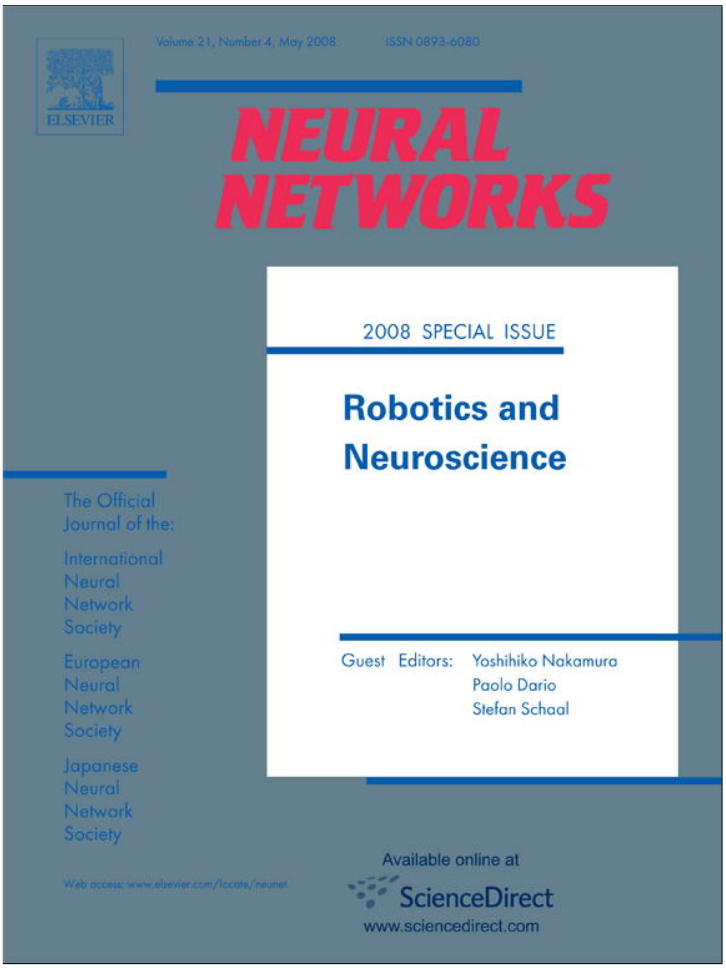

Neural Networks 21 (2008) 553–561

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2008 Special Issue

Embodied models of delayed neural responses: Spatiotemporal categorization and predictive motor control in brain based devices

Jeffrey L. McKinstry *, Anil K. Seth, Gerald M. Edelman, Jeffrey L. Krichmar

The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, United States

ARTICLE INFO

Article history:
Received 4 December 2006
Received in revised form
25 January 2008
Accepted 25 January 2008

Keywords:
Somatosensory cortex
Whisker
Cerebellum
Motor control
Robots
Eligibility traces

ABSTRACT

In order to respond appropriately to environmental stimuli, organisms must integrate over time spatiotemporal signals that reflect object motion and self-movement. One possible mechanism to achieve this spatiotemporal transformation is to delay or lag neural responses. This paper reviews our recent modeling work testing the sufficiency of delayed responses in the nervous system in two different behavioral tasks: (1) Categorizing spatiotemporal tactile cues with thalamic "lag" cells and downstream coincidence detectors, and (2) Predictive motor control was achieved by the cerebellum through a delayed eligibility trace rule at cerebellar synapses. Since the timing of these neural signals must closely match real-world dynamics, we tested these ideas using the brain based device (BBD) approach in which a simulated nervous system is embodied in a robotic device. In both tasks, biologically inspired neural simulations with delayed neural responses were critical for successful behavior by the device.

1. Introduction

In order to respond to real world stimuli at the appropriate time, the nervous system must transform stimuli in space and time or hold a signal for a time period before responding behaviorally. For example, consider a rodent that sweeps its whiskers along an object to determine its shape. How are the sensory impulses combined across time and across multiple whiskers to categorize the object? In another instance, when a rabbit receives a puff of air in the eye, which is paired with a stimulus predicting the occurrence of the air puff, it learns over time to close its eyelid precisely in time to protect the eye from the noxious stimulus. After the noxious stimulus arrives, how is the predictive stimulus maintained long enough to allow associative learning to take place?

Nervous systems integrate signals over durations ranging from microseconds (e.g. delay lines in the owl auditory system, (Carr & Konishi, 1990)) to seconds (e.g. persistent firing during working memory tasks, (Funahashi, Bruce, & Goldman-Rakic, 1989; Fuster & Alexander, 1971)). A mechanism for integrating signals over time is provided by lag cells, found in the visual thalamus of the cat, which respond to a visual stimulus with a characteristic delay that varies from cell to cell (Saul & Humphrey, 1992). This type of cell can function in a similar manner to delay lines, and has been proposed to provide a mechanism for direction selectivity in simple cells in the visual cortex of the cat (Jagadeesh, Wheat, Kontsevich, Tyler, & Ferseter, 1997).

Additional mechanisms for maintaining a signal over time have been found at the synapse. For example, long term depression of parallel fiber synapses onto Purkinje cells is maximal when parallel fibers are stimulated within a time window from 125 to 250 ms prior to climbing fiber activation (Chen & Thompson, 1995). A candidate "eligibility trace" mechanism is embodied in a nonlinear calcium response in these synapses which is maximal when the parallel fiber input precedes climbing fiber activation from 50 to 200 ms (Wang, Denk, & Hausser, 2000).

To test models of delayed neural responses, we used an approach employing brain based devices (BBDs), in which a simulated nervous system is embodied in a robotic device, to test mechanisms of delayed neuronal responses during behavior (Almassy, Edelman, & Sporns, 1998; Edelman et al., 1992; Krichmar & Edelman, 2002, 2005; Krichmar, Nitz, Gally, & Edelman, 2005; Seth, McKinstry, Edelman, & Krichmar, 2004a, 2004b). The BBD approach forces the modeler to consider how the timescale of neural mechanisms matches the timescale of behavior. Physical embodiment is critical for understanding issues of timing in the real world. Virtual sensory input and simulated motor output are designed by the modeler and can inadvertently bias a neural simulation. However, when embedding a nervous system simulation in a behaving device, the device's behavioral response must match its sensorimotor signals. Moreover, physical embodiment in such a device emphasizes many of the challenging aspects of discrimination in the real world: noisy sensors, movement variation, and complexity of a real-world environment.

* Corresponding author. Tel.: +1 858 626 2071; fax: +1 858 626 2099.
E-mail address: jeffmckinstry@nsi.edu (J.L. McKinstry).

0893-6080/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2008.01.004

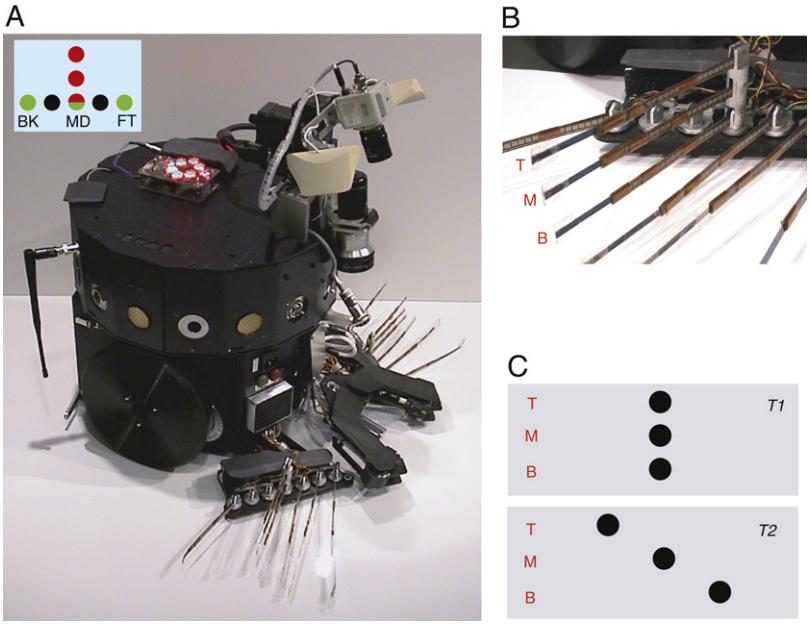

Fig. 1. A. Darwin IX with its left and right whisker arrays. The arrangement of a whisker array is shown in the inset. Each array has 7 whiskers arranged in a row of 5 and a column of 3. Whiskers used for wall following are marked in white (FT, MD, BK). Whiskers that provide input to the neural simulation are marked in black. Note that one whisker (white/black) is used for both purposes. B. Detail of a whisker array: The top (T), middle (M), and bottom (B) whiskers in the column are labeled; these whiskers provide input to the neural simulation. C. Schematic of textures T1 and T2. Each texture consists of pegs embedded in a wall; pegs are aligned in rows corresponding to the whiskers in a column. Pegs in the top row deflect the top whisker (T), and similarly for pegs in the middle row (M) and the bottom row (B).

By using a real-world environment, not only is the risk of biases reduced, but the experimenter is also freed from the burden of constructing a highly complex simulated environment.

This paper describes recent work testing biologically inspired mechanisms of delayed neural responses that facilitate categorization of spatiotemporal tactile cues and predictive cerebellar motor control using the BBD approach. We find that a population of lag-cell-like neuronal units that respond to artificial whisker deflections in a moving device is sufficient to support texture discrimination of whisker–barrel responses lasting approximately one second. In a task where the device's own movement causes visual optic flow, we show that a delayed eligibility trace mechanism at simulated Purkinje cell and deep cerebellar nuclei cell synapses allows for an association in which the visual cue predicts a future collision.

2. Spatiotemporal pattern discrimination in Darwin IX

Haptic sensory information provided by mystacial vibrissae (whiskers) of the rat allows the animal to discriminate among different textures in its environment (Harvey, Bermejo, & Zeigler, 2001; Prigg, Goldreich, Carvell, & Simons, 2002). This requires the integration of sensory input from the whiskers across time and space, providing an excellent model system for exploring spatiotemporal pattern categorization. To explore how haptic data may be integrated into perceptual categories, we equipped a BBD, Darwin IX, with artificial whiskers and a simulated nervous system based on the neuroanatomy of the rat somatosensory system.

In our experiments with Darwin IX, the device autonomously explored a walled environment containing two distinct textures each consisting of various patterns of pegs embedded in the walls. It became conditioned to avoid one of the textures by association of this texture with an innately aversive stimulus (i.e. a change in reflectivity of the environment's flooring). This aversive stimulus was used in an experimental paradigm analogous to fear-conditioning with a 'foot-shock' at particular locations in the environment. Similar to a rodent in such a conditioning paradigm, Darwin IX demonstrated its aversive behavior by stopping ("freezing") and then moving away from noxious stimuli.

We tested the idea that a diverse population of neuronal units with varying sensory response delays could bridge the temporal gaps brought about by moving a tactile sensor across a spatial pattern. Such a scheme has been found in the visual system of the cat (Saul & Humphrey, 1992) and may be a mechanism for direction selectivity in the primary visual cortex (Jagadeesh et al., 1997). Delayed neuronal responses, which have been found in the perirhinal cortex of the rat, can be as long as four seconds from stimulus onset (Beggs, Moyer, McGann, & Brown, 2000).

2.1. Darwin IX: Construction and experimental paradigm

Darwin IX is based on a mobile robotic platform (Nomadic Technologies) augmented by a whisker array on each side (Fig. 1A). Each array consists of seven whiskers arranged in a single column of three and single row of five where one of the whiskers was both the row and column (see Fig. 1A, inset). The whisker column supplied input to the simulated nervous system, while whisker row supported innate avoidance and wall-following behaviors. The whiskers are made of two polyamide strips, placed back to back, that emit a signal proportional to the bending of the strip (Abrams, Gentile Entertainment).

Darwin IX's default behavior was to move forward in a straight line at a speed of ∼8 cm/s. Darwin IX also had an innate wall-following capability based on signals from the first, third, and fifth whiskers in the whisker row. The innate wall-following behavior was programmed with a simple feedback controller that maintained these three whiskers within a desired range (see Seth et al. (2004b) for details).

Darwin IX had an innate freezing/avoidance response which was triggered upon detection of a simulated foot-shock by a downward pointing infra-red sensor that measured changes in reflectivity of floor surface. Construction paper, which was the same color as the floor but more reflective, was placed upon the floor in locations to trigger an innate aversive response. This

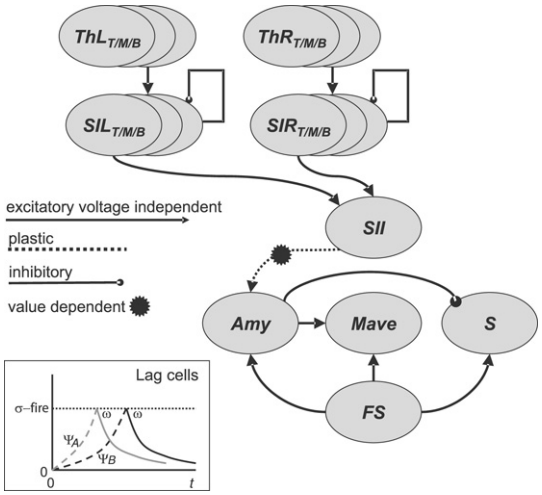

Fig. 2. Neuroanatomy of Darwin IX. The simulated nervous system contains 17 neuronal areas, 1161 neuronal units, and ∼8400 synaptic connections. Aversive avoidance responses are evoked by activity in area M_{ave}. Areas $ThL_{T/M/B}$ *and* $ThR_{T/M/B}$ receive input from the corresponding whiskers in the whisker columns on the appropriate side. These areas contain 'lag cells' with temporal response properties. The operation of two idealized lag cells (A and B) is shown in the inset. The poststimulus internal state of cell $A(s_A^{\text{in}})$ rises quickly (gray dashed line), at a rate determined by φ_A. The internal state of B (s_B^{in}) rises more slowly (black dashed line, φ_B). When the internal state of each cell reaches a threshold (σ^{fire}), output is generated (solid lines) which decays at a rate determined by ω.

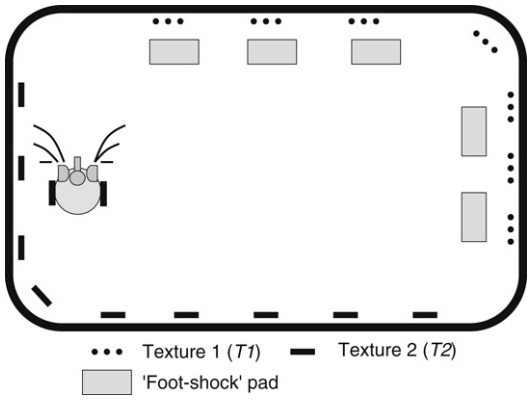

Fig. 3. Experimental setup for Darwin IX. Darwin IX explored a walled enclosure (2.41 m × 2.95 m) with textures T1 and T2 on the walls. Instances of each texture were regularly spaced along the walls at intervals of ∼30 cm. Located on the floor adjacent to T1 patterns were 'foot-shock' pads made of reflective construction paper. Training and testing was repeated for each Darwin IX subject after exchanging the positions of T1 and T2.

response consisted of stopping for ∼4 s followed by a turn away from the noxious stimuli. The signal from the infra-red sensor activated Darwin IX's value system resulting in a neuromodulatory alteration of synaptic strength in the simulated nervous system (see below).

Darwin IX's simulated nervous system contained 17 neural areas, 1101 neuronal units, and ∼8400 synaptic connections (see Fig. 2). A mean-firing rate model was used to simulate neuronal units in Darwin IX where the activity of a neuronal unit represents the average activity of ∼100 neurons over ∼100 ms. Each simulation cycle, that is, the time to update every neuronal unit and plastic synapse, took roughly 100 ms of real time. Darwin IX contained areas analogous to the somatosensory pathway in the rat brain, specifically the ventromedial nuclei of the thalamus, and primary and secondary somatosensory areas (in our model, $Th \rightarrow S1 \rightarrow S2$). Areas $S1$ and *Th* were subdivided into left (L) and right (R) regions and further into 'top' (T), 'middle' (M) and 'bottom' (B) 'barrel' regions, such that each barrel received input from a single whisker in the column on the corresponding side (see Fig. 2). This one-to-one mapping between 'barrel' regions and whiskers is analogous to the whisker barrel formations observed in rat thalamus and $S1$ (Jensen & Killackey, 1987; Woolsey & Van der Loos, 1970).

Neuronal units in the *Th* barrels projected topographically to corresponding barrels of $S1$. Each barrel of $S1$ had local inhibitory connections which served to increase the activity contrast among neuronal units. All barrels in $S1$ projected to area $S2$ such that each neuronal unit in $S2$ took input from 3 neuronal units, each of which was in a *different* barrel of either the left sub-area or the right sub-area of $S1$. This arrangement ensured that synaptic input to a neuronal unit in $S2$ was sparse and balanced. By this arrangement, a deflection of a particular sequence of Darwin IX's whiskers led to a spatiotemporal pattern of activity in $S2$. Such a dynamic sequence was comparable to that observed in the rat brain (Ghazanfar & Nicolelis, 1999).

Darwin IX's nervous system also contained areas supporting conditioning. The neural area *FS* was activated by detection of a simulated 'foot-shock' and projected to areas *S, Amy,* and M_{ave}. Area *Amy* is analogous to the amygdala, a neural area which has been widely implicated in the acquisition of conditioned fear (LeDoux, 1995; Maren & Fanselow, 1996). Area M_{ave} is analogous to a motor cortical area, in which activity elicited an innate aversive freezing/avoidance response. Synaptic plasticity in Darwin IX between neuronal units in areas $S2$ and *Amy* (see Fig. 2) was carried out using a modified Hebbian rule. This process is 'value dependent', i.e. the degree of change is modulated by activity in the simulated value system (area S). Connections were strengthened when the value system activity was above the baseline, and weakened when below the baseline. For details, see Seth et al. (2004b).

Each barrel in area *Th* contained 20 'lag' cells; neuronal units which have varying, time-lagged response properties similar to those found in the lateral geniculate nucleus of the cat (Saul & Humphrey, 1992; Wolfe & Palmer, 1998). Each lag cell neuronal unit responded to a whisker deflection with a specific delay, ranging from 1 to 20 simulation cycles. Details of the lag cell model can be found in the Appendix A and are illustrated in the inset in Fig. 2. The three whiskers aligned in a column provided input to the corresponding simulated whisker barrels in area *Th*.

Darwin IX's environment was a rectangular arena with black flooring and textures along the walls of its environment (see Fig. 3). One texture ($T1$) consisted of a vertically aligned column of pegs, the other ($T2$) consisted of a vertically staggered column of pegs with offsets between pegs of ∼6 cm. Two adjacent walls contained $T1$, the other two contained $T2$; either $T1$ or $T2$ was associated with a simulated aversive foot-shock. The aversive response was triggered by Darwin IX's downward-facing IR sensor when it detected the reflective construction paper placed on the floor of the arena near the aversive texture.

Experiments were divided into training and testing stages. During training, one of the two textures was paired with the simulated foot-shock (the other texture was neutral). Darwin IX autonomously explored its enclosure for 25,000 simulation cycles, corresponding to ∼48 encounters with each wall and ∼24 aversive responses to the simulated foot-shock. During testing, the foot-shock pads were removed and Darwin IX was allowed to explore its enclosure for ∼15,000 simulation cycles. Training and testing were repeated using three different Darwin IX "subjects" initialized with different random seeds, pairing both $T1$ and $T2$ with foot-shock (six training/testing episodes in total). During training and testing of each subject, responses of all neuronal units were recorded and saved for analysis. The position of Darwin IX was continuously recorded by an overhead camera that detected an array of LEDs

positioned on the top surface of the device, the images from which were time stamped for analysis.

2.2. Texture discrimination and spatiotemporal categorization

After training, texture discrimination by Darwin IX subjects was assessed by removing the reflective construction paper and measuring the number of conditioned responses to the texture associated with the simulated foot-shock locations. During testing, Darwin IX subjects which were trained to avoid $T1$ made aversive responses on 96.6% (S.E. = 0.18%) of encounters with $T1$. When trained to avoid $T2$, these subjects made aversive responses on 97.9% (S.E. = 0.14%) of encounters with $T2$. Only 3.2% of all aversive responses to both textures occurred inappropriately, i.e. in response to whisker deflections by walls or by the texture *not* associated with foot-shock. This behavioral assay demonstrated that Darwin IX was able to discriminate between the two spatiotemporal texture patterns of whisker deflection.

Darwin IX's ability to categorize textural stimuli is supported by spatiotemporal patterns of neural activity in $S2$. Each texture deflected the array of whiskers in a column in a specific temporal order. The lag cells in area *Th* and neural units downstream in $S1$ presented a pattern of activity with both a spatial component (i.e. the particular whisker) and a temporal component (i.e. the time since deflection). $S2$ neuronal units acted as coincidence detectors responding to particular combinations of this spatiotemporal activity in S1.

The population response of $S2$ to a texture was specific and repeatable, and was supported by quantitative measures of pattern similarity for all possible pairs of texture responses. The similarity metric was the dot product between the neuronal unit activity vectors (normalized to length 1) for each neuron. There was high similarity between activity patterns representing the same texture (mean similarity was 0.72, S.E. = 0.20), but not between activity patterns representing different textures (mean similarity was 0.34, S.E. = 0.12).

Perceptual categorization by Darwin IX involved multiple spatiotemporal transformations. First, a spatially defined stimulus, such as a texture, was transformed into temporally arranged sensory input as a result of whisker deflections during movement. Second, this input was reformed into spatial patterns of activity in area $S2$, corresponding to perceptual categories, as a result of the response properties of the "lag" cells in the barrel regions $Th \rightarrow S1 \rightarrow S2$. The response properties of units in $S2$ were similar to the complex spatiotemporal receptive fields that have been found in rat somatosensory cortex (Ghazanfar & Nicolelis, 1999).

Observations of Darwin IX showed that time-lagged neuronal responses to somatosensory input provide a plausible mechanism for spatiotemporal pattern recognition. The rodent brain may use multiple pathways in order to achieve ecologically significant discriminations, and it is worth noting that the present model is not incompatible with alternatives such as the phase locked loop model of whisker processing (Ahissar & Arieli, 2001). The significance of Darwin IX's performance for the present discussion rests not so much in a particular solution to the problem of texture discrimination but in the illustration of a neural mechanism that is able to integrate information across space and across time. The key insight is that a population of neurons with variable response latencies constitutes a repertoire that provides a mapping between temporal interval patterns and spatial response patterns. The proposed 'lag' cells that mediate this mapping provide a modality-specific 'memory' of stimulus events over short timescales, which in Darwin IX permits texture discrimination and selective conditioning to textures.

In the next section we discuss how a similar emphasis on temporal processing using delays can support prediction in the context of motor control.

3. Predictive motor control in a BBD

Recent theories of motor control suggest that the cerebellum learns to replace primitive reflexes with a predictive motor signal. The idea is that the reflexive motor commands provide an error signal for a predictive controller, which then learns to produce a correct motor control signal prior to the less adaptive reflex response (Kawato & Gomi, 1992; Kettner et al., 1997; Medina, Carey, & Lisberger, 2005; Wolpert, Miall, & Kawato, 1998; Worgotter & Porr, 2005). Synaptic eligibility traces in the cerebellum have been proposed as a mechanism to bridge the temporal gap between predictive signals and subsequent reflexive motor responses (Kettner et al., 1997). In this idea, suprathreshold presynaptic activity causes a synapse to be eligible for plasticity, and the amount of potential synaptic change decays over time until an error signal arrives at the synapse resulting in synaptic strength modification. We tested an alteration of this idea by introducing a delayed eligibility trace learning rule, in which synapses become eligible for plasticity only after a fixed delay from the onset of suprathreshold presynaptic activity. In our model, these synapses were from neuronal units that responded to visuomotor stimuli onto Purkinje cells or onto cells in the deep cerebellar nucleus.

We tested the generality of this proposed aspect of cerebellar function in a BBD with separate reflexes for turning and for braking. The BBD's task was to traverse a curved course outlined by traffic cones, without collisions. Initially, collisions or near collisions with the cones generated a reflexive movement away from the obstacle and a reflexive braking response. These reflex commands were also used as error signals to the cerebellar model via simulated climbing fiber inputs. Success in this task required the BBD's cerebellum to associate predictive visual motion cues, which came from optic flow generated by self-movement, with the correct movements to avoid collisions with the cone boundaries.

3.1. Predictive motor control: Construction and experimental paradigm

The BBD used for the motor control experiments was built on the Segway Robotic Mobility Platform (RMP), a commercially available robotic version of the Segway Human Transporter scooter (Fig. 4A). The device received sensory input from a color camera and banks of short range IR proximity sensors that are mounted low around the device to detect nearby objects (Fleischer et al., in press). An aluminum chassis on the commercial base contained a cluster of six compact Pentium IV PCs and enough battery capacity to operate for approximately 45 min.

The device moved forward in a straight line at a maximum of 1.25 m per second unless a cone was detected directly in front of the device by the IR detectors, a wall was detected, or activity of the neural simulation caused the device to slow down and/or turn. An array of IR proximity detectors signaled the presence of cones within 6 inches of the device. These signals drove the turning and braking reflexes, causing the device to slow down and move away from the obstacle, and sent an error signal via the simulated inferior olive to the cerebellum (see *IO-Turn* and *IO-Velo* in Fig. 5). The RMP turn rate was set based upon the population activity in the *Motor-Turn* neural area; activity on the left resulted in a turn to the right and vice versa (see Fig. 5). The speed of the robot was controlled based on the activity of the *Motor-Velo* neural area (see Fig. 5). The *Motor-Velo* area slowed down the RMP based on the number of IR detectors signaling an obstacle; the more the IR detectors that were activated, the lower the velocity. These primitive reflexes were sufficient to drive the BBD down the course, but not without collisions.

Fig. 5 shows a high-level diagram of the simulated nervous system including the various neural areas and the arrangement

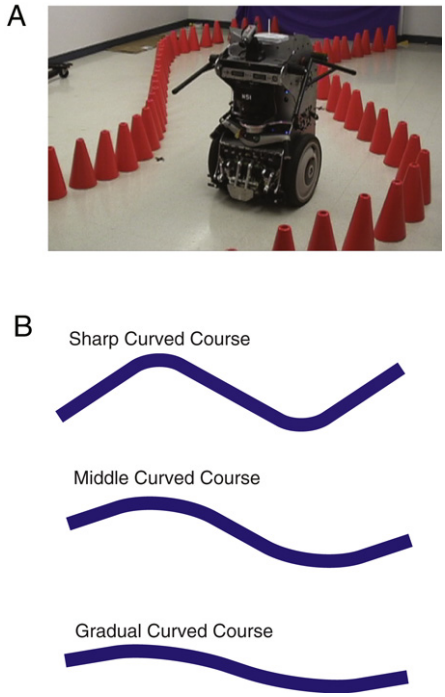

Fig. 4. The Segway BBD and its environment. A. The BBD is built on the Segway Robotic Mobility Platform. The device navigated a path dictated by the traffic cones that were spaced apart by a few inches. B. The diagram shows the layout of the different courses. The lane dictated by the cones was five feet wide and roughly 25 feet long. The device itself was approximately two feet in diameter.

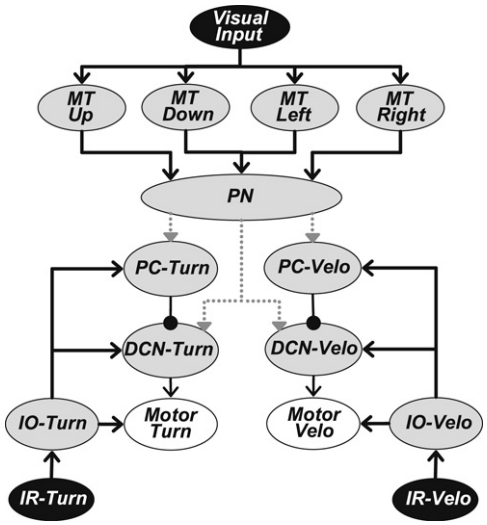

Fig. 5. Schematic of the regional and functional neuroanatomy of the BBD. Gray ellipses denote different neural areas, black ellipses denote sensory input areas, and white ellipses denote motor areas. Arrows denote synaptic projections from one area to another. Black arrows ending in open arrowheads denote excitatory connections, black arrows ending in a circular endpoint denote inhibitory connections, and gray arrows ending in filled arrowheads with dotted lines denote plastic connections. Visual input from a camera on the BBD projected to cortical area *MT*. The simulated cerebellar region consisted of a precerebellar nuclei (*PN*), Purkinje cells (*PC-Turn* and *PC-Velo*), deep cerebellar nuclei (*DCN-Turn* and *DCN-Velo*), and input from the inferior olive (*IO-Turn* and *IO-Velo*) where "*Velo*" refers to velocity. Neuronal units in the inferior olive were driven by the IR proximity detectors, which in turn drove motor neurons for turning (*Motor-Turn*) and braking (*Motor-Velo*). Motor neurons were also driven by *DCN*.

of synaptic connections. Similar to Darwin IX, each area contains neuronal units that can be either excitatory or inhibitory, each of which represents a local population of neurons (Edelman, 1987),

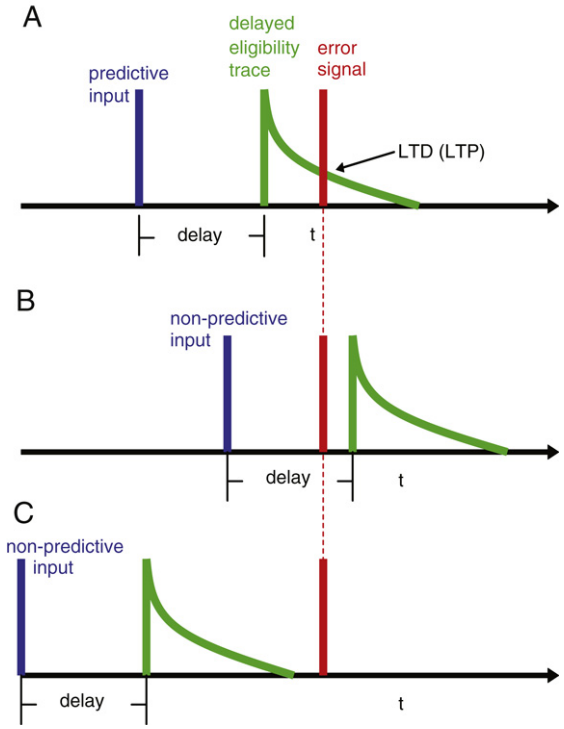

Fig. 6. The delayed eligibility trace learning rule creates a temporal window during which predictive inputs may arrive in order to be linked to motor commands. A. Predictive input arrives within the time window for synaptic change. B. Input arrives too late to effect synaptic change. C. Input arrives too early to effect synaptic change.

in which the mean-firing rate variable of each unit corresponds to the average activity of a group of roughly 100 real neurons during a time period of approximately 40 ms. The network contained 28 neural areas, 27,688 neuronal units, and 1.6 million synapses, and was updated in real time. For complete details, see McKinstry, Edelman, and Krichmar (2006).

The simulated cerebellum had precerebellar nuclei (*PN*) that received input from cortical areas that responded to visual motion (*MT* → *PN* in Fig. 5) and output to the cerebellum (*PN* → *PC*, *PN* → *DCN* in Fig. 5). A simulated cerebellar cortex contained Purkinje cells that inhibited deep cerebellar nuclei affecting turning and velocity (*PC* → *DCN* in Fig. 5), and an inferior olive that simulated climbing fiber input to the cerebellum (*IO* → *PC*, *IO* → *DCN* in Fig. 5). The deep cerebellar nuclei area projected to the motor neural area resulting in RMP motor commands (*DCN* → *Motor* in Fig. 5).

In order to turn a reflex into a predictive motor response, or "preflex", a mechanism was necessary to maintain predictive information until an error signal arrived. A simple mechanism would be to maintain a decaying trace of activity at each synapse (Sutton & Barto, 1998). However, such a mechanism strengthens the synapses with prior input in proportion to their temporal proximity to the arrival time of the error signal. This heuristic may be inadequate in many situations where temporally distant sensory-motor signals predict the error signal.

As a modification to the eligibility trace, we implemented a delayed eligibility trace, which created a window of time during which predictive input could occur in order to be associated with a motor response. Fig. 6 demonstrates schematically how such a delayed eligibility trace learning rule works. When a suprathreshold input arrives in the precerebellar nucleus, an eligibility trace is triggered after a certain delay, which was varied parametrically in our experiments. Fig. 6A illustrates how an input arriving at the proper time before an error signal arrives will trigger

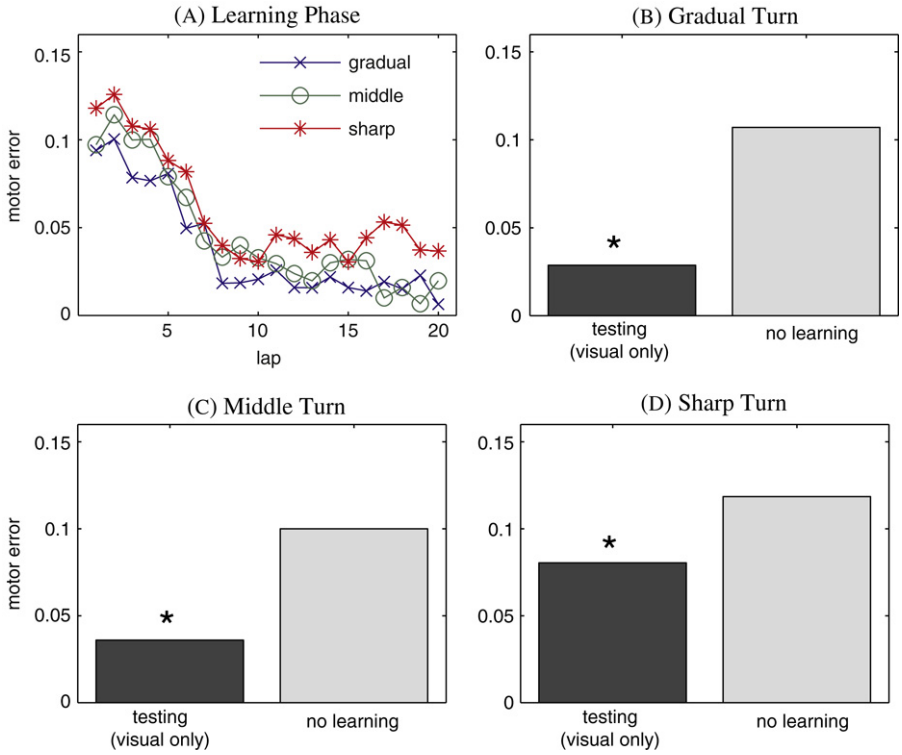

Fig. 7. Training and testing the device on different curved courses. The mean motor error from five subjects is shown in the plots. A. Learning curves during training on the gradual turn, sharp turn, and middle turn courses. B–D. Motor errors were significantly lower in the test group, which had access to only visual cues, than in the control or "no learning" groups on the Gradual course (B), Middle course (C), and Sharp course (D). $*$ denotes $p < 0.01$ Wilcoxon Rank Sum test.

learning. This happens when the delayed trace, shown in green, is above zero. When the signal arrives too early or too late as in Fig. 6B and C, synaptic change is diminished. For specific details on synaptic strength modification with the delayed eligibility trace, see Appendix B.

3.2. Predictive motor control through delayed synaptic eligibility

After learning, the BBD's visual responses predicted future collisions and resulted in smooth movement down the middle of the pathways marked by the cones. Successful performance across the three different courses with varying turns (see Fig. 4b) required a combination of braking and turning of the proper magnitude at the proper time. A 4 s delay incorporated into the delayed eligibility trace learning rule was sufficient for successful navigation on all three courses (Fig. 7). Significantly longer or shorter delays were not predictive in this particular task. Subjects learned to slow down prior to and during turns, and they learned to turn in the proper direction at the proper time. Subjects on the sharp course, which contained roughly 90° turns, had slightly worse performance than on the other courses. Nevertheless, in the testing phase, subjects with cerebellar learning performed significantly better on all three courses than did subjects without learning (Fig. 7B–D).

Subjects adapted their behaviors to the particulars of each course. For example, subjects were faster on the gradual course than on the sharp course. Success on the sharp course required slower speed and more frequent turning to the left or the right. Subjects on the gradual course typically proceeded at maximum velocity on the straightaway, and simultaneously slowed and turned slightly on the curves. Learning on one course generalized to others. For example, subjects trained on the sharp course were tested on the gradual course and vice versa. In both cases, initially trained subjects showed significantly better performance on the early training laps (e.g. laps 1–6) than naïve subjects ($p < 0.005$ one-tailed t-test). Adapting from the gradual to the sharp course, however, required additional training to reach peak performance.

The delayed eligibility trace allowed the BBD's behavior to move from reflexive control to predictive control and bridge the temporal gap between predictive signals and the arrival of a subsequent error signal. Evidence for such a mechanism has been found in the cerebellum (Chen & Thompson, 1995; Wang et al., 2000), and Kettner et al. (1997) have used a similar rule in computer simulations of the vestibular-occular reflex. We tested the sufficiency of this mechanism in the real world where successful performance required adaptation of turning and braking commands. Due to the delayed eligibility trace, experience resulted in a forward shift in time in neuronal responses triggering a conditioned response: initially, collision error signals drove the motor neurons via reflex pathways. After learning, visual input drove motor neurons prior to any error signal, similarly to the well-known eye-blink conditioning paradigm (Medina & Mauk, 2000). The BBD's ability to traverse curved courses was attributed to the delayed eligibility trace with the appropriate delay and the system was also able to generalize its learning from one course to another. Our findings provide additional support for the theory that the cerebellum can learn to replace an arbitrary reflexive neural control system with an adaptive, predictive controller or "preflex".

4. Discussion

Biological systems cannot afford to store all signals indefinitely. Thus a recurring theme in neuroscience is how the nervous system maintains and integrates relevant signals across different timescales.

In this paper we have described two models that respond to real-world signals by maintaining those signals across temporal delays using two different mechanisms found in the nervous system. One strategy is through delays in neuronal responses such

that spatiotemporal stimuli can be transformed into categorical patterns of activity (Beggs et al., 2000; Saul & Humphrey, 1992; Wolfe & Palmer, 1998). This strategy was effective in converting artificial whisker signals into patterns of activity that facilitate texture discrimination. Another strategy is to delay the eligibility of synaptic plasticity such that an environmental cue is associated with a prior signal (Kawato & Gomi, 1992; Kettner et al., 1997; Medina et al., 2005; Wolpert et al., 1998; Worgotter & Porr, 2005). This mechanism allowed a visuomotor stimulus to predict a future collision and resulted in a preemptive motor action. In both cases, neural signals with lags or delays bridged temporal gaps and allowed the BBD to adapt its behavior appropriately.

Lagged responses in the somatosensory system. Spatiotemporal transformations are fundamental to neural operations underlying adaptive behavior. Nowhere is this more evident than in the somatosensory system of the rat in which spatiotemporal signals arise due to self-movement, the texture pattern, and the dynamics of the tactile sensor. Darwin IX illustrates a neural mechanism, based on the rat somatosensory system, which is capable of orchestrating spatiotemporal transformations of whisker signals to allow texture discrimination, and selective conditioning to textures, in a real world environment.

Perceptual categorization by Darwin IX involved multiple spatiotemporal transformations. First, a spatially defined stimulus (a texture) is transformed into temporally arranged sensory input as a result of whisker deflections during movement. Second, this input is reformed into spatial patterns of activity in the simulated somatosensory area *S*2, corresponding to perceptual categories, as a result of the response properties of the 'lag' cells in the barrel regions of simulated ventromedial thalamus (*Th*) and the connectivity in the pathway $Th \rightarrow S1 \rightarrow S2$. As a result, neuronal units in *S*2 respond to specific combinations of whisker deflections with particular poststimulus delays. Analysis of neural activity in *S*2 revealed the formation of spatiotemporal activity patterns corresponding to specific haptic perceptual categories. The response properties of these units are analogous to cells with complex spatiotemporal receptive fields that have been found in rat somatosensory cortex (Ghazanfar & Nicolelis, 1999) as well as in cat visual cortex (DeAngelis, Ohzawa, & Freeman, 1995).

Delayed synaptic change in the cerebellum. Building a device that can move with the grace and dexterity of an animal requires smooth movements that preempt the awkward reflexive movements that arise from collisions or errors. The idea that a simple, feedback controller can be used for these reflexes, and eventually replaced, through experience, by a predictive controller, was experimentally confirmed in a BBD for two arbitrary reflexes, turning and braking. Predictive cerebellar activity resulted in smooth motor responses that precluded awkward reflex responses and collisions, supporting a recent theory of cerebellar function (Wolpert et al., 1998).

An important issue is the mechanism by which the predictive signal is associated with the motor response. Since the predictive signal disappears before the arrival of the reflexive error signal in our path following task, we assumed an eligibility trace at each synapse which maintained predictive signals for a certain time window long enough to bridge the temporal gap between predictive signals and error signals. In our experiments, delaying the onset of the eligibility trace was critical, preventing the unwanted association between signals occurring at the time of the reflex (which were too late), and the motor response. Behavioral support for a delay exists in the eye-blink conditioning literature. Pairing a tone with an unconditioned stimulus is ineffective if the tone occurs less than 80 ms prior to the unconditioned stimulus (Medina & Mauk, 2000).

As with Darwin IX, it was necessary to adjust the neural timescale of the cerebellar model to match the behavioral timescale which was on the order of seconds. Although the optimal eligibility trace delay found for our system was 4 s, it is likely that different optimal delays would be found depending on the reflex and the environmental conditions. Since the reflex involved visual feedback for a slow-moving device, longer eligibility trace delays were necessary. Indeed, we have recently used the same delayed eligibility trace in a much faster robot during a sense and avoid task, and found that the optimal eligibility trace delay was only one second (unpublished data). In both cases, the cerebellar system formed an association between a motor response and sensory signals such that the device avoided potential collisions.

An alternative to the notion that the cerebellum has fixed delays of varying durations to handle multiple timescales is the proposal that the cerebellum always uses the same 80 ms delay due to the use of discrete, pulsatile control signals occurring at a frequency of 8–12 Hz (Lang, Sugihara, & Llinas, 2006; Vallbo & Wessberg, 1993). A delay of 80 ms would be roughly accurate if the cerebellum had to predict the next motor response given the prior motor command and other predictive inputs.

Summary. Both the whisker and the cerebellar model take into account delays between input stimuli and motor output. Another method for bridging this gap is reinforcement learning which is often used in the robotics community (Schaal, 2002; Sutton & Barto, 1998) and has a neurobiological correlate (Schultz, Dayan, & Montague, 1997). In reinforcement learning, a value function mapping the state of the motor plant and world to value is learned. This function can then be used to derive an optimal control policy mapping the current state to motor control signals. The method of temporal differences along with an eligibility trace is often employed to overcome the delay between the reinforcement (such as a collision) and the commands which lead to it. Our model of the cerebellum as a predictive controller differs from reinforcement learning in several ways. First, the eligibility trace has a delay before onset which we found to be necessary for improved performance in our task. Second, rather than learning a separate value function, the system learned a mapping from the state of the environment directly to a motor control signal. This was possible since each reflex generated the correct command, and all that was necessary was to elicit the command earlier in the form of a "preflex". This general approach may explain why so many learned reflexes in animals are dependent upon the cerebellum. The model of the whisker system used in Darwin IX used "lag cells" to maintain a record of stimuli long enough for spatiotemporal categorization. However, the task requires an immediate motor response (freezing) once the spatiotemporal pattern was categorized, therefore a simple error correction rule was sufficient to form the correct association between the stimulus and the response.

Arguably, the BBD methodology provides the ideal test for integrating the dynamic signals from the environment into coherent percepts since the device interacts with a real environment in real time thus forcing the critical issue of bridging the temporal gap between unfolding events. We have used this approach to provide stringent tests of two of these mechanisms, lag cells and delayed eligibility traces. The two types of memory tested in our experiments represent two classes of mechanisms. Lag cells represent the class of mechanisms that maintain neuronal output across time, while the delayed eligibility trace rule is an example of a mechanism for preserving synaptic input over time.

Acknowledgements

We thank J. Gally for useful suggestions and discussion; and J. Snook, D. Moore, and D. Hutson for building the Brain-Based Device. This work was supported by grants from the Office of Naval Research, Defense Advanced Research Projects Agency, and the Neurosciences Research Foundation.

Appendix A. Thalamic lag cells

Each lag cell is characterized by an internal state (s_i^{in}), an output (s_i), and a cell-specific lag parameter set to be $\psi_i = \frac{0.2}{i}, i \in \{1, 2, \ldots, 20\}$ for cell i in each barrel. When triggered by a whisker deflection, the internal s_i^{in} state of cell i in the corresponding barrel increases at a rate determined by ψ_i. When this internal state reaches a threshold, the cell begins to emit an output signal and s_i^{in} is reset to zero. Due to differences in ψ_i among lag cells, each whisker deflection evokes a wave of activity in the corresponding barrel, with some cells firing shortly after deflection and the remainder firing with gradually increasing delays (see Fig. 2; inset).

Specifically, the internal state of each lag cell i, in the barrel corresponding to whisker k, is updated according to:

$$s_{ki}^{\text{in}}(t+1) = \begin{cases} 0.2; & s_{ki}^{\text{in}}(t) < 0.2, \overline{\text{diff}}_k(t) > 3.0 \\ 0; & s_{ki}^{\text{in}}(t) \geqslant \sigma_i^{\text{fire}} \\ (1+\psi_i)(s_{ki}^{\text{in}}(t)); & \text{otherwise} \end{cases} \quad (1)$$

where $\overline{\text{diff}}_k(t)$ is the difference between successive whisker readings averaged over the last four samples (a value exceeding 3.0 signifies a whisker deflection), and σ_i^{fire} is a firing threshold set to 0.3.

The output s_{ki} is calculated using:

$$s_{ki}(t+1) = \begin{cases} \tanh\left(10\left(\omega_i(s_{ki}(t))\right)\right); & s_{ki}^{\text{in}}(t) < \sigma_i^{\text{fire}} \\ \tanh\left(10\left(\omega_i(s_{ki}(t)) + (1-\omega_i)s_{ki}^{\text{in}}(t)\right)\right); & \\ \quad \text{otherwise} & \end{cases} \quad (2)$$

where $\omega_i = 0.8$ determines the persistence of unit activity. This value is fed as input into neuronal units in the corresponding barrels of $S1$.

Appendix B. Delayed eligibility trace

More specifically, synaptic strengths were subject to modification according to a synaptic rule that depends on the presynaptic, postsynaptic, and inferior olive activities.

Synaptic changes are given by:

$$\Delta c_{ij}(t+1) = \eta s_i(t) \cdot \text{trace}_{\text{eligibility}}(t) \cdot (IO_i(t) - 0.02); \quad (3)$$

where c_{ij} is the connection strength from unit j to unit i, $s_i(t)$ is the activity of the postsynaptic unit, $IO_i(t)$ is the activity of the inferior olive unit corresponding to unit i, η is a fixed learning rate, and $\text{trace}_{\text{eligibility}}(t)$ is the eligibility trace of synapse j. The eligibility trace (see below) determines the amount of efficacy change at a specific synapse for a given time. This learning rule supports both potentiation and depression at *PC* and *DCN* synapses (Hansel, Linden, & D'Angelo, 2001). When η was negative (e.g. in $PN \rightarrow PC$ synapses), the learning rule induced depression when the *IO* was active above a baseline firing rate, and potentiation when *IO* was below the baseline (Ohyama, Medina, Nores, & Mauk, 2002). Note that this learning rule supported extinction of learned responses when the error from the *IO* was absent (Mauk, Medina, Nores, & Ohyama, 2000; Seth et al., 2004b).

In the model, the change in synaptic efficacy was based on the delayed eligibility trace rule, according to which an eligibility trace ($\text{trace}_{\text{eligibility}}$) determined the amount of synaptic change at that synapse when eligible:

$$\text{trace}_{\text{eligibility}}(t+1) = \begin{cases} 0 & \text{if } t < \text{delay}, \\ s(t-\text{delay}) & \text{if } s(t-\text{delay}) \geqslant \sigma, \\ 0.90 \cdot \text{trace}_{\text{eligibility}}(t) & \text{otherwise} \end{cases}, \quad (4)$$

where $s(t)$ is the presynaptic input to the synapse at time t, $\sigma = 0.15$, and *delay* is a time offset from the previous simulation cycle. (Note that once $s(t - \text{delay}) \geq \sigma$ was used to trigger synaptic eligibility, further input was ignored until $\text{trace}_{\text{eligibility}}(t + \Delta) < 0.1$ where Δ is the time offset after *delay*). When presynaptic input exceeds a threshold, the synapse becomes eligible for modification after a set delay, at which time, the eligibility declines exponentially.

References

Ahissar, E., & Arieli, A. (2001). Figuring space by time. *Neuron*, *32*, 185–201.

Almassy, N., Edelman, G. M., & Sporns, O. (1998). Behavioral constraints in the development of neuronal properties: A cortical model embedded in a real-world device. *Cerebral Cortex*, *8*, 346–361.

Beggs, J. M., Moyer, J. R., Jr., McGann, J. P., & Brown, T. H. (2000). Prolonged synaptic integration in perirhinal cortical neurons. *Journal of Neurophysiology*, *83*, 3294–3298.

Carr, C. E., & Konishi, M. (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *Journal of Neuroscience*, *10*, 3227–3246.

Chen, C., & Thompson, R. F. (1995). Temporal specificity of long-term depression in parallel fiber—Prukinje synapses in rat cerebellar slice. *Learning and Memory*, *2*, 185–198.

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in Neurosciences*, *18*, 451–458.

Edelman, G. M. (1987). *Neural darwinism: The theory of neuronal group selection*. New York: Basic Books, Inc.

Edelman, G. M., Reeke, G. N., Gall, W. E., Tononi, G., Williams, D., & Sporns, O. (1992). Synthetic neural modeling applied to a real-world artifact. *Proceedings of the National Academy of Sciences of the United States of America*, *89*, 7267–7271.

Fleischer, J., Szatmary, B., Hutson, D., Moore, D., Snook, J., Edelman, G. M., et al. (2006). A neurally controlled robot competes and cooperates with humans in Segway soccer. In *Robotics and Automation, Proceedings 2006 IEEE International Conference* (pp. 3673–3678).

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*, 331–349.

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*, 652–654.

Ghazanfar, A. A., & Nicolelis, M. A. (1999). Spatiotemporal properties of layer V neurons of the rat primary somatosensory cortex. *Cerebral Cortex*, *9*, 348–361.

Hansel, C., Linden, D. J., & D'Angelo, E. (2001). Beyond parallel fiber LTD: The diversity of synaptic and non-synaptic plasticity in the cerebellum. *Nature Neuroscience*, *4*, 467–475.

Harvey, M. A., Bermejo, R., & Zeigler, H. P. (2001). Discriminative whisking in the head-fixed rat: Optoelectronic monitoring during tactile detection and discrimination tasks. *Somatosensory and Motor Research*, *18*, 211–222.

Jagadeesh, B., Wheat, H. S., Kontsevich, L. L., Tyler, C. W., & Ferseter, D. (1997). Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *Journal of Neurophysiology*, *78*, 2772–2789.

Jensen, K. F., & Killackey, H. P. (1987). Terminal arbors of axons projecting to the somatosensory cortex of the adult rat. I. The normal morphology of specific thalamocortical afferents. *Journal of Neuroscience*, *7*, 3529–3543.

Kawato, M., & Gomi, H. (1992). A computational model of four regions of the cerebellum based on feedback-error learning. *Biological Cybernetics*, *68*, 95–103.

Kettner, R. E., Mahamud, S., Leung, H. C., Sitkoff, N., Houk, J. C., Peterson, B. W., et al. (1997). Prediction of complex two-dimensional trajectories by a cerebellar model of smooth pursuit eye movement. *Journal of Neurophysiology*, *77*, 2115–2130.

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., & Edelman, G. M. (2005). Brain-based devices for the study of nervous systems and the development of intelligent machines. *Artificial Life*, *11*, 63–77.

Krichmar, J. L., Nitz, D. A., Gally, J. A., & Edelman, G. M. (2005). Characterizing functional hippocampal pathways in a brain-based device as it solves a spatial memory task. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 2111–2116.

Lang, E. J., Sugihara, I., & Llinas, R. (2006). Olivocerebellar modulation of motor cortex ability to generate vibrissal movements in rat. *Journal of Physiology*, *571*, 101–120.

LeDoux, J. E. (1995). Emotion: Clues from the brain. *Annu Rev Psychol*, *46*, 209–235.

Maren, S., & Fanselow, M. S. (1996). The amygdala and fear conditioning: Has the nut been cracked? *Neuron*, *16*, 237–240.

Mauk, M. D., Medina, J. F., Nores, W. L., & Ohyama, T. (2000). Cerebellar function: Coordination, learning or timing? *Current Biology*, *10*, R522–525.

McKinstry, J. L., Edelman, G. M., & Krichmar, J. L. (2006). A cerebellar model for predictive motor control tested in a brain-based device. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 3387–3392.

Medina, J. F., Carey, M. R., & Lisberger, S. G. (2005). The representation of time for motor learning. *Neuron*, *45*, 157–167.

Medina, J. F., & Mauk, M. D. (2000). Computer simulation of cerebellar information processing. *Nature Neuroscience*, *3*(Suppl), 1205–1211.

Ohyama, T., Medina, J. F., Nores, W. L., & Mauk, M. D. (2002). Trying to understand the cerebellum well enough to build one. *Annals of New York Academi of Sciences*, *978*, 425–438.

Prigg, T., Goldreich, D., Carvell, G. E., & Simons, D. J. (2002). Texture discrimination and unit recordings in the rat whisker/barrel system. *Physiology & Behaviour*, *77*, 671–675.

Saul, A. B., & Humphrey, A. L. (1992). Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *Journal of Neurophysiology*, *68*, 1190–1208.
Schaal, S. (2002). Learning robot control. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd edition) (pp. 983–987). MIT Press.
Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
Seth, A. K., McKinstry, J. L., Edelman, G. M., & Krichmar, J. L. (2004a). Visual binding through reentrant connectivity and dynamic synchronization in a brain-based device. *Cerebral Cortex*, *14*, 1185–1199.
Seth, A. K., McKinstry, J. L., Edelman, G. M., & Krichmar, J. L. (2004b). Active sensing of visual and tactile stimuli by brain-based devices. *International Journal of Robotics and Automation*, *19*, 222–238.
Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: The MIT Press.
Vallbo, A. B., & Wessberg, J. (1993). Organization of motor output in slow finger movements in man. *Journal of Physiology*, *469*, 673–691.
Wang, S. S., Denk, W., & Hausser, M. (2000). Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience*, *3*, 1266–1273.
Wolfe, J., & Palmer, L. A. (1998). Temporal diversity in the lateral geniculate nucleus of cat. *Visual Neuroscience*, *15*, 653–675.
Wolpert, D., Miall, C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*, 338–347.
Woolsey, T. A., & Van der Loos, H. (1970). The structural organization of layer IV in the somatosensory region (SI) of mouse cerebral cortex. The description of a cortical field composed of discrete cytoarchitectonic units. *Brain Research*, *17*, 205–242.
Worgotter, F., & Porr, B. (2005). Temporal sequence learning, prediction, and control: A review of different models and their relation to biological mechanisms. *Neural Computing*, *17*, 245–319.