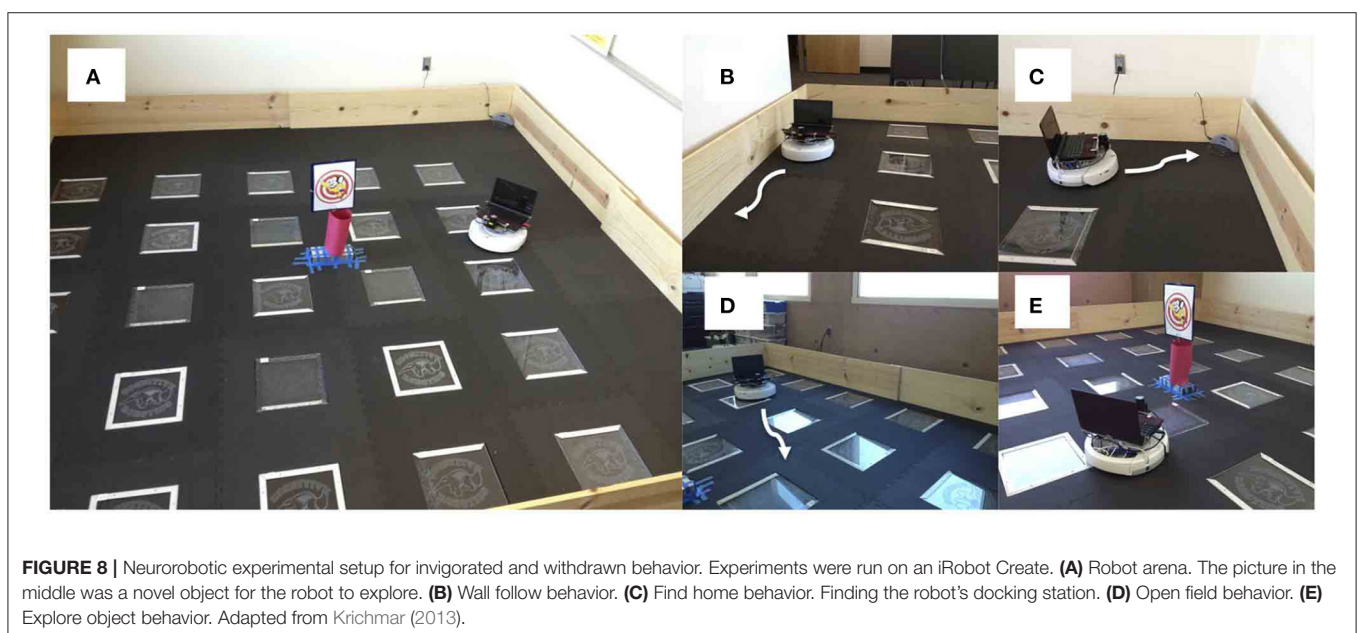
Consider the open-field test that is used to measure anxiety in rodents (Fonio et al., 2009). Usually, when a mouse is placed in an unfamiliar open area it will first stay near the borders of the environment in which it might be concealed. The mouse may hide in a nest area if available. After some time the mouse decides the environment is safe and becomes curious. The mouse will then proceed to explore the environment by moving more and investigating the middle of the area. Serotonin levels can alter this behavior. For example, Heisler et al. (1998) showed that mice with increased serotonin spent less time in the center of the open-field arena. In contrast, cocaine, which increases the level of dopamine, increased locomotive activity and the exploration of novel objects (Carey et al., 2008).

A neurorobot experiment took this trade-off into consideration by modeling the interactions between the serotonergic and dopaminergic systems (Krichmar, 2013). **Figure 8** shows the experimental setup and behaviors. A neural network controlled the behavior of an autonomous robot and tested in the open-field paradigm. When simulated serotonin levels were high, sensory events led to withdrawn anxious behavior such as wall following and finding its nest (i.e., the robot's charging station). When simulated dopamine levels were high, sensory events led to curious behavior such as locomotion to the middle of the enclosure or exploring a novel object.

The robot responded appropriately to sensory events in its environment. Novel objects resulted in its exploring the

environment, and stressful events caused the robot to seek safety. When the environment was unfamiliar, serotonergic activity dominated, resulting in anxious behavior such as WallFollow and FindHome actions. However, once the robot had become more familiar with its environment (approximately 60 s into the trial) DA levels were higher and there was more curious or exploratory behavior. At approximately 120 s into the trial, there was an unexpected light event due to flashing the lights on and off, which resulted in a phasic 5-HT response and a longer tonic increase in 5-HT. This caused the robot to respond with withdrawn or anxious behavior until approximately 210 s into the trial when a pair of object events triggered exploration of the center of the environment. Specifically, tonic levels of 5-HT had decayed, and the object events caused an increase in DA levels triggering a change in behavioral state (see **Figure 9A**). **Figure 9B** shows the proportion of curious behavior (OpenField and ExploreObject) and anxious behavior (FindHome and WallFollow) for five experimental trials. Each bar was the average proportion of time spent in either curious (green bars) or anxious (red bars) behaviors. The error bar denoted the standard error. **Figure 9C** shows the behavior time-locked to the light event. The light event, which occurred at approximately the halfway point in the trial, was introduced to cause a stress response. After the light event, the neurorobot's behavior rapidly switched to anxious behavior until roughly 60 s later when it became curious again. Variation occurred due to different times of the light event and random variations in other sensory events.

Similar to the rodent experiments, changing the serotonin and dopamine levels affected the robot's behavior. Increasing the tonic serotonin levels in the model caused the robot to respond to a stressful event such as a bright light to stay near the walls or its charging station indefinitely. Increasing the tonic DA levels resulted in more curiosity and risk taking. In effect, it took more risks by venturing into the middle of the environment during



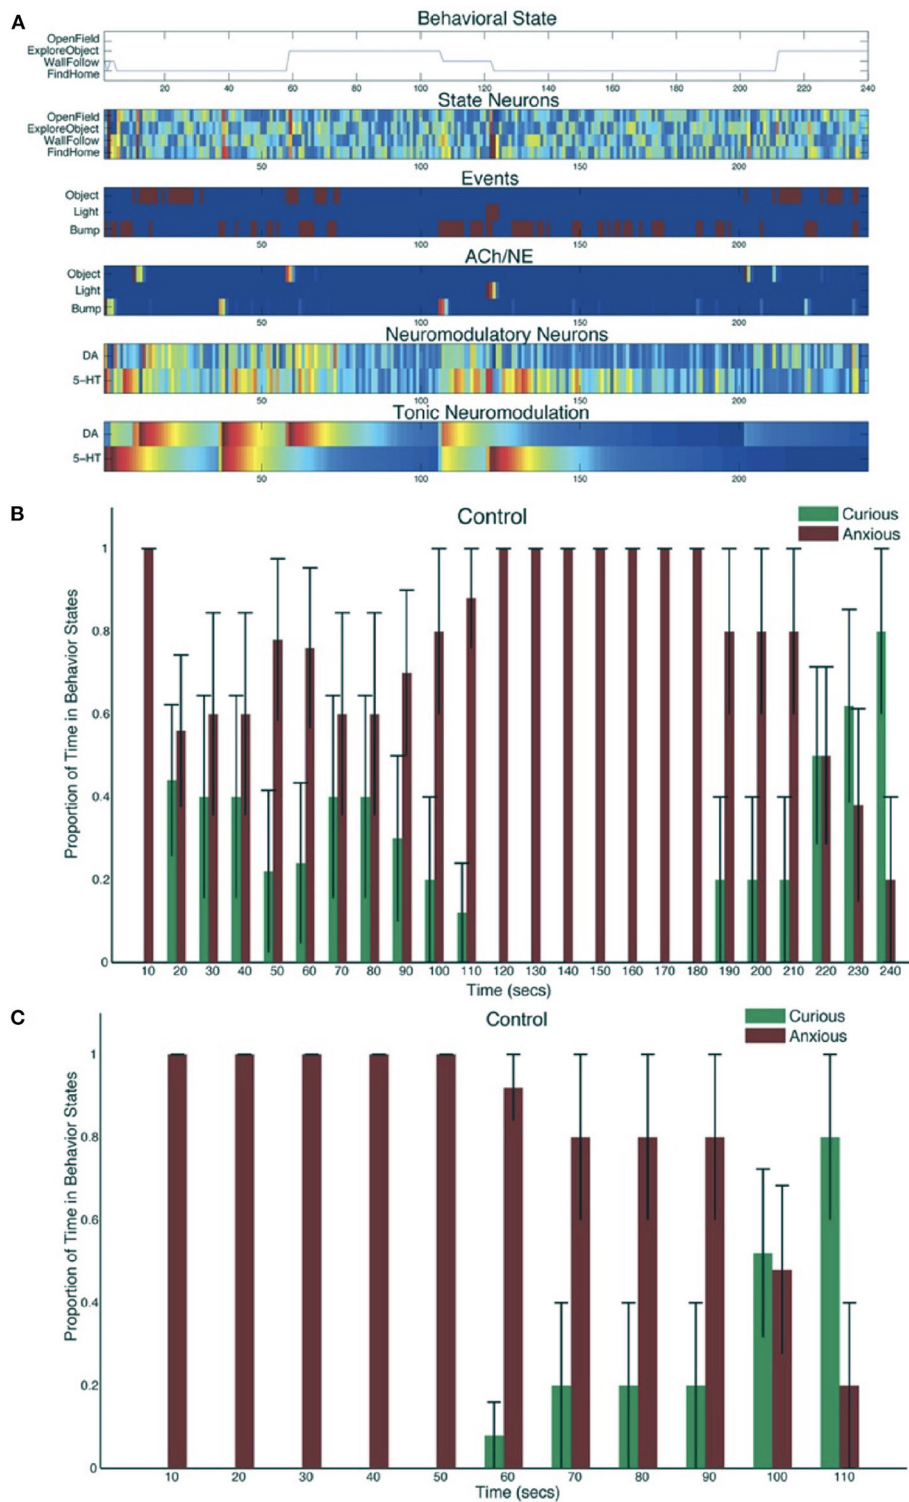


FIGURE 9 | Neurobot results for invigorated and withdrawn behavior. **(A)** A representative example of behavioral and neural responses in a single trial. The x-axis shows time progression of the trial in seconds. The subplot for “Behavioral State” denotes the state of the robot across time. The subplots for “State Neurons,” “Events,” “ACh/NE,” and “Neuromodulatory Neurons” show neural activity over the trial with pseudocolors ranging from dark blue (no activity) to bright red (maximal activity). The subplot for “Tonic Neuromodulation” denotes the level of tonic activation contributing to DA and 5-HT neurons. **(B)** Proportion of time in Curious (ExploreObject and OpenField) and Anxious (FindHome and WallFollow) behavior, averaged over 5 trials. Error bars represent standard error. **(C)** Similar to B except the behaviors were time-locked to the Light event. Adapted from Krichmar (2013).

or right after the stressful light event. Despite the simplicity of the neural network model, the robot's behavior looked quite natural and similar to that of a mouse in the same situation. Since the neuroroboticist was able to precisely control the neuromodulation activity, the experiment could shed light on neurological issues such as anxiety, depression, and obsessive compulsive disorders.

4.3. Expected Uncertainty vs. Unexpected Uncertainty

The world is full of uncertainty with which we must cope in our daily lives. Sometimes the uncertainty is expected, forcing us to increase our concentration on a task. Other times the uncertainty is unexpected, forcing us to divert our attention. How we deal with these types of uncertainty can be thought of as a behavioral trade-off. Once again, neuromodulators influence this trade-off of how we apply our attention. Yu and Dayan (2005) suggested that cholinergic neuromodulation tracks expected uncertainty (i.e., the known unreliability in the environment), and noradrenergic neuromodulation tracks unexpected uncertainty (i.e., observations that violate prior expectations). The basal forebrain, where cholinergic neurons reside, encodes the uncertainty of prior information and this can modulate attention to different features. The locus coeruleus, where noradrenergic neurons reside, is involved in cognitive shifts in response to novelty. When there are strong violations of expectations, locus coeruleus activity may induce a "network reset" that causes a reconfiguration of neuronal activity that clears the way to adapt to these changes (Bouret and Sara, 2005). In modeling and in experimental work, it has been shown that the cholinergic system mediates uncertainty seeking (Naude et al., 2016; Belkaid and Krichmar, 2020). Uncertainty seeking is especially advantageous in situations when reward sources are uncertain. The trade-off between expected and unexpected uncertainty can also be observed in how we apply our attention (Avery et al., 2012). Top-down attention or goal-driven attention, which ramps up our attention to look for something, is related to expected uncertainty. Bottom-up or stimulus-driven attention occurs when a surprise or unexpected uncertainty diverts our attention.

The Toyota HSR neurobot experiment discussed in Section 3.2 and shown in **Figure 7** explored this trade-off between expected and unexpected uncertainty by modeling the cholinergic and noradrenergic system ability to regulate attention (Zou et al., 2020). Because the user's goals could be uncertain, simulated cholinergic neurons tracked how likely the user would be to choose any of these goals (i.e., expected uncertainty). When the user interacting with the robot changed their goals (i.e., unexpected uncertainty), the noradrenergic system in the model responded by resetting prior beliefs and rapidly adapting to the new goal. **Figure 10** shows how the robot correctly guessed the user's goals, which then drove attention to the object associated with the goal (e.g., eat leads to attention to an apple or orange). Note how quickly the noradrenergic (NE) system responded to goal changes, which led to the cholinergic system (ACh) increasing attention to objects related to the new goal. In this

way, the robot was able to monitor a trade-off between the known and unknown uncertainties in the world to rapidly respond to the user's changing needs. Similar to the example in the schema experiments described in Section 3.1, such goal-driven attention could provide benefits for assistive robot technologies.

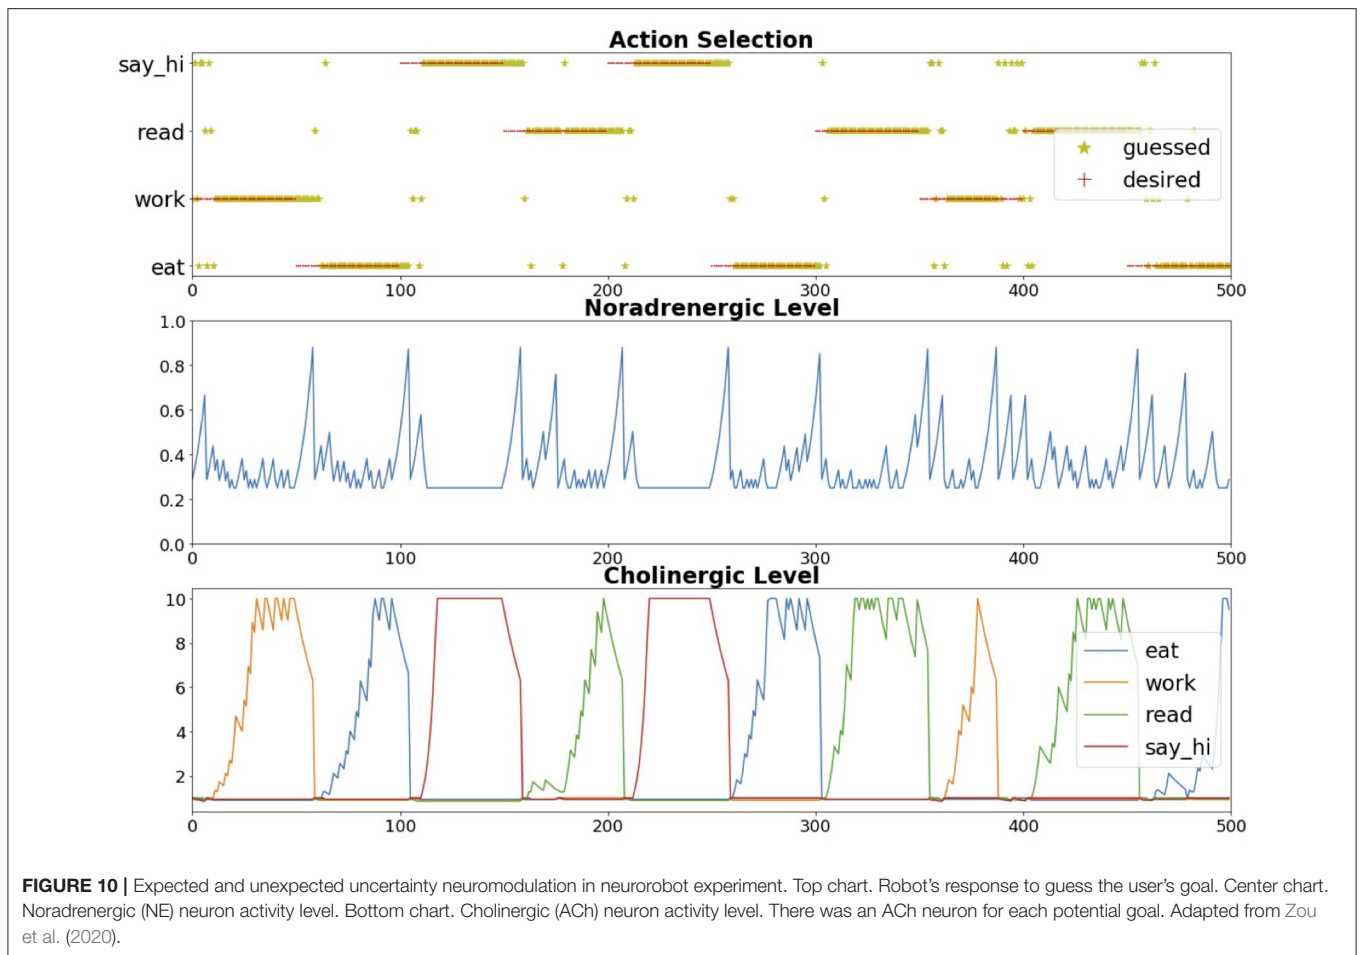
4.4. Exploration vs. Exploitation

During decision making or information gathering, there exists a trade-off between exploration and exploitation in which it is sometimes best to *explore* new options and other times it is best to *exploit* opportunities that have paid off in the past. A framework was presented in which neuromodulation controlled the exploration/exploitation trade-off (Aston-Jones and Cohen, 2005). When neuromodulators have tonic activity, the animal's behavior is exploratory and somewhat arbitrary. However, when the neuromodulator has a burst of phasic activity, the animal is decisive and exploits the best potential outcome at the given time. The CARL robot described in Section 3.2 (see **Figure 6**) incorporated tonic and phasic neuromodulation (Cox and Krichmar, 2009). When there was tonic neuromodulation, the robot randomly explored its environment by looking at different colored panels. If one of the colors became salient due to a reward or punishment signal, the robot's neuromodulatory systems responded with a phasic burst of activity. This phasic neuromodulation caused a rapid exploitative response to investigate the colored panel. It also triggered learning to approach positive-value objects and avoid negative-value objects.

4.5. Foraging vs. Defending

Hormones are chemical messengers in the body that can affect the brain and other organs. They regulate a number of bodily functions such as body temperature, thirst, hunger, and sleep. Like neuromodulators, hormones can be triggered by environmental events and can broadly change neural activity. For example, the hormone orexin regulates hunger levels. This can lead to a behavioral trade-off in which animals foraging for food are less willing to defend a territory (Padilla et al., 2016). Foraging for food may cause an animal to leave its nest exposed to predators. However, defending one's territory requires energy expenditure, which if prolonged requires food for replenishment.

Hormones can be modeled and embodied in robots to explore interesting naturalistic tradeoffs (Canamero, 1997). For example, Cañamero's group modeled hormones that tracked a robot's health (Lones et al., 2018); one hormone was related to the battery level, and another hormone monitored the robot's internal temperature, which was related to how much the robot moved and the climate of its environment. The robot's tasks were to maintain health and gather food resources, which might require aggressive action. This neurorobotics study demonstrated how maintaining health requires behavioral trade-offs. Searching for food increased the robot's internal temperature and reduced the robot's battery level. Being aggressive to obtain food also reduced battery levels. However, not searching for food would lead to starvation. Modeling the secretion and decay of hormones allowed the robot to maintain a comfortable energy and internal temperature level and at times led to an aggressive behavior of pushing objects away to get at food resources. However,



as the environments became more complex an epigenetic system, which monitored and controlled the hormones, became necessary for the robot's comfort level to be maintained satisfactorily. Their epigenetic system acted in a similar way to the hypothalamus, a subcortical brain region that regulates many of our bodily functions. Their experiments showed that an epigenetic mechanism significantly and consistently improved the robot's adaptability and might provide a useful general mechanism for adaptation in autonomous robots.

4.6. Stress vs. Calm

In his book *Why Zebras Don't Get Ulcers*, Robert Sapolsky describes how a zebra, which is calmly grazing, responds when it encounters a lion (Sapolsky, 2004). The zebra quickly runs away from this stressful situation. Once clear of danger, the zebra is calm again. This fight-or-flight response is mediated by the stress hormones known as glucocorticoids, which increase blood flow and awareness. However, this stress response does come at the expense of regulating long-term health and short-term memory (Chiba and Krichmar, 2020). Unlike zebras, people sometimes remain in a constant state of stress due to elevated glucocorticoids, which can cause damage to the hippocampus and memory.

Although there has been little work to date on neurorobots that regulate their stress level, downregulation of behavior could be useful for autonomous systems far from power sources, which might have a stress-like response to carry out a mission and then switch to a low-power calm mode after the mission has been accomplished. In an interesting paradigm that explores the stress vs. calm trade-off, experiments have shown that rats appear to be capable of empathy and prosocial behavior (Ben-Ami Bartal et al., 2011). In one study a rat was trapped in a cage and clearly stressed. Another rat observing this behavior became stressed, too. The observing rat, feeling bad for its trapped friend, found a lever that opened the cage and released the trapped rat. This study suggested that rats can feel another's pain (i.e., feel empathy) and are willing to act on the other's behalf (i.e., can be prosocial).

In a robotic variation of the empathy experiment, a rat was trapped in a cage interacted with two different robots, one of which was helpful and opened the cage and the other of which was uncooperative and ignored the trapped rat (Quinn et al., 2018). Interestingly, the rat remembered who its robot friends were. When the helpful robot was trapped, the rat freed that robot but did not free the robot that was uncooperative. This could have implications for rescue robotics. A robot that can identify and relieve stress or

anxiety could have applications for robotic caretakers or for disaster relief.

4.7. Social vs. Solitary

Hormones can regulate a trade-off between social bonding and independence. Estrogen, progesterone, oxytocin, and prolactin can influence a number of neural systems to ensure maternal nurturing, bonding, and protection of young (Rilling and Young, 2014). In particular, oxytocin has been shown to regulate social and paternal bonding (Young and Wang, 2004). An interesting example of oxytocin's effect on bonding has been observed in voles. Whereas, prairie voles are polygamous and the male does not assist in the nurturing of young pups, meadow voles are monogamous and both parents participate in the pup rearing. Interestingly, meadow voles have more oxytocin receptors than prairie voles. Furthermore, inhibiting oxytocin prevents pair bonding in meadow voles (Young and Wang, 2004). However, social bonding requires devoting energy to another, possibly at the expense of one's own health. Therefore, it could be argued that there should be a balance between social and solitary behavior.

Neurobot experimenters have investigated the balance between social and solitary behavior. By simulating social hormones, Cañamero's group investigated attachment bonds between a robot and a "caregiver" (Canamero et al., 2006; Hiolle et al., 2009, 2012). Although they did not explicitly model oxytocin, their system did simulate the type of bonding observed between a parent and a child. The robot found a good balance between asking the caregiver for help and learning on its own. Too much interaction with a caregiver led to stress and rejection by the robot. Not enough interaction with a caregiver resulted in isolation. As with humans, a proper balance is important for learning and development.

5. DISCUSSION

In this article, we discussed a number of principles to consider when designing neurobots and experiments using robots to test brain theories. These principles are strongly inspired by Pfeifer and Bongard's design principles for intelligent agents. We build upon these design principles by grounding them in neurobiology and by adding principles based on neuroscience research. We highlight the importance in neurobotics for designing systems that are reactive, adaptive, predictive, able to manage behavioral tradeoffs, and capable of learning from experience.

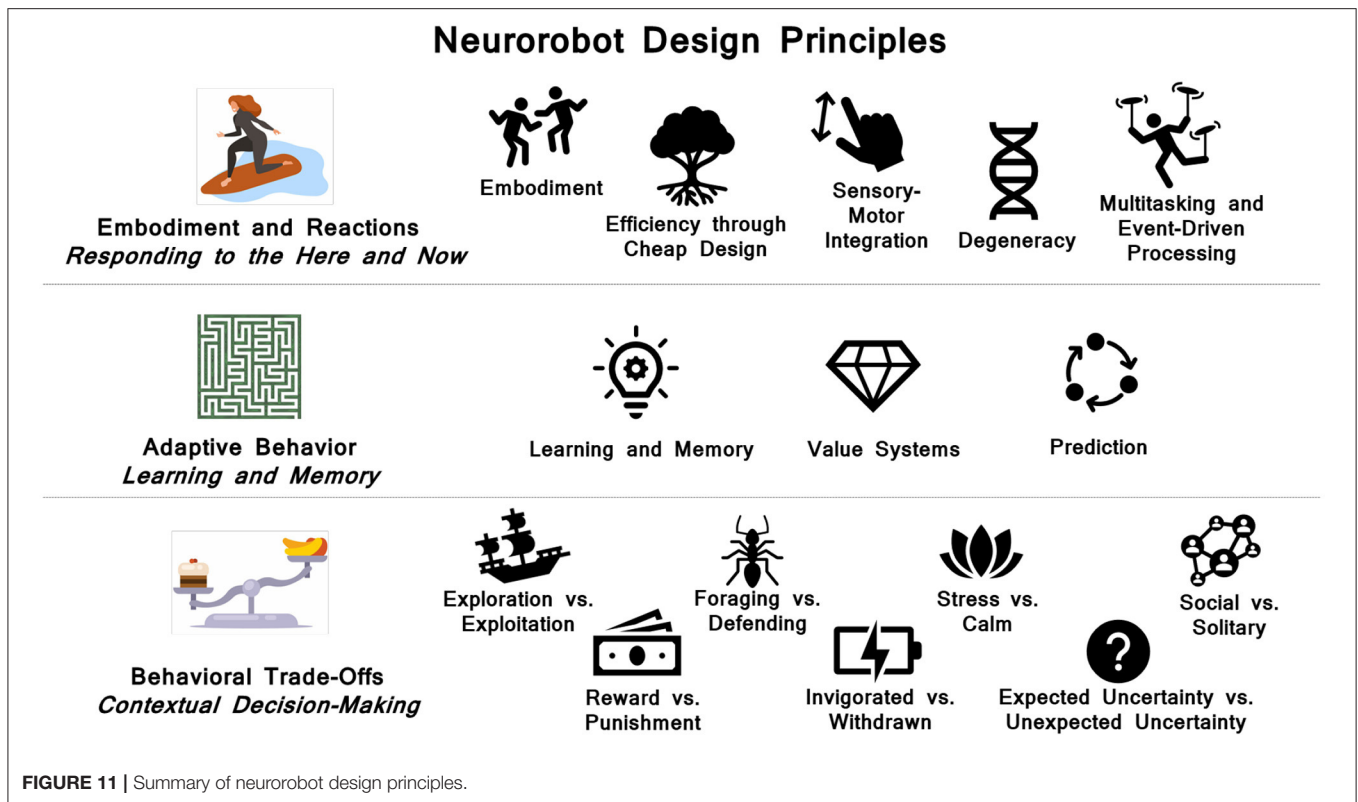
As summarized in **Figure 11**, the principles fall into three broad categories: 1) *Embodiment and reactions*. These are reactive, reflexive and rapid responses. They are often carried out without involving the central nervous system. Rather they emerge through the body's interaction with the environment or are handled by the peripheral nervous system and reflex arcs involving the spinal cord. They can have short-term adaptive properties and lead to behavioral repertoires. 2) *Adaptive behavior*. Biological organisms have the ability to learn continually over the lifetime of the organism. A major property of the brain is its plasticity. In particular, hippocampal interactions with the cortex lead to long-term contextual memory. Neuromodulatory systems can signal value, which

shapes learning and triggers adaptive behavior. Another hallmark of the brain is its ability to predict outcomes. This requires the construction, maintenance and updating of memory systems. 3) *Behavioral trade-offs*. To survive in a dynamic world, organisms must make decisions based on their needs and environmental context. Oftentimes, these needs are a trade-off between opposing motivations (e.g, taking a risk for a reward vs. playing it safe to avoid punishment). These trade-offs can lead to interesting behavior. Many of these trade-offs are regulated by sub-cortical neuromodulator and hormone levels.

5.1. Importance of Low-Level Processes and Model Organisms

Although the examples in this article focused primarily on vertebrates, the principles could be applied to modeling other organisms. Studies of the insect visual system have led to elegant, efficient solutions for robot navigation that could be deployed on neuromorphic hardware (Galluppi et al., 2014; Schoepe et al., 2021). The emphasis on vertebrates and especially the mammalian brain has been data-driven in part. There are numerous studies of rodents, non-human primates and humans that have provided modelers with anatomical, behavioral, and neurophysiological data to make their simulations more biologically accurate. However, the complexity of these organisms and their brains makes holistic modeling difficult. A promising avenue may be to study organisms that have less complex brains and behaviors. Recent work on model organisms such as *Drosophila* and the nematode *C. Elegans* are providing rich data sets for neuroboticists (Sarma et al., 2018; Scheffer et al., 2020). The OpenWorm project (Sarma et al., 2018) provides the biological data and the simulation tools, including a robotics sub-project, to create interesting neurobots that follow many of the design principles introduced here.

There is a tendency in neurobotics, which is a subarea of cognitive robotics (Cangelosi and Asada, 2022), to model cognitive functions such as attention, decision-making, planning, etc. Although modeling human cognition may be an ultimate goal for the field, many of the design principles introduced here concentrate on low-level processes such as how the body shapes behavior, motivations, homeostatic control of body functions to name a few. As we emphasized in this paper and in Chiba and Krichmar (2020), many of these processes are driven by sub-cortical brain regions and neurochemicals. Their interaction with the environment and bodily functions lead to interesting behavior that could be described as cognitive. Early examples of neurobots and behavior-based robots demonstrated that intelligent behavior could arise from interactions between the robot and the environment without complex nervous systems or control systems (Braitenberg, 1986; Brooks, 1991; Holland, 2003). More work needs to be done to first build a foundation of these low-level processes, upon which higher-order cognitive processes can be added. Moreover, each cognitive model should carefully choose an appropriate level of abstraction and state assumptions about the lower level processes that support the model.



5.2. Next Steps and Future Directions

The principles of neurorobotics introduced here can address major challenges facing artificial intelligence and robotics research. In general, neurorobotics explores these challenges in embodied settings, providing a fresh perspective that extends beyond simulations and algorithms performed on computers. Neurorobotics incorporates features of the brain that may begin to address lifelong continual learning, efficient computing, operating on scarce knowledge, and human-computer interaction.

We expect that new neuroscience discoveries will further inform neurobots, and vice versa. Progress in learning and memory may lead to applications capable of continual learning. Advances in our understanding of multimodal sensory systems may be incorporated into neurorobotics that not only classify but also understand the meaning of what they are perceiving. Given the recent achievements of artificial intelligence, hybrid systems combining machine learning and deep learning with neurobotic design principles could lead to interesting applications in autonomous driving, assistive robots, and manufacturing.

In the intermediate term, we believe that neurobiological concepts in learning and memory, navigation, decision making, social behavior, and more, will have found their way into practical applications. Progress in neuromorphic computing and

algorithms will lead to applications that can run at the edge with little human intervention. This may lead to advances in search and rescue robots and in robots capable of autonomously exploring unknown environments such as the deep sea or extraterrestrial planets.

In the long-term we hope that neurorobotics will achieve more general intelligence rather than being designed for specific tasks. In fact the delineation between conventional robotics and neurorobotics may be blurred, with all robots possessing some neurobiologically inspired aspects. With the rapid advances in computing and other technologies, it is hard to predict far into the future. However, we do believe that neurorobotics and cognitive machines, in some form, will seamlessly be a part of our everyday lives.

6. CONCLUSION

In closing, we believe that neurobots that follow many of the design principles discussed in this article will have more interesting, naturalistic behavior. This not only allows the robot to be a better model for understanding the complex behaviors observed in biology, it also could lead to better robots that show more intelligence and that are more natural in their interactions with other agents.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

REFERENCES

- Ahmadi, A., and Tani, J. (2019). A novel predictive-coding-inspired variational rnn model for online prediction and recognition. *Neural Comput.* 31, 2025–2074. doi: 10.1162/neco_a_01228
- Arkin, R. (1998). *Behavior-Based Robotics*. Cambridge, MA: MIT Press.
- Ascoli, G. A., Donohue, D. E., and Halavi, M. (2007). Neuromorpho.org: a central resource for neuronal morphologies. *J. Neurosci.* 27, 9247–9251. doi: 10.1523/JNEUROSCI.2055-07.2007
- Aston-Jones, G., and Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450. doi: 10.1146/annurev.neuro.28.061604.135709
- Avery, M. C., and Krichmar, J. L. (2017). Neuromodulatory systems and their interactions: a review of models, theories, and experiments. *Front. Neural Circ.* 11, 108. doi: 10.3389/fncir.2017.00108
- Avery, M. C., Nitz, D. A., Chiba, A. A., and Krichmar, J. L. (2012). Simulation of cholinergic and noradrenergic modulation of behavior in uncertain environments. *Front. Comput. Neurosci.* 6, 5. doi: 10.3389/fncom.2012.00005
- Belkaid, M., and Krichmar, J. L. (2020). Modeling uncertainty-seeking behavior mediated by cholinergic influence on dopamine. *Neural Netw.* 125, 10–18. doi: 10.1016/j.neunet.2020.01.032
- Ben-Ami Bartal, I., Decety, J., and Mason, P. (2011). Empathy and pro-social behavior in rats. *Science* 334, 1427–1430. doi: 10.1126/science.1210789
- Beyeler, M., Rounds, E. L., Carlson, K. D., Dutt, N., and Krichmar, J. L. (2019). Neural correlates of sparse coding and dimensionality reduction. *PLoS Comput. Biol.* 15, e1006908. doi: 10.1371/journal.pcbi.1006908
- Boureau, Y. L., and Dayan, P. (2011). Opponency revisited: competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology* 36, 74–97. doi: 10.1038/npp.2010.151
- Bouret, S., and Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenergic function. *Trends Neurosci.* 28, 574–582. doi: 10.1016/j.tins.2005.09.002
- Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press. A Bradford Book.
- Brooks, R. A. (1991). Intelligence without representation. *Artif. Intell.* 47, 139–159. doi: 10.1016/0004-3702(91)90053-M
- Canamero, D. (1997). “Modeling motivations and emotions as a basis for intelligent behavior,” in *Proceedings of the First International Conference on Autonomous Agents* (Marina del Rey, CA), 148–155.
- Canamero, L., Blanchard, A. J., and Nadel, J. (2006). Attachment bonds for human-like robots. *Int. J. Humanoid Rob.* 3, 301–320. doi: 10.1142/S0219843606000771
- Cangelosi, A., and Asada, M. (2022). *Cognitive Robotics*. Cambridge, MA: The MIT Press.
- Carey, R. J., Damianopoulos, E. N., and Shanahan, A. B. (2008). Cocaine effects on behavioral responding to a novel object placed in a familiar environment. *Pharmacol. Biochem. Behav.* 88, 265–271. doi: 10.1016/j.pbb.2007.08.010
- Carpenter, G. A., and Grossberg, S. (1987). Art 2: self-organization of stable category recognition codes for analog input patterns. *Appl. Opt.* 26, 4919–4930. doi: 10.1364/AO.26.004919
- Chame, H. F., Ahmadi, A., and Tani, J. (2020). A hybrid human-neurobotics approach to primary intersubjectivity via active inference. *Front. Psychol.* 11, 584869. doi: 10.3389/fpsyg.2020.584869
- Chiba, A., and Krichmar, J. (2020). Neurobiologically inspired self-monitoring systems. *Proc. IEEE* 108, 976–986. doi: 10.1109/JPROC.2020.2979233
- Chiel, H. J., and Beer, R. D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends Neurosci.* 20, 553–557. doi: 10.1016/S0166-2236(97)01149-1
- Chou, T. S., Bucci, L. D., and Krichmar, J. L. (2015). Learning touch preferences with a tactile robot using dopamine modulated stdp in a model of insular cortex. *Front. Neurobot.* 9, 6. doi: 10.3389/fnbot.2015.00006
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204. doi: 10.1017/S0140525X12000477
- Collins, S., Ruina, A., Tedrake, R., and Wisse, M. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science* 307, 1082–1085. doi: 10.1126/science.1107799
- Cox, B. R., and Krichmar, J. L. (2009). Neuromodulation as a robot controller: a brain inspired design strategy for controlling autonomous robots. *IEEE Rob. Autom. Mag.* 16, 72–80. doi: 10.1109/MRA.2009.933628
- Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature* 521, 503–507. doi: 10.1038/nature14422
- Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Daw, N. D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Netw.* 15, 603–616. doi: 10.1016/S0893-6080(02)00052-7
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw.* 15, 495–506. doi: 10.1016/S0893-6080(02)00044-8
- Edelman, G. M. (1993). Neural darwinism: selection and reentrant signaling in higher brain function. *Neuron* 10, 115–125. doi: 10.1016/0896-6273(93)90304-A
- Edelman, G. M., and Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proc. Natl. Acad. Sci. U.S.A.* 98, 13763–13768. doi: 10.1073/pnas.231499798
- Fiore, V. G., Sperati, V., Mannella, F., Miroli, M., Gurney, K., Friston, K., et al. (2014). Keep focussing: striatal dopamine multiple functions resolved in a single mechanism tested in a simulated humanoid robot. *Front. Psychol.* 5, 124. doi: 10.3389/fpsyg.2014.00124
- Fitzpatrick, P., and Metta, G. (2003). Grounding vision through experimental manipulation. *Philos. Trans. A Math. Phys. Eng. Sci.* 361, 2165–2185. doi: 10.1098/rsta.2003.1251
- Fleischer, J., Szatmary, B., Hutson, D., Moore, D., Snook, J., Edelman, G., et al. (2006). “A neurally controlled robot competes and cooperates with humans in segway soccer,” in *IEEE International Conference on Robotics and Automation* (Orlando, FL: IEEE), 3673–3678.
- Fleischer, J. G., Gally, J. A., Edelman, G. M., and Krichmar, J. L. (2007). Retrospective and prospective responses arising in a modeled hippocampus during maze navigation by a brain-based device. *Proc. Natl. Acad. Sci. U.S.A.* 104, 3556–3561. doi: 10.1073/pnas.0611571104
- Fonio, E., Benjamini, Y., and Golani, I. (2009). Freedom of movement and the stability of its unfolding in free exploration of mice. *Proc. Natl. Acad. Sci. U.S.A.* 106, 21335–21340. doi: 10.1073/pnas.0812513106
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 11, 127–138. doi: 10.1038/nrn2787

FUNDING

This work was supported by NSF award IIS-1813785 and by the United States Air Force Award FA9550-19-1-0306.

ACKNOWLEDGMENTS

The authors would like to thank the present and former members of the CARL lab at UC Irvine for many fruitful discussions and inspiring robot experiments.

- Fuster, J. M. (2004). Upper processing stages of the perception-action cycle. *Trends Cogn. Sci.* 8, 143–145. doi: 10.1016/j.tics.2004.02.004
- Gallagher, S. (2001). The practice of mind: theory, simulation or interaction? *J. Consciousness Stud.* 8, 83–108. Available online at: <https://www.ingentaconnect.com/contentone/imp/jcs/2001/00000008/f0030005/1207#>
- Galluppi, F., Denk, C., Meiner, M. C., Stewart, T. C., Plana, L. A., Eliasmith, C., et al. (2014). “Event-based neural computing on an autonomous mobile platform,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)* (Hong Kong: IEEE), 2862–2867.
- Gaussier, P., Banquet, J. P., Cuperlier, N., Quoy, M., Aubin, L., Jacob, P.-Y., et al. (2019). Merging information in the entorhinal cortex: what can we learn from robotics experiments and modeling? *J. Exp. Biol.* 222, 1–13. doi: 10.1242/jeb.186932
- George, D., and Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. *PLoS Comput. Biol.* 5, e1000532. doi: 10.1371/journal.pcbi.1000532
- Grossberg, S. (2013). Adaptive resonance theory: how a brain learns to consciously attend, learn, and recognize a changing world. *Neural Netw.* 37, 1–47. doi: 10.1016/j.neunet.2012.09.017
- Heisler, L. K., Chu, H. M., Brennan, T. J., Danao, J. A., Bajwa, P., Parsons, L. H., et al. (1998). Elevated anxiety and antidepressant-like responses in serotonin 5-HT_{1A} receptor mutant mice. *Proc. Natl. Acad. Sci. U.S.A.* 95, 15049–15054. doi: 10.1073/pnas.95.25.15049
- Hickok, G., Houde, J., and Rong, F. (2011). Sensorimotor integration in speech processing: computational basis and neural organization. *Neuron* 69, 407–422. doi: 10.1016/j.neuron.2011.01.019
- Hiolle, A., Bard, K. A., and Canamero, L. (2009). “Assessing human reactions to different robot attachment profiles,” in *Ro-Man 2009: The 18th IEEE International Symposium on Robot and Human Interactive Communication, Vols. 1, 2* (Toyama), 824–829.
- Hiolle, A., Canamero, L., Davila-Ross, M., and Bard, K. A. (2012). “Eliciting caregiving behavior in dyadic human-robot attachment-like interactions,” in *ACM Transactions on Interactive Intelligent Systems, Vol. 2*, Article 3.
- Holland, O. (2003). Exploration and high adventure: the legacy of grey walter. *Philos. Trans. A Math. Phys. Eng. Sci.* 361, 2085–2121. doi: 10.1098/rsta.2003.1260
- Hwu, T., Kashyap, H. J., and Krichmar, J. L. (2020). “A neurobiological schema model for contextual awareness in robotics,” in *2020 International Joint Conference on Neural Networks (IJCNN)* (Glasgow), 1–8.
- Hwu, T., and Krichmar, J. L. (2020). A neural model of schemas and memory encoding. *Biol. Cybern.* 114, 169–186. doi: 10.1007/s00422-019-00808-7
- Hwu, T. J., and Krichmar, J. L. (2022). *Neurorobotics: Connecting the Brain, Body, and Environment*. Cambridge, MA: The MIT Press.
- Krichmar, J. L. (2008). The neuromodulatory system - a framework for survival and adaptive behavior in a challenging world. *Adapt. Behav.* 16, 385–399. doi: 10.1177/1059712308095775
- Krichmar, J. L. (2013). A neurobotic platform to test the influence of neuromodulatory signaling on anxious and curious behavior. *Front. Neurobot.* 7, 1. doi: 10.3389/fnbot.2013.00001
- Krichmar, J. L., Nitz, D. A., Gally, J. A., and Edelman, G. M. (2005). Characterizing functional hippocampal pathways in a brain-based device as it solves a spatial memory task. *Proc. Natl. Acad. Sci. U.S.A.* 102, 2111–2116. doi: 10.1073/pnas.0409792102
- Krichmar, J. L., Severa, W., Khan, M. S., and Olds, J. L. (2019). Making bread: biomimetic strategies for artificial intelligence now and in the future. *Front. Neurosci.* 13, 666. doi: 10.3389/fnins.2019.00666
- Kumaran, D., Hassabis, D., and McClelland, J. L. (2016). What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends Cogn. Sci.* 20, 512–534. doi: 10.1016/j.tics.2016.05.004
- Laughlin, S. B., and Sejnowski, T. J. (2003). Communication in neuronal networks. *Science* 301, 1870–1874. doi: 10.1126/science.1089662
- Lones, J., Lewis, M., and Cañamero, L. (2018). A hormone-driven epigenetic mechanism for adaptation in autonomous robots. *IEEE Trans. Cogn. Dev. Syst.* 10, 445–454. doi: 10.1109/TCDS.2017.2775620
- McClelland, J. L., McNaughton, B. L., and O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* 102, 419–457. doi: 10.1037/0033-295X.102.3.419
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). Artificial brains. a million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Merrick, K. (2017). Value systems for developmental cognitive robotics: a survey. *Cogn. Syst. Res.* 41, 38–55. doi: 10.1016/j.cogsys.2016.08.001
- Milford, M., Jacobson, A., Chen, Z., and Wyeth, G. (2016). Ratslam: Using models of rodent hippocampus for robot navigation and beyond. *Rob. Res. ISRR* 114, 467–485. doi: 10.1007/978-3-319-28872-7_27
- Miyazaki, K., Miyazaki, K. W., Yamanaka, A., Tokuda, T., Tanaka, K. F., and Doya, K. (2018). Reward probability and timing uncertainty alter the effect of dorsal raphe serotonin neurons on patience. *Nat. Commun.* 9, 2048. doi: 10.1038/s41467-018-04496-y
- Murata, S., Arie, H., Ogata, T., Sugano, S., and Tani, J. (2014). Learning to generate proactive and reactive behavior using a dynamic neural network model with time-varying variance prediction mechanism. *Adv. Rob.* 28, 1189–1203. doi: 10.1080/01691864.2014.916628
- Naude, J., Tolu, S., Dongelmans, M., Torquet, N., Valverde, S., Rodriguez, G., et al. (2016). Nicotinic receptors in the ventral tegmental area promote uncertainty-seeking. *Nat. Neurosci.* 19, 471–478. doi: 10.1038/nn.4223
- Olshausen, B., and Field, D. (2004). Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* 14, 481–487. doi: 10.1016/j.conb.2004.07.007
- Oudeyer, P., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evolutionary Comput.* 11, 265–286. doi: 10.1109/TEVC.2006.890271
- Oudeyer, P. Y., and Kaplan, F. (2007). What is intrinsic motivation? a typology of computational approaches. *Front. Neurobot.* 1, 6. doi: 10.3389/neuro.12.006.2007
- Padilla, S. L., Qiu, J., Soden, M. E., Sanz, E., Nestor, C. C., Barker, F. D., et al. (2016). Agouti-related peptide neural circuits mediate adaptive behaviors in the starved state. *Nat. Neurosci.* 19, 734–741. doi: 10.1038/nn.4274
- Pfeifer, R., and Bongard, J. (2006). *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: The MIT Press.
- Prescott, T. J., Redgrave, P., and Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adapt. Behav.* 7, 99–127. doi: 10.1177/105971239900700105
- Quinn, L. K., Schuster, L. P., Aguilar-Rivera, M., Arnold, J., Ball, D., Gygi, E., et al. (2018). When rats rescue robots. *Anim. Behav. Cogn.* 5, 368–379. doi: 10.26451/abc.05.04.04.2018
- Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580
- Renaudo, E., Girard, B., Chatila, R., and Khamassi, M. (2015). Respective advantages and disadvantages of model-based and model-free reinforcement learning in a robotics neuro-inspired cognitive architecture. *Procedia Comput. Sci.* 71(Suppl. C), 178–184. doi: 10.1016/j.procs.2015.12.194
- Richert, M., Fisher, D., Piekiewicz, F., Izhikevich, E. M., and Hylton, T. L. (2016). Fundamental principles of cortical computation: unsupervised learning with prediction, compression and feedback. *arXiv:1608.06277*. doi: 10.48550/arXiv.1608.06277
- Rilling, J. K., and Young, L. J. (2014). The biology of mammalian parenting and its effect on offspring social development. *Science* 345, 771. doi: 10.1126/science.1252723
- Sapolsky, R. M. (2004). *Why Zebras Don’t Get Ulcers: The Acclaimed Guide to Stress, Stress-Related Disease and Coping, 3rd Edn*. New York, NY: Holt Paperbacks.
- Sarma, G. P., Lee, C. W., Portegys, T., Ghayoomi, V., Jacobs, T., Alicea, B., et al. (2018). Openworm: overview and recent advances in integrative biological simulation of caenorhabditis elegans. *Philos. Trans. R Soc. Lond. B Biol. Sci.* 373, 1758. doi: 10.1098/rstb.2017.0382
- Scheffer, L. K., Xu, C. S., Januszewski, M., Lu, Z., Takemura, S. Y., Hayworth, K. J., et al. (2020). A connectome and analysis of the adult drosophila central brain. *Life* 9, e57443. doi: 10.7554/eLife.57443.sa2
- Schoepe, T., Janotte, E., Milde, M., Bertrand, O., Egelhaaf, M., and Chicca, E. (2021). Finding the gap: neuromorphic motion vision in cluttered environments. *arXiv:2102.08417 [cs.NE]*. doi: 10.48550/arXiv.2102.08417
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593

- Shadmehr, R., and Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Exp. Brain Res.* 185, 359–381. doi: 10.1007/s00221-008-1280-5
- Solway, A., and Botvinick, M. M. (2012). Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychol. Rev.* 119, 120–154. doi: 10.1037/a0026435
- Sporns, O. (2010). *Networks of the Brain*. Cambridge, MA: MIT Press.
- Sporns, O., and Alexander, W. H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Netw.* 15, 761–774. doi: 10.1016/S0893-6080(02)00062-X
- Swanson, L. W. (2007). Quest for the basic plan of nervous system circuitry. *Brain Res. Rev.* 55, 356–372. doi: 10.1016/j.brainresrev.2006.12.006
- Tani, J. (2016). *Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. New York, NY: Oxford University Press.
- Teyler, T. J., and DiScenna, P. (1986). The hippocampal memory indexing theory. *Behav. Neurosci.* 100, 147–154. doi: 10.1037/0735-7044.100.2.147
- Tops, M., Russo, S., Boksem, M. A., and Tucker, D. M. (2009). Serotonin: modulator of a drive to withdraw. *Brain Cogn.* 71, 427–436. doi: 10.1016/j.bandc.2009.03.009
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., et al. (2007). Schemas and memory consolidation. *Science* 316, 76–82. doi: 10.1126/science.1135935
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., et al. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science* 333, 891–895. doi: 10.1126/science.1205274
- van Kesteren, M. T., Ruiters, D. J., Fernandez, G., and Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends Neurosci.* 35, 211–219. doi: 10.1016/j.tins.2012.02.001
- Wheeler, D. W., White, C. M., Rees, C. L., Komendantov, A. O., Hamilton, D. J., and Ascoli, G. A. (2015). Hippocampome.org: a knowledge base of neuron types in the rodent hippocampus. *Elife* 4, e09960. doi: 10.7554/eLife.09960
- Yamamoto, T., Nishino, T., Kajima, H., Ohta, M., and Ikeda, K. (2018). “Human support robot (hsr),” in *Siggraph'18: Acm Siggraph 2018 Emerging Technologies* (Vancouver, BC).
- Young, L. J., and Wang, Z. (2004). The neurobiology of pair bonding. *Nat. Neurosci.* 7, 1048–1054. doi: 10.1038/nn1327
- Yu, A. J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692. doi: 10.1016/j.neuron.2005.04.026
- Ziemke, T. (2003). “What’s that thing called embodiment?” in *Proceedings of the Annual Meeting of the Cognitive Science Society, Vol. 25* (Boston, MA), 1305–1310.
- Zou, X., Kolouri, S., Pilly, P. K., and Krichmar, J. L. (2020). Neuromodulated attention and goal-driven perception in uncertain domains. *Neural Netw.* 125, 56–69. doi: 10.1016/j.neunet.2020.01.031

Conflict of Interest: TH was employed by Riot Games.

The remaining author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Krichmar and Hwu. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.