The Neuromodulatory System: A Framework for Survival and Adaptive Behavior in a Challenging World

Jeffrey L. Krichmar

Adaptive Behavior 2008; 16; 385
DOI: 10.1177/1059712308095775

Adaptive **Behavior**

The Neuromodulatory System: A Framework for Survival and Adaptive Behavior in a Challenging World

Jeffrey L. Krichmar
Department of Cognitive Sciences, University of California, Irvine, USA

Biological organisms have the ability to respond quickly to an ever-changing world. Because this adaptability is so critical for survival, all vertebrates have sub-cortical structures, which comprise the neuromodulatory systems, to regulate fundamental behavior and drive decision making in response to environmental events. In the vertebrate, there are separate neuromodulators that respond to threats, reward anticipation, novelty, and attentional effort. However, each of these neuromodulatory systems has a similar effect, that is, to cause 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. In this article, it is proposed that principles of the neuromodulatory system could provide a framework for controlling artificial agents that may improve current artificial agent behavior. These agents would operate autonomously, effectively explore their environment, and be decisive when environmental conditions call for action.

Keywords autonomous systems · decision making · exploration/exploitation · neuromodulation · neurorobotics

1 Introduction

Vertebrates have sub-cortical structures, known as neuromodulatory systems, which regulate fundamental behavior, set the organism's internal states, and are critical for an organism's survival. When an important environmental event occurs, it is the neuromodulatory system that triggers the organism to respond quickly and accurately to that event. There are separate neuromodulators that alter responses to risks, rewards, novelty, and effort. Moreover, the neuromodulatory systems provide the foundation for cognitive function in higher organisms. Attention, emotion, goal-directed behavior, and decision making all derive from the interaction between the neuromodulatory systems and areas such as the amygdala, frontal cortex, and hippocampus. Therefore, understanding neuromodulatory function may provide a basis for the construction of cognitive machines and the control of autonomous systems.

Although there have been great advances in autonomous systems (Cho, 2007; Gibbs, 2004; Squyres, 2005; Yenne, 2004), the controllers of these machines are still very much tailored to specific missions and do not have the behavioral repertoire we normally associate with that of biological organisms. Behavior-based robotics (Jones & Roth, 2003) do not learn from experience and cannot adapt to environmental change. Probabilistic robot controllers (Thrun, Burgarde, &

Correspondence to: Jeffrey L. Krichmar, Department of Cognitive Sciences, SSPA 2175, 3151 Social Science Plaza, University of California, Irvine, CA 92697-5100, USA.
E-mail: jkrichma@uci.edu
Tel.: +1 949 824 5888; *Fax*: +1 949 824 2307

(2008), Vol 16(6): 385–399.
DOI: 10.1177/1059712308095775
Figures 1, 4 appear in color online: http://adb.sagepub.com

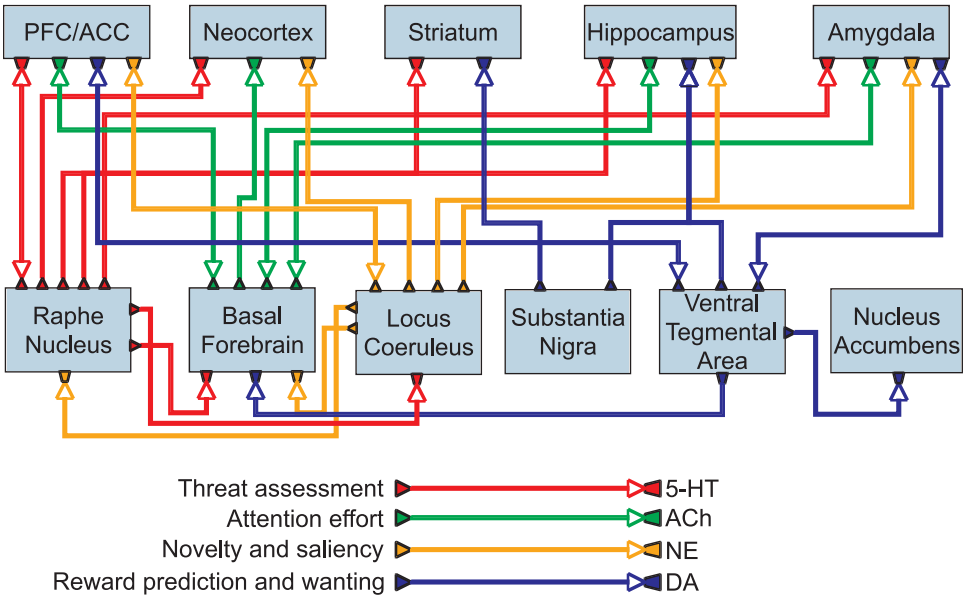

Figure 1 Architecture of the neuromodulatory systems. The raphe nucleus is the source of serotonin (5-HT), the basal forebrain is the source of acetylcholine (ACh), the locus coeruleus is the source of norepinephrine (NE), and the substantia nigra and ventral tegmental area are the sources of dopamine (DA). The prefrontal cortex (PFC), anterior cingulate cortex (ACC), neocortex, striatum, hippocampus, and amygdala are reciprocally connected with the neurmodulatory systems. Arrows projecting back to the neuromodulatory systems depict glutamatergic connections. 5-HT, shown in red, is thought to signal threat assessment. ACh, shown in green, is thought to signal attentional effort. NE, shown in light orange, is thought to signal vigilance and arousal. DA, shown in blue, is thought to signal reward prediction and wanting.

Fox, 2005) need an accurate model of their sensors and actuators. Robots controlled by reinforcement learning (Sutton & Barto, 1998) adapt their behavior through a reward prediction error (Doya & Uchibe, 2005; Guenter, Hersch, Calinon, & Billard, 2007; Iida, Kuwayama, Kanoh, Kato, & Itoh, 2004; Kondo & Ito, 2004; Nakamura, Mori, Sato, & Ishii, 2007; Stone, Sutton, & Kuhlmann, 2005). Some of these reinforcement learning robot controllers are neurally inspired by reward expectation signals found in the dopaminergic system of the brain (Alexander & Sporns, 2002; Arleo, Smeraldi, & Gerstner, 2004; Krichmar & Edelman, 2002; Sporns & Alexander, 2002). However, these models and robots do not address other aspects of adaptive behavior, such as attention, novelty, and threat assessment.

In this article, I present a framework for designing autonomous systems that is based on principles of the vertebrate neuromodulatory system (see Figure 1). Ascending neuromodulatory systems include noradrenergic, serotonergic, dopaminergic, and cholinergic projections from the brainstem and basal forebrain regions to broad areas of the of the central nervous system, the thalamus, and hypothalamus (Briand, Gritton, Howe, Young, & Sarter, 2007). Each of these neuromodulatory systems consists of small pools of neurons (on the order of thousands in the rodent and tens of thousands in the human) located in the brainstem, pontine nucleus, and basal forebrain.

Despite the different origins and chemical signatures of neuromodulatory systems, there are several commonalities among them:

1. The origination of these systems is sub-cortical.
2. Each of these neuromodulatory systems is the locus of a particular chemical transmitter that is projected to broad areas of the brainstem, thalamus, and cortex.
3. All of these neuromodulatory systems are reciprocally connected with the frontal cortex and parts of the limbic system.
4. The effect of these neuromodulatory systems on downstream targets is similar.

From the evidence, it appears that the common effect of the neuromodulatory system is to increase the

signal-to-noise ratio of downstream neuronal targets such that the organism can make quick and decisive choices. Indeed, the major targets of the neuromodulators shown in Figure 1 are areas noted for driving behavior (K. C. Berridge, 2004), conditioning responses (McGaugh, 2004), attention (Sarter, Gehring, & Kozak, 2006), and making decisions (Schall, 2001; Smith & Ratcliff, 2004). In general, the effect of activation of the neuromodulatory system on post-synaptic targets is to increase responses to stimuli and suppress responses to noise (Aston-Jones & Cohen, 2005).

Furthermore, it will be argued that the main differences between neuromodulatory systems are the environmental stimuli that activate them. For example, the serotonergic (5-HT) system appears to be driven by stress or threats (Millan, 2003), the cholinergic (ACh) system by attentional effort (Baxter & Chiba, 1999), the dopaminergic (DA) system by reward anticipation (Schultz, Dayan, & Montague, 1997) and "wanting" (K. C. Berridge, 2004), and the noradrenergic (NE) system by novelty and saliency (C. W. Berridge & Waterhouse, 2003; Yu & Dayan, 2005).

The main purpose of this article is to review the vertebrate neuromodulatory system with the goal of showing how principles of neuromodulation could be used as inspiration for the control of autonomous systems. In the remainder of this article, each neuromodulatory system will be described separately in more detail. This will be followed by an illustrative demonstration of how the action of neuromodulation could affect downstream neuronal targets, which in turn affect behavior. Finally, I will present a computational framework based on the neuromodulatory system, and briefly describe how such a framework might be applied to the control of an artificial agent.

2 Neuromodulatory Systems

The neuromodulators 5-HT, ACh, DA, and NE all originate in the brainstem and basal forebrain (see Figure 1). All of these neuromodulators appear to be important for arousal, but in different ways that are highly specific to cognitive functions (Briand et al., 2007; Robbins et al., 1998). Moreover, because these neuromodulatory systems have reciprocal projections to the amygdala and forebrain structures, neuromodulators can influence "cognitive" networks, and in turn be influenced by "cognitive" areas (see Figure 1).

2.1 Threat Assessment: Raphe Nucleus and Serotonin

Serotonergic projections, which originate in the raphe nuclei of the brainstem, extend to almost all forebrain areas (Barnes & Sharp, 1999). The expression of 5-HT occurs throughout the cortex, ventral striatum, hippocampus, and amygdala (Harvey, 2003; Meneses & Perez-Garcia, 2007). The raphe receives connections from the prefrontal cortex and the anterior cingulate cortex (Briand et al., 2007). The 5-HT receptors are found on both excitatory pyramidal cells and inhibitory interneurons. 5-HT receptors are also found on ACh and glutamatergic (Glu) axon terminals where an increase of ACh and Glu could enhance learning, acquisition, and consolidation (Meneses & Perez-Garcia, 2007).

Release of 5-HT appears to be related to cognitive control of stress. The structures that are innervated by 5-HT and their connecting circuits modulate the behavioral response to threats and risks—behaviors that are typically thought to reflect the anxiety state of the organism (Millan, 2003). Recent work by Amat and colleagues has provided support for the role of the prefrontal–raphe feedback projection in the cognitive control of stress (Amat, Paul, Zarza, Watkins, & Maier, 2006).

2.2 Attentional Effort: Basal Forebrain and Acetylcholine

Acetylcholine (ACh) originates from the basal forebrain and projects to the cortex, amygdala, and hippocampus. Basal forebrain projections are distinct bands that could be area, modality, and task specific. At the cortical level, the medial prefrontal and orbitofrontal cortices, which receive projections from the basal forebrain, share reciprocal connections with the basal forebrain. These feedback loops to the basal forebrain have been hypothesized to be a component of the neuronal mechanisms that serve to enhance input processing and the allocation of attentional resources to behaviorally significant stimuli under challenging conditions (Sarter et al., 2006).

Basal forebrain cholinergic neurons appear to modulate attention and optimize information processing (Baxter & Chiba, 1999). Removal of cholinergic projections to the parietal and frontal cortex impairs the ability to increase attentional effort (Bucci, Hol-

land, & Gallagher, 1998). Another interpretation is that ACh release is related to the expected uncertainty—the known unreliability of predictive relationships in the environment (Yu & Dayan, 2005). It appears that the basal forebrain increments attentional processing when prediction accuracy is reduced, but not when predictive value is held constant (Chiba, Bucci, Holland, & Gallagher, 1995). It has also been shown that the central nucleus of the amygdala has a direct projection to the basal forebrain that mediates increased attention as a consequence of expectancy violation (Han, Holland, & Gallagher, 1999).

2.3 Reward Prediction and Wanting: Ventral Tegmental Area and Dopamine

Dopamine (DA) is produced by two groups of cell bodies in the mesencephalon: the substantia nigra (SN) and the ventral tegmental area (VTA). The VTA projects to the nucleus accumbens (NAc) and is the pathway implicated in mediating reward related behaviors (Hyman, Malenka, & Nestler, 2006). The SN is the source of dopamine in the basal ganglia. Both the SN and VTA project to the hippocampus (Scatton, Simon, Le Moal, & Bischoff, 1980). The VTA also projects to the frontal cortex and predominantly innervates the medial prefrontal cortex (Fluxe et al., 1974). Similar to the other neuromodulatory systems, the VTA receives input from the prefrontal cortex (Sesack & Pickel, 1992).

Current theories posit that dopamine is important for facilitating learning through the influence of prediction (Schultz, 1997; Schultz et al., 1997), and for incentive salience or "wanting" (K. C. Berridge, 2004). "Wanting" refers to the motivation process in acquiring an object, which differs from "liking" where the pleasure is derived from the object itself (K. C. Berridge, 2004). A recent proposal ties the prediction error to wanting by suggesting that incentive salience is the expected future reward that maps actions to rewards (McClure, Daw, & Montague, 2003). Alternatively, it has been proposed that DA is involved with the discovery of new actions and it influences action–outcome contingencies (Redgrave & Gurney, 2006). In all of these variants, it appears that DA is an important signal for the acquisition of salient, value-laden objects.

2.4 Novelty and Saliency: Locus Coeruleus and Norepinephrine

Norepinephrine (NE) in the central nervous system is produced by the locus coeruleus, which projects to virtually all brain regions with the exception of basal ganglia (C. W. Berridge & Waterhouse, 2003). Similar to other neuromodulatory systems, the prefrontal cortex activates the locus coeruleus (C. W. Berridge & Waterhouse, 2003). In particular, the orbitofrontal cortex, which is related to the evaluation of reward, and the anterior cingulate cortex, which is related to the evaluation of cost, project to the locus coeruleus and drive phasic responses (Aston-Jones & Cohen, 2005). There is a feedback loop where the amygdala affects stress hormones, which then act on the nucleus of the solitary tract, which then acts on the locus coeruleus, which then releases NE in the amygdala. Norepinephrine activation in amygdala helps to consolidate and modulate memory in other brain regions (McGaugh, 2004).

Norepinephrine and the locus coeruleus are sensitive to novel and salient objects in the environment and task relevant stimuli that cannot be fully predicted (e.g., recognizing a conditioned stimulus, or an oddball stimulus). In other words, NE release is related to unexpected uncertainty—gross changes in the environment that violate top-down expectations (Yu & Dayan, 2005). However, it is not related to anticipation and is independent of valence (C. W. Berridge & Waterhouse, 2003). Norepinephrine appears to be important for accuracy of an action and a trade-off between distractibility and vigilance (Robbins et al., 1998). At low tonic levels of locus coeruleus activity subjects are inattentive and at high tonic levels subjects are distracted (Aston-Jones & Cohen, 2005). However, at moderate levels subjects are engaged in a task, respond to task relevant stimuli and perform well.

3 Neuromodulatory Effect on Cortical Networks

A major premise of this article is that although each neuromodulatory system is triggered by different environmental stimuli, the effect of the neuromodulatory systems on the nervous system is the same: it results in decisive responses when necessary and arbitrary

responses when desirable. All neuromodulatory systems appear to have both tonic and phasic activity responses (Briand et al., 2007). In tonic mode, the baseline activity of a neuromodulatory system is elevated but not bursting. In phasic mode, the neuromodulatory system exhibits short bursts of activity. When the neuromodulatory systems have low tonic activity, the signal-to-noise ratio is low and the organism's behavior is distracted. However, such distracted behavior may be favorable by allowing the organism to explore new actions and break out of local minima. When the neuromodulatory systems have phasic activity, the signal-to-noise ratio increases and the system is in an exploitive mode. In this mode, the organism is more decisive and is not bothered by distractions.

3.1 Tonic Versus Phasic Modes of Neuromodulatory Activity

All neuromodulatory systems appear to have both a tonic mode, where activity ranges between 1 and 6 Hz, and a phasic mode where there is a transient burst of activity (Briand et al., 2007). Phasic signals have a cognitive function for increasing signal detection and driving decision-making processes (Aston-Jones & Cohen, 2005). In the tonic mode, neuromodulatory activity is correlated with distractibility and poor performance. In the tonic mode, the animal is more likely to consider distracters. In the phasic mode, the animal is more attentive and decisive (Aston-Jones & Cohen, 2005).

Although Aston-Jones was describing the locus coeruleus system, this description of the tonic and phasic modes can be extended to the raphe nucleus, basal forebrain, and the ventral tegmental areas effect on downstream targets. For example, VTA DA neurons discharge in both tonic and phasic fashions and these firing patterns result in tonic and phasic release of DA in the prefrontal cortex (Lapish, Kroener, Durstewitz, Lavin, & Seamans, 2007). It appears that the DA phasic signal makes the prefrontal cortex more responsive to behaviorally relevant stimuli. Emerging evidence has also identified tonic and phasic modes of cholinergic activity (Briand et al., 2007).

3.2 Phasic Mode: Suppress Distracters and Increase Signal-to-Noise Ratio

The effect of phasic activity on downstream targets is to increase the signal-to-noise ratio (SNR) in neural circuits such that the organism increases the discrimination between optimal and nonoptimal stimuli. Neuromodulators may achieve this increase in the SNR by amplifying thalamocortical inputs, increasing inhibitory currents, and suppressing associational inputs (Gu, 2002; Hasselmo & McGaughy, 2004; Kobayashi et al., 2000; Lapish et al., 2007). The effect of an increased SNR due to neuromodulation is to change cortical sensory maps such that processing of behaviorally relevant stimuli is enhanced (Sarter, Hasselmo, Bruno, & Givens, 2005). For example, it has been shown that receptive fields of pyramidal cells sharpen with the application of 5-HT and DA (Williams & Goldman-Rakic, 1995; Williams, Rao, & Goldman-Rakic, 2002). Figure 2 shows how altering the excitatory and inhibitory currents through phasic neuromodulation could cause a winner-take-all (WTA) network response.

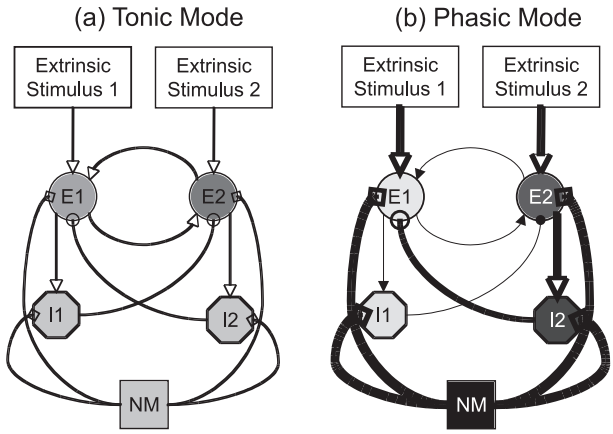

Figure 2 Effect of neuromodulators on network properties. (a) In the tonic mode, the neuromodulator, NM, is tonically active, both extrinsic and associative (*E1*←→*E2*) connections drive excitatory pyramidal cells (*E1*, *E2*) and inhibitory interneurons (*I1*, *I2*). There is only a slight bias toward *E2* over *E1*. (b) In the phasic mode, NM is phasically active, extrinsic connections are enhanced, and associative connections are suppressed. The phasic activity drives the network toward a winner-take-all situation and there is a strong bias of *E2* over *E1*.

3.3 Decision Making and the Trade-Off Between Exploring and Exploiting the Environment

A simple network model was constructed to illustrate how the neuromodulatory systems can potentially drive behavior to exploit previously successful actions,

or explore novel and potentially favorable actions. Although the model is simple, it does demonstrate a neurobiologically plausible mechanism, as well as a possible behavioral outcome of switching between tonic and phasic neuromodulation. During phasic neuromodulation, inhibitory inputs and extrinsic inputs from the thalamus are amplified relative to the intrinsic or associational inputs, whereas during tonic neuromodulation, the intrinsic synaptic connections have a stronger effect (Gu, 2002; Hasselmo & McGaughy, 2004; Kobayashi et al., 2000; Lapish et al., 2007).

The model's architecture followed the connectivity in Figure 2. The activation function for the two excitatory neurons, *E1* and *E2* are:

$$\begin{aligned} E_1(t+1) &= Ext_1 w_{ext} + E_2(t) w_{int} - E_2(t) w_{inh} + rnd(-1,+1) \\ E_2(t+1) &= Ext_2 w_{ext} + E_1(t) w_{int} - E_1(t) w_{inh} + rnd(-1,+1) \end{aligned} \quad (1)$$

where Ext_1 and Ext_2 are inputs from extrinsic neurons, w_{ext} is the extrinsic connection weight, w_{int} is the intrinsic connection weight, w_{inh} is the inhibitory weight, and *rnd* returns a random number between negative one and positive one.

To simulate the effect of phasic neuromodulation, the extrinsic and inhibitory synaptic weights were amplified at the mid-point of the simulation. During the first 50 simulation cycles, the synaptic weights, w_{ext}, w_{int}, and w_{inh} were all set to be equal (i.e., 0.10). During the last 50 simulation cycles, the extrinsic and inhibitory synaptic weights were increased. Figure 3 shows a representative example in which the network behavior was prone to random fluctuations in the tonic mode, but shifted to strongly biased in the phasic mode. In the example shown in Figure 3, Ext_1 was set to 0.50 and Ext_2 was set to 0.60. At time step 50, the inhibitory and extrinsic synaptic weights were set to 2.5 to simulate phasic neuromodulation.

To further explore the effect of phasic neuromodulation, simulations were run over a range of input values and synaptic weights (see Figure 4). Each simulation was run 100 times with a given set of parameter values. Figure 4a shows how robust the effect of phasic neuromodulation is over a wide range of values. During tonic neuromodulation, *E2* was higher than *E1* in only 56% of the trials, whereas after phasic neuromodulation, *E2* was higher than *E1* in 81% of the trials ($p << .00001$; *t* test). In contrast, when the intrinsic connections were modulated (see Figure 4b),

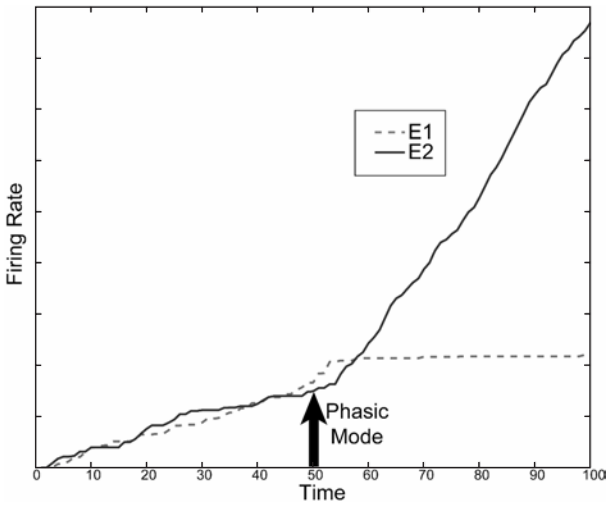

Figure 3 Simulation of the effect of phasic activity on decision making. The lines represent accumulated activity of two simulated neurons. The network connectivity follows Figure 2. Neuron *E2* (solid black line) is receiving slightly higher extrinsic stimulus input (0.60 vs. 0.50) than neuron *E1* (dashed gray line). In the tonic mode (prior to black arrow), the two processes were subject to random fluctuations with no clear winner. In the phasic mode (after the black arrow), the extrinsic (w_{ext}) and inhibitory (w_{inh}) synaptic weights were increased from 0.10 to 2.5. As a result of this neuromodulation, the activity of *E2* is enhanced and the activity of the *E1* is suppressed.

the differences between *E1* and *E2* were relatively arbitrary (54% before the weight change and 53% after the weight change) and insignificant ($p > .85$; *t* test).

Another possible neuromodulatory mechanism is an alteration of the synaptic gain in target neurons (Bogacz, Usher, Zhang, & McClelland, 2007; Servan-Schreiber, Printz, & Cohen, 1990). In this computational model, phasic neuromodulation occurs through change in synaptic gain that increases the contrast between activated and inhibited units. This drives the neurons to be more binary in function when the gain is high, and more arbitrary when the gain is low. The activation function in Equation 2 can achieve this modulation through changes in the gain parameter.

$$activation = \frac{1}{1 + e^{-(gain * input)}} \quad (2)$$

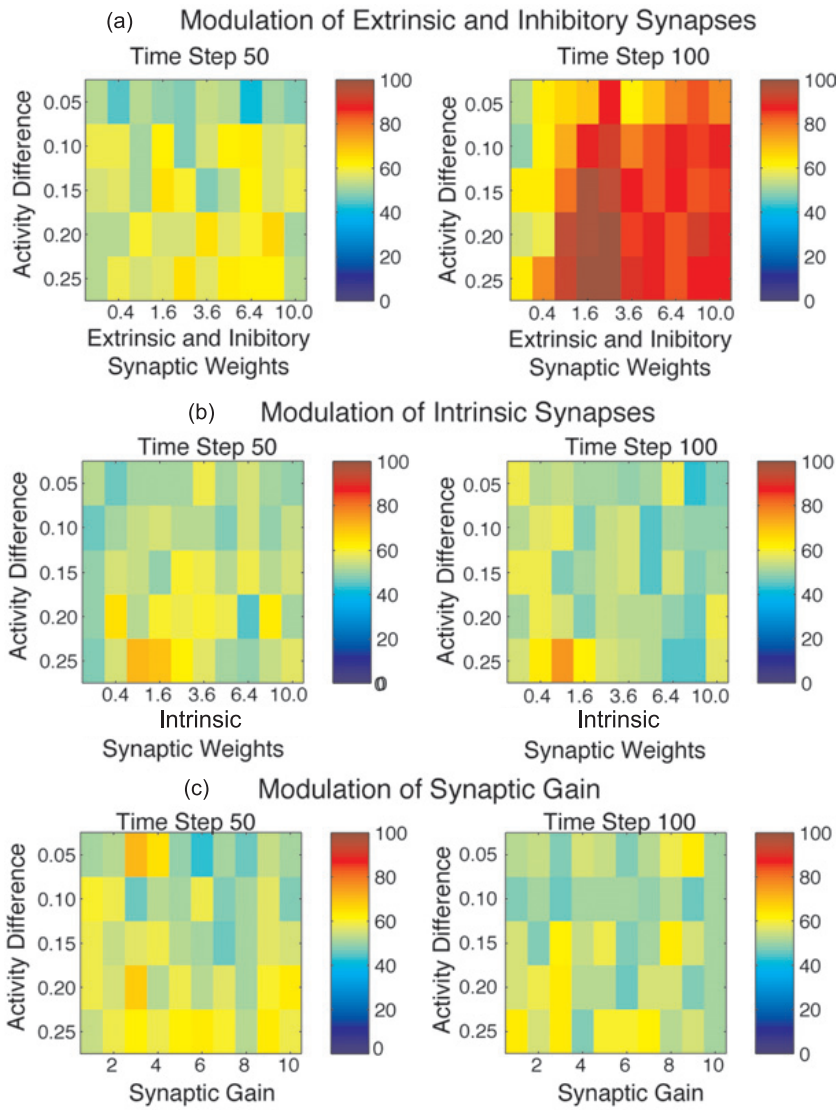

Figure 4 Simulations of neuromodulation under varying conditions. The same protocol shown in Figure 3 was used in these simulations. Each pixel in the figures shows the number of trials out of 100 where the activity of *E2* was greater than *E1* at the mid-point of the trial (left column) and at the end of the trial (right column). The difference in input activity was given by setting *E1* to 0.50, and *E2* to 0.50 plus the value shown in the *y*-axis. Prior to time step 50, all weights were set to 0.10. (a) Simulation of phasic neuromodulation by increasing the extrinsic and inhibitory synaptic weights. At time step 50, the extrinsic and inhibitory weights were set to the value given in the *x*-axis. (b) Simulation of tonic neuromodulation by increasing the intrinsic synaptic weights. At time step 50, intrinsic weights were set to the value given in the *x*-axis. (c) At time step 50, the synaptic gain was set to the value given in the *x*-axis (see Equation 2).

where input is the total synaptic input to the neuron, which is given in Equation 1, and gain is an open parameter.

Although this change in synaptic gain drove the network to be binary in its response, the choice was not necessarily advantageous (see Figure 4c). The differences between *E1* and *E2* were relatively arbitrary (54% for tonic and 54% for phasic) and insignificant ($p > .90$; t test). In these simulations, the neuron with higher activation at the time of phasic neuromodulation tended to be the winner.

The simulations shown in Figures 3 and 4 illustrate a biologically plausible network mechanism for the exploration/exploitation trade-off. It appears from

these simulations that altering the relative weighting of synaptic drives is more effective in exploiting sensory stimuli differences than changing the synaptic gain to all inputs.

This notion of neuromodulation through altering synaptic signaling could be expanded from the simple examples above to a more complete network model. Differences in sensory stimuli would make some decision neurons more active than others. Phasic neuromodulation would amplify those decision neurons that were slightly more active such that there is a winner-take-all response. Initially, neuromodulator levels could be driven by environmental stimuli as described in Section 2. But, neuromodulator levels can be driven from higher order cortical and limbic areas (see Figure 1). Therefore, these decision neurons, through experience-dependent plasticity, may associate specific neuromodulators with extrinsic stimuli and drive phasic neuromodulation in response to sensory cues.

4 Computational Framework for Neuromodulatory Systems

In order to fully understand the neuromodulatory systems, such that they can be applied to artificial systems, a general theory of how neuromodulation affects the nervous system and the organism is necessary. There are many open issues, at a basic systems neuroscience level, in understanding how these neuromodulatory systems control behavior. Many of these systems have been studied individually, but few have studied the interactions between these systems. A computational framework for the neuromodulatory systems is presented based on the following premises:

- The common effect of the neuromodulatory systems is to drive an organism to be decisive when environmental conditions call for such actions, and to allow the organism to be more exploratory when there are no pressing events.
- The main differences between neuromodulatory systems are the environmental stimuli that activate them. The serotonergic (5-HT) system sets the threat level for risk aversion (Millan, 2003), the cholinergic (ACh) system sets the level of attentional effort (Baxter & Chiba, 1999), the dopaminergic (DA) system drives reward anticipation (Schultz et al., 1997) and motivation (K. C. Berridge, 2004), and the noradrenergic (NE) system sets the level of response to novel and salient objects (C. W. Berridge & Waterhouse, 2003; Yu & Dayan, 2005).

These premises are novel because they describe a function for neuromodulation in general and because such a method could be applied to the control of artificial agents.

4.1 Related Work

This article presents a framework for controlling the behavior of artificial systems by specifically examining the function of the vertebrate neuromodulatory system. In this section, I review work from the fields of autonomous agents and computational neuroscience that is related to topics of this article. Of particular relevance are models of neuromodulation, as well as models of affect, emotion, motivation, action selection, and exploration/exploitation trade-offs.

In the field of autonomous agents, "affect" has been used to shape the behavior of both simulated and robotic agents. For example, Blanchard and Cañamero (2006) examined trade-offs between exploration and exploitation based on the notions of well-being and affect. They define "affect" as the immediate or instinctive evaluation of a situation (positiveness or negativeness). In their experiments, the robot's well-being was related to the agent's internal value judgment of its distance to a box, and its affect was related to the agent's evaluation of safety based on its familiarity with the objects it was sensing. The robot's behavior is dictated by a dynamical system and the affect term modulates the robot's motivation to continue. Similar to affect is the notion of comfort or safety, which has also been proposed to influence exploration behavior in robots (Likhachev & Arkin, 2000). Affect has been used in evolutionary algorithms to develop exploration/exploitation strategies in dynamic choice trials (McMahon, Scott, Baxter, & Browne, 2006), and affect has been embedded into the reinforcement-learning algorithm where reward is based on the happiness and sadness of the agent (Salichs & Malfaz, 2006).

In an interesting model that combines emotional affect with internal anticipation, Broekens and colleagues modeled action selection through internal simulation of behavior (Broekens, Kosters, & Verbeek, 2007). The learning by their agents was influenced by

affect: the positiveness and negativeness of a situation. The affect mechanism was based on how well the agent was doing compared with what it was used to. The internal simulation mechanism used artificial affect to select potential interactions and fed these simulated interactions into a reinforcement-learning (RL) algorithm. Reinforcement learning treated these inputs as real and selected behavior based on the predicted value of action–state pairs. Their agents were simulated in two different gridworlds. One in which they had to cope with changes in the environment that made goals less attainable, and another where the meaning of cues in the environment could change drastically. Positive affect was more applicable to the former task, whereas negative affect appeared to be important for the latter when a new task had to be learned.

Of particular interest to the present work are models of emotions because emotions may be linked to the action of neuromodulation. However, real biological emotions are extraordinarily complex and beyond the scope of this article (Damasio, 2005; Rolls, 2000). Models of emotions have been used to control and affect behavior selection (Cañamero, 1997; Moffat, Phaf, & Frijda, 1993; Pfeifer, 1994). For example, Cañamero proposed that emotions could send "hormones" that affect action selection, attention, and perception. In Cañamero's model, emotions included fear, anger, happiness, sadness, boredom, and interest (Cañamero, 1997). The agents in this model also had motivations such as aggression, cold, curiosity, fatigue, hunger, self-protection, thirst, and warmth. Some of these emotions and motivations described by Cañamero may have neural correlates to the actions of the neuromodulators that were described in Section 2. The agents, called Abbotts, had to avoid enemies and find food and drink to survive in a dynamic simulated environment. Inspired by Cañamero's model of emotion released hormones, Cañamero and colleagues investigated how hormonal feedback and affect can modulate action selection and the exploration/exploitation trade-off by introducing a hormone-like mechanism that modulated the perception of cues in a robot (Avila-Garcia & Cañamero, 2004). Their "hormone-like" mechanism modified the perception of incentive cues by acting on the motivation parameter in the equations dictating the robot's behavior. The hormonal feedback modulated the robot's action selection in an environment where it competed for resources with another robot.

In the autonomous agent models discussed above, brain-inspired terms such as anticipation, affect, emotions, hormones, and modulation are simulated to facilitate action selection and exploration/exploitation trade-offs. However, these models do not specifically address the mechanism by which the nervous system gives rise to behaviors, such as action selection, exploration, and exploitation. The framework presented here, provides a neural description of how the neuromodulatory effect on neural circuitry could account for such adaptive behavior. A possible advantage of this approach is that it provides a model that can be directly tested against animal models; both in its behavioral response and in its neuronal response. This synergy between empirical and simulated data, which can lead to improvements in the model and predictions in the modeled organism, is a goal of computational neuroscience.

In the field of computational neuroscience, theoretical models have been proposed on neuromodulation, but they have not considered all of the neuromodulatory systems and their interactions with cortical and subcortical areas. The phasic response of the dopamine system has been proposed to signal temporal difference error (Schultz et al., 1997; Sutton & Barto, 1998). Following this idea, the phasic response of dopamine has been modeled to shape action selection and reward anticipation behavior in both simulated agents and neurorobots (Alexander & Sporns, 2002; Frank & Claus, 2006; Frank, Seeberger, & O'Reilly, 2004; Krichmar & Edelman, 2002; Krichmar, Seth, Nitz, Fleischer, & Edelman, 2005; Sporns & Alexander, 2002). A recent theoretical model speculated that cholinergic release is related to "expected uncertainty" and noradrenergic release is related to "unexpected uncertainty" (Yu & Dayan, 2005). Their statistical model predicted ACh and NE levels, as well as the behavioral responses in a probabilistic cued attention task. Frank and O'Reilly have developed models of decision making based on the anatomical and physiological properties of the basal ganglia and its interaction with frontal cortex (Frank & Claus, 2006; Frank et al., 2004). In these models, they explicitly model neurons, neuroanatomy, synaptic plasticity, and the levels of neuromodulators such as dopamine and norepinephrine. Recently, they developed a model of frontal-striatal dopamine and noradrenergic function to understand attention deficit/hyperactivity disorder (ADHD) in a Go/NoGo learning task (Frank, Santamaria, O'Reilly,

& Willcutt, 2007). Their model predicts that dopamine is important for motivation and updating working memory, while the noradrenergic system is involved in response inhibition and variability.

Neural models of action selection do not necessarily address neuromodulation, but are pertinent to the ideas presented here. For example, one biologically inspired model of action selection is called the leaky competing accumulator (LCA) model of choice (Bogacz et al., 2007; Usher & McClelland, 2001). Similar to the examples given in Figures 3 and 4, there is an accumulation of evidence and competition between neuronal populations corresponding to different alternatives prior to a decision. However, the mechanism of neuromodulation is different than that proposed here. Bogacz uses an alteration of the gain function similar to Equation 2 that was used generate the results shown in Figure 4c (Bogacz et al., 2007). In contrast, the present article proposes a specific shift in synaptic signaling (see Figure 4a). Prescott and his colleagues have developed models of the basal ganglia and more recently the brainstem that demonstrate action selection and switching behavior (Gurney, Prescott, & Redgrave, 2001a, 2001b; Gurney, Prescott, Wickens, & Redgrave, 2004; Humphries, Gurney, & Prescott, 2007; Prescott, Montes Gonzalez, Gurney, Humphries, & Redgrave, 2006). Although these models do not explicitly include neuromodulation or plasticity, they show that action selection and switching behavior can emerge from the unique anatomy and dynamics of these structures.

Kenji Doya and colleagues have proposed a theory that encompasses multiple neuromodulatory systems (Doya, 2002, 2008). In this theory, Doya suggests that each neuromodulator has a specific parametric role in the temporal difference learning; that is, dopamine represents the global learning signal for the prediction of rewards, serotonin controls the balance between short-term and long-term prediction or temporal discounting, norepinephrine controls the balance between exploration and exploitation, and acetylcholine controls the speed of memory update or learning rate. In their "Cyber-Rodent" project, autonomous robots explore an environment containing battery packs for consumption, and conspecifics with which to exchange genetic information. In these robotic experiments they have shown that agents explore new behaviors when their battery packs are low, but take more exploitative behavior when their battery packs are nearly empty (Doya & Uchibe, 2005). They suggest that this behavior has correlates with features of the noradrenergic system. In their recent neurobiology work, they have shown that serotonin levels are related to reward discounting (Schweighofer et al., 2008; Schweighofer, Tanaka, & Doya, 2007; Tanaka et al., 2007). Although this theory is elegant in encapsulating the neuromodulatory systems, it does not take into account their effect on neuronal targets and their specific ability to shape behavioral responses. Also, despite their intriguing modeling and experimental evidence, it is not clear from the behavioral and neuroscience literature that the role of neuromodulation is to calculate temporal difference error and reinforcement learning. For example, it has been proposed that dopamine is involved with the discovery of new actions as opposed to reward prediction (Redgrave & Gurney, 2006).

4.2 Neuromodulation as a Controller for an Artificial Agent

In this article, the effect of neuromodulation on neuronal decision making is specifically addressed, a neural mechanism has been put forward, and a role for each neuromodulator has been identified. All neuromodulators have the same effect on downstream targets, that is, neuromodulators change the synaptic influences of extrinsic and intrinsic inputs to post-synaptic targets (see Figures 2 to 4). Phasic neuromodulation changes the organism's behavior to be more exploitive or decisive, whereas tonic neuromodulation causes the organism to be more exploratory. This is in agreement with the idea of cholinergic modulation of attention (Pauli & O'Reilly, 2008) and noradrenergic modulation of decision making (Aston-Jones & Cohen, 2005), but extends it to the other neuromodulators. Specific neuromodulator levels are driven by different environmental stimuli, but could also be influenced by limbic or cortical areas through experience-dependent plasticity and associative learning.

A novel means of controlling an autonomous system could be developed by following the above principles of vertebrate neuromodulation. A controller for an artificial agent, based on the neuromodulatory system, could be advantageous for autonomous robots carrying out tasks in the face of environmental challenges. Environmental signals could trigger different simulated neuromodulators causing the system to pay particular attention to specific stimuli (see Figure 5).

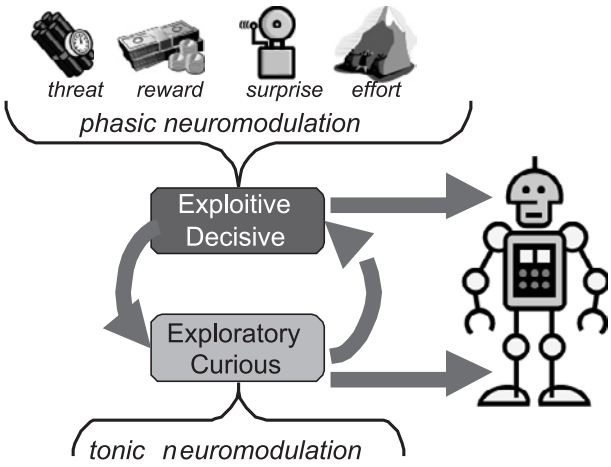

Figure 5 Schematic of the effect of neuromodulation on the behavior of an agent. Phasic neuromodulation drives the agent toward more exploitive and decisive behavior, and tonic neuromodulation drives the agent toward more exploratory or curious behavior. The activity of each neuromodulator is related to environmental stimuli. For example, serotonin levels appear to be related to threat assessment, dopamine levels appear to be related to reward anticipation, norepinephrine levels appear to be related to surprise or novelty, and acetylcholine levels appear to be related to attentional effort.

For example, a loud noise or looming object could trigger the serotonergic system causing the robot to be more risk averse. An object that suddenly appears could trigger the noradrenergic system causing the robot to be more inquisitive. The need to recharge batteries or get maintenance could trigger the dopaminergic system causing the robot to be more intent on acquiring its goal. Sensor confusion or a noisy environment could trigger the cholinergic system causing the robot to focus on the most impending signals and ignore distractions.

4.3 Future Work

In the future, we will expand the illustrative simulations and networks of Figures 2 to 4 into a complete network model containing the different neuromodulators and cortical and limbic areas that have simulated neuromodulator receptors. Initially, different environmental stimuli would drive specific neuromodulators as described in Section 2. After experience-dependent learning, the cortical and limbic regions may become associated with the sensory stimuli that trigger neuromodulation. In this manner, sensory and decision-making neurons could drive the neuromodulatory regions prior to an environmental event. For example, if the color red has become associated with a noxious stimulus (e.g., a bad odor), the agent's visual system, upon seeing red, may drive the serotonergic system causing an escape response prior to interacting with the red object. In our previous modeling work, we have shown how this predictive signal can arise through experience-dependent plasticity between cortical areas and reward systems and how it can result in decisions being made earlier and more decisively (Krichmar & Edelman, 2002; Seth, McKinstry, Edelman, & Krichmar, 2004). Our previous work concentrated on the dopaminergic system. The present framework expands this idea to other neuromodulators that signal risks, threats, novelty, surprise, and attentional effort. It will be of interest to test this network model based on the architecture in Figure 1 in a more diverse environment. For example, many of the dynamic environments from the artificial intelligence and autonomous agent community may provide a challenging test bed for a neurobiologically inspired model, as well as an interesting comparison of approaches (Avila-Garcia & Cañamero, 2004; Broekens et al., 2007; Cañamero, 1997; Prescott et al., 2006).

5 Summary

The main idea put forth in this article was that neuromodulators have the common effect of driving a neural network to switch between exploratory and exploitive decision making by the alteration of synaptic signaling and that each neuromodulator is driven by different stimuli. Furthermore, it was argued that these notions of neuromodulation could be applied to the control of autonomous agents and robots.

Specifically, the notion put forth by Aston-Jones and colleagues (Aston-Jones & Cohen, 2005) was extended to all neuromodulators. That is, in each of the neuromodulatory systems, the tonic mode results in exploratory behavior, and the phasic mode causes exploitative behavior. Each neuromodulatory system is triggered by specific environmental cues: (1) The raphe nucleus signifies a threat to the organism. (2) The locus coeruleus signifies novelty or saliency; interactions with the raphe nucleus (see Figure 1) may cause an organism to be more vigilant when a

threat is imminent. (3) The ventral tegmental area signifies the motivational need or desire for objects. (4) The basal forebrain signifies the need to increase attentional effort. Note that this increase may be driven in part by the other neuromodulatory systems (see Figure 1).

A controller modeled after the vertebrate neuromodulatory system, in which the robot's behavior approaches the complexity and flexibility associated with higher order animals, could improve the current state of autonomous system design. Such a system would be particularly sensitive to a wide range of changes in an organism's environment and would allow the organism to quickly and decisively act on these changes. Moreover, the neuromodulatory system alters plasticity in areas that are important for conditioning, memory, and planning (e.g., amygdala, hippocampus, and prefrontal cortex) allowing the system to adapt responses over a longer timescale. In this way, the neuromodulatory systems are well tuned to carry out important functions for an organism's survival. The neuromodulatory system is not only crucial for survival behavior, but it also provides the foundation from which cognitive function in higher organisms arises. This neurally inspired control system would be flexible, experience dependent, and autonomous; just like a biological organism.

References

Alexander, W. H., & Sporns, O. (2002). An embodied model of learning, plasticity, and reward. *Adaptive Behavior*, *10*, 143–159.

Amat, J., Paul, E., Zarza, C., Watkins, L. R., & Maier, S. F. (2006). Previous experience with behavioral control over stress blocks the behavioral and dorsal raphe nucleus activating effects of later uncontrollable stress: role of the ventral medial prefrontal cortex. *Journal of Neuroscience*, *26*, 13264–13272.

Arleo, A., Smeraldi, F., & Gerstner, W. (2004). Cognitive navigation based on nonuniform Gabor space sampling, unsupervised growing networks, and reinforcement learning. *IEEE Transactions on Neural Networks*, *15*, 639–652.

Aston-Jones, G., & Cohen, J. D. (2005). Adaptive gain and the role of the locus coeruleus-norepinephrine system in optimal performance. *Journal of Comparative Neurology*, *493*, 99–110.

Avila-Garcia, O., & Cañamero, L. (2004). Using hormonal feedback to modulate action selection in a competitive scenario. In *From Animals to Animats 8, Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB'04)* (pp. 243–252). Cambridge, MA: MIT Press.

Barnes, N. M., & Sharp, T. (1999). A review of central 5-HT receptors and their function. *Neuropharmacology*, *38*, 1083–1152.

Baxter, M. G., & Chiba, A. A. (1999). Cognitive functions of the basal forebrain. *Current Opinion in Neurobiology*, *9*, 178–183.

Berridge, C. W., & Waterhouse, B. D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Research Reviews*, *42*, 33–84.

Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience. *Physiology & Behavior*, *81*, 179–209.

Blanchard, A., & Cañamero, L. (2006). Modulation of exploratory behavior for adaptation to the context. In *Proceedings of the AISB'06 Symposium on Biologically Inspired Robotics (Biro-net)* (pp. 131–137). Bristol, UK: AISB Press.

Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society of London*, *362*, 1655–1670.

Briand, L. A., Gritton, H., Howe, W. M., Young, D. A., & Sarter, M. (2007). Modulators in concert for cognition: Modulator interactions in the prefrontal cortex. *Progress in Neurobiology*, *83*, 69–91.

Broekens, J., Kosters, W. A., & Verbeek, F. J. (2007). Affect, anticipation, and adaptation: Affect-controlled selection of anticipatory simulation in artificial adaptive agents. *Adaptive Behavior*, *15*, 397–422.

Bucci, D. J., Holland, P. C., & Gallagher, M. (1998). Removal of cholinergic input to rat posterior parietal cortex disrupts incremental processing of conditioned stimuli. *Journal of Neuroscience*, *18*, 8038–8046.

Cañamero, L. D. (1997). Modeling motivations and emotions as a basis for intelligent behavior. In *Proceedings of the First International Conference on Autonomous Agents* (pp. 148–155). New York: ACM Press.

Chiba, A. A., Bucci, D. J., Holland, P. C., & Gallagher, M. (1995). Basal forebrain cholinergic lesions disrupt increments but not decrements in conditioned stimulus processing. *Journal of Neuroscience*, *15*, 7315–7322.

Cho, A. (2007). Robotics. Robotic cars tackle crosstown traffic—and not one another. *Science*, *318*, 1060–1061.

Damasio, A. (2005). *Descartes' error: Emotion, reason, and the human brain*. New York: Penguin.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, *15*, 495–506.

Doya, K. (2008). Modulators of decision making. *Nature Neuroscience*, *11*, 410–416.

Doya, K., & Uchibe, E. (2005). The Cyber Rodent project: Exploration of adaptive mechanisms for self-preservation and self-reproduction. *Adaptive Behavior*, *13*, 149–160.

Fluxe, K., Hokfelt, T., Johansson, O., Jonsson, G., Lidbrink, P., & Ljungdahl, A. (1974). The origin of the dopamine nerve terminals in limbic and frontal cortex. Evidence for mesocortico dopamine neurons. *Brain Research*, *82*, 349–355.

Frank, M. J., & Claus, E. D. (2006). Anatomy of a decision: striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychological Review*, *113*, 300–326.

Frank, M. J., Santamaria, A., O'Reilly, R. C., & Willcutt, E. (2007). Testing computational models of dopamine and noradrenaline dysfunction in attention deficit/hyperactivity disorder. *Neuropsychopharmacology*, *32*, 1583–1599.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, *306*, 1940–1943.

Gibbs, W. W. (2004). From finish to start. Was the grand challenge robot race in March the fiasco it appeared to be? Hardly, argues William "Red" Whittaker. The annual event is pushing mobile robotics to get real. *Scientific American*, *291*, 33–34.

Gu, Q. (2002). Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. *Neuroscience*, *111*, 815–835.

Guenter, F., Hersch, M., Calinon, S., & Billard, A. (2007). Reinforcement learning for imitating constrained reaching movements. *Advanced Robotics*, *21*, 1521–1544.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, *84*, 401–410.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biological Cybernetics*, *84*, 411–423.

Gurney, K., Prescott, T. J., Wickens, J. R., & Redgrave, P. (2004). Computational models of the basal ganglia: from robots to membranes. *Trends in Neurosciences*, *27*, 453–459.

Han, J. S., Holland, P. C., & Gallagher, M. (1999). Disconnection of the amygdala central nucleus and substantia innominata/nucleus basalis disrupts increments in conditioned stimulus processing in rats. *Behavioral Neuroscience*, *113*, 143–151.

Harvey, J. A. (2003). Role of the serotonin 5-HT(2A) receptor in learning. *Learning & Memory*, *10*, 355–362.

Hasselmo, M. E., & McGaughy, J. (2004). High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. *Progress in Brain Research*, *145*, 207–231.

Humphries, M. D., Gurney, K., & Prescott, T. J. (2007). Is there a brainstem substrate for action selection? *Philosophical Transactions of the Royal Society of London*, *362*, 1627–1639.

Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: the role of reward-related learning and memory. *Annual Review of Neuroscience*, *29*, 565–598.

Iida, S., Kuwayama, K., Kanoh, M., Kato, S., & Itoh, H. (2004). A dynamic allocation method of basis functions in reinforcement learning. In *Proceedings of AI 2004*, *Advances in Artificial Intelligence*, Lecture notes in computer science (Vol. 3339, pp. 272–283). Berlin: Springer.

Jones, J., & Roth, D. (2003). *Robot programming : A practical guide to behavior-based robotics*. New York: McGraw-Hill.

Kobayashi, M., Imamura, K., Sugai, T., Onoda, N., Yamamoto, M., Komai, S., & Watanabe, Y. (2000). Selective suppression of horizontal propagation in rat visual cortex by norepinephrine. *European Journal of Neuroscience*, *12*, 264–272.

Kondo, T., & Ito, K. (2004). A reinforcement learning with evolutionary state recruitment strategy for autonomous mobile robots control. *Robotics and Autonomous Systems*, *46*, 111–124.

Krichmar, J. L., & Edelman, G. M. (2002). Machine psychology: Autonomous behavior, perceptual categorization, and conditioning in a brain-based device. *Cerebral Cortex*, *12*, 818–830.

Krichmar, J. L., Seth, A. K., Nitz, D. A., Fleischer, J. G., & Edelman, G. M. (2005). Spatial navigation and causal analysis in a brain-based device modeling cortical-hippocampal interactions. *Neuroinformatics*, *3*, 197–222.

Lapish, C. C., Kroener, S., Durstewitz, D., Lavin, A., & Seamans, J. K. (2007). The ability of the mesocortical dopamine system to operate in distinct temporal modes. *Psychopharmacology*, *191*, 609–625.

Likhachev, M., & Arkin, R. (2000). Robotic comfort zones. In *Proceedings of SPIE: Sensor Fusion and Decentralized Control in Robotic Systems III* (pp. 27–41).

McClure, S. M., Daw, N. D., & Montague, P. R. (2003). A computational substrate for incentive salience. *Trends in Neurosciences*, *26*, 423–428.

McGaugh, J. L. (2004). The amygdala modulates the consolidation of memories of emotionally arousing experiences. *Annual Review of Neuroscience*, *27*, 1–28.

McMahon, A., Scott, D., Baxter, P., & Browne, W. (2006). An autonomous explore/exploit strategy. In *AISB'06 Symposium on Nature Inspired Systems* (pp. 192–201). Bristol, UK: AISB Press.

Meneses, A., & Perez-Garcia, G. (2007). 5-HT(1A) receptors and memory. *Neuroscience & Biobehavioral Reviews*, *31*, 705–727.

Millan, M. J. (2003). The neurobiology and control of anxious states. *Progress in Neurobiology*, *70*, 83–244.

Moffat, D., Phaf, R. H., & Frijda, N. H. (1993). Analysis of a model of emotions. In *Prospects for Artificial Intelligence* (pp. 219–228). Oxford: IOS Press.

Nakamura, Y., Mori, T., Sato, M. A., & Ishii, S. (2007). Reinforcement learning for a biped robot based on a CPG-actor-critic method. *Neural Networks*, *20*, 723–735.

Pauli, W. M., & O'Reilly, R. C. (2008). Attentional control of associative learning-A possible role of the central cholinergic system. *Brain Research*, *1202*, 43–53.

Pfeifer, R. (1994). The "fungus eater approach" to emotion: a view from artificial intelligence. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*, *1*, 42–57.

Prescott, T. J., Montes Gonzalez, F. M., Gurney, K., Humphries, M. D., & Redgrave, P. (2006). A robot model of the basal ganglia: behavior and intrinsic processing. *Neural Networks*, *19*, 31–61.

Redgrave, P., & Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*, *7*, 967–975.

Robbins, T. W., Granon, S., Muir, J. L., Durantou, F., Harrison, A., & Everitt, B. J. (1998). Neural systems underlying arousal and attention. Implications for drug abuse. *Annals of the New York Academy of Sciences*, *846*, 222–237.

Rolls, E. T. (2000). *The brain and emotion*. New York: Oxford University Press.

Salichs, M. A., & Malfaz, M. (2006). Using emotions on autonomous agents. The role of happiness, sadness and fear. In *Proceedings of the AISB'06 Symposium on Integrative Approaches to Machine Consciousness* (pp. 157–164). Bristol, UK: AISB Press.

Sarter, M., Gehring, W. J., & Kozak, R. (2006). More attention must be paid: the neurobiology of attentional effort. *Brain Research Reviews*, *51*, 145–160.

Sarter, M., Hasselmo, M. E., Bruno, J. P., & Givens, B. (2005). Unraveling the attentional functions of cortical cholinergic inputs: interactions between signal-driven and cognitive modulation of signal detection. *Brain Research Reviews*, *48*, 98–111.

Scatton, B., Simon, H., Le Moal, M., & Bischoff, S. (1980). Origin of dopaminergic innervation of the rat hippocampal formation. *Neuroscience Letters*, *18*, 125–131.

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, *2*, 33–42.

Schultz, W. (1997). Dopamine neurons and their role in reward mechanisms. *Current Opinion in Neurobiology*, *7*, 191–197.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.

Schweighofer, N., Bertin, M., Shishida, K., Okamoto, Y., Tanaka, S. C., Yamawaki, S., & Doya, K. (2008). Low-serotonin levels increase delayed reward discounting in humans. *Journal of Neuroscience*, *28*, 4528–4532.

Schweighofer, N., Tanaka, S. C., & Doya, K. (2007). Serotonin and the evaluation of future rewards: theory, experiments, and possible neural mechanisms. *Annals of the New York Academy of Sciences*, *1104*, 289–300.

Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science*, *249*, 892–895.

Sesack, S. R., & Pickel, V. M. (1992). Prefrontal cortical efferents in the rat synapse on unlabeled neuronal targets of catecholamine terminals in the nucleus accumbens septi and on dopamine neurons in the ventral tegmental area. *Journal of Comparative Neurology*, *320*, 145–160.

Seth, A. K., McKinstry, J. L., Edelman, G. M., & Krichmar, J. L. (2004). Visual binding through reentrant connectivity and dynamic synchronization in a brain-based device. *Cerebral Cortex*, *14*, 1185–1199.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*, 161–168.

Sporns, O., & Alexander, W. H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks*, *15*, 761–774.

Squyres, S. (2005). *Roving Mars : Spirit, opportunity, and the exploration of the Red Planet*. New York: Hyperion.

Stone, P., Sutton, R. S., & Kuhlmann, G. (2005). Reinforcement learning for RoboCup soccer keepaway. *Adaptive Behavior*, *13*, 165–188.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Tanaka, S. C., Schweighofer, N., Asahi, S., Shishida, K., Okamoto, Y., Yamawaki, S., & Doya, K. (2007). Serotonin differentially regulates short- and long-term prediction of rewards in the ventral and dorsal striatum. *PLoS ONE*, *2*, e1333.

Thrun, S., Burgarde, W., & Fox, D. (2005). *Probabilistic robotics*. Cambridge, MA: MIT Press.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.

Williams, G. V., & Goldman-Rakic, P. S. (1995). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. *Nature*, *376*, 572–575.

Williams, G. V., Rao, S. G., & Goldman-Rakic, P. S. (2002). The physiological role of 5-HT2A receptors in working memory. *Journal of Neuroscience*, *22*, 2843–2854.

Yenne, B. (2004). *Attack of the drones*. St. Paul, MN: Zenith Press.

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*, 681–692.

About the Author

Jeffrey L. Krichmar received a B.S. in computer science in 1983 from the University of Massachusetts at Amherst, an M.S. in computer science from The George Washington University in 1991, and a Ph.D. in computational sciences and informatics from George Mason University in 1997. He spent 15 years as a software engineer on projects ranging from the PATRIOT Missile System at the Raytheon Corporation to Air Traffic Control for the Federal Systems Division of IBM. In 1997, he became an assistant professor at The Krasnow Institute for Advanced Study at George Mason University. From 1999 to 2007, he was a senior fellow in theoretical neurobiology at The Neurosciences Institute. He is currently an assistant professor in the Department of Cognitive Sciences at the University of California, Irvine. His research interests include neurorobotics, embodied cognition, biologically plausible models of learning and memory, and the effect of neural architecture on neural function.