Object recognition, Adaptive Behavior and Learning in Brain-Based Devices

Jeffrey L. Krichmar, Douglas A. Nitz, Gerald M. Edelman
The Neurosciences Institute, 10640 John J. Hopkins Drive, San Diego, CA, USA
krichmar@nsi.edu, nitz@nsi.edu, edelman@nsi.edu

Abstract

An open question in neuroscience is how animals combine the various attributes of stimuli in their environments into coherent perceptual categories and how they discriminate among objects in a scene. Testing a theory of visual binding would require the simultaneous study of brain function at many levels of organization. Present day electrophysiology only allows the recording of at most hundreds of neurons while an animal is performing a behavioral task. Because of this limitation and the sheer complexity of the nervous system, computational modeling has become essential for investigating theories of brain function. Accordingly, our group has constructed a series of brain-based devices (BBDs); i.e. physical devices with simulated nervous systems that guide behavior, to serve as heuristic bases for testing theories of brain function. Unlike animal models, BBDs permit analysis of activity at all levels of its nervous system as the device behaves in its real environment. We present a possible solution to the binding problem based on synchronous activity across neuronal groups brought about by reentrant connectivity. We first show the sufficiency of this theory in a laboratory setting and then demonstrate that these principles can be transferred to a more real-world application: robots capable of playing a game of soccer with humans.

1. Introduction

Animals effortlessly combine the various attributes of visual stimuli to form coherent perceptual categories and to discriminate among multiple objects in a scene. Yet the visual brain is functionally segregated: Separate cortical regions are specialized to respond to features such as shape, color, and object motion, and no single region has superordinate control. This poses the so-called binding problem [1]: How do these functionally segregated regions coordinate their activities to link various features of individual objects while distinguishing among different objects? Most proposed mechanisms for solving the binding problem fall into one of two general classes: 1) binding through the influence of attentional processes, executive mechanisms, or superordinate maps [2, 3]. 2) binding

through the selective synchronization of dynamically formed neuronal groups [4-6]. Advocates of neural synchrony suggest that binding is an automatic, dynamic, and pre-attentive process arising from low-level neural dynamics. For example, the linkage of neuronal groups by reentry, the recursive exchange of signals across multiple, parallel and reciprocal connections [7], can lead to selective synchronization [8-11]. Synchronization of activity among neuronal groups can form coherent circuits corresponding to perceptual categories [8]. A fundamental challenge for proponents of neural synchrony is to show how such emergent functional circuits contribute to an organism's adaptive behavior, especially in situations that require preferential behavior towards one object among many in a scene.

Elucidation of brain mechanisms underlying behavior, such as visual binding followed by discriminatory action, requires simultaneous measurements across multiple levels. The heuristic value of synthetic modeling using brain-based devices (BBDs), which are described here, is supported by the fact that such types of measurements are difficult to obtain and compare in living animals. Given the successful construction of BBDs [12-14], we observe their overall behavior while simultaneously recording the state of all components of their simulated nervous systems. Since our purpose is to test theories of real nervous systems in order to arrive at a better understanding of brain function, we base the BBD's organization on real neuroanatomy and physiology.

We argue that a BBD should be constrained by the following design principles: 1) The device needs to engage in a behavioral task. 2) The device's behavior must be controlled by a simulated nervous system having a design that reflects the brain's architecture and dynamics. 3) The device needs to be situated in the real-world [15, 16]. 4) The behavior of the device and the activity of its simulated nervous system must allow comparisons with empirical data. Because of these constraints, BBD simulations tend to require large-scale networks of neuronal elements that reflect vertebrate brain architecture and dynamics, high performance computing to run the network in real-time, and the engineering of specialized physical devices to embody the network.

BBDs are not programmed by instructions like computers, but instead, like biological systems, they operate according to selectional principles that allow them to adapt to the environment [7]. Their design, which possesses neuroanatomical structure and large-scale neural dynamics, differs fundamentally from that of robots. Robotic approaches using classical artificial intelligence are based on data representation, rule-driven algorithms, and the manipulation of formal symbol systems.

BBDs must have a morphology or body plan that allows for active exploration in a real environment with a brain simulation controlling the BBD's behavior. Changes in the nervous system that result in lasting modifications of the device's behavior are realized through a neuromodulatory value system that signals the salience of environmental cues triggering broad changes to the BBD's nervous system. These features yield a system that generalizes signals from the environment into perceptual categories and adapts its behavior so that it becomes increasingly successful in coping with its environment. The BBDs have been designated the Darwin series of automata. Over the last 12 years, various Darwin automata have been shown to develop perceptual categorization, invariant visual object recognition, integration of scenes containing multiple visual shapes with overlapping features, fusion of different sensory modalities, and learning in the form of operant conditioning [12-14, 17, 18].

In this paper we describe two recent BBDs that address the problem of visual binding, scene segmentation, and motor behavior. First, we describe Darwin VIII, a BBD that demonstrated visual binding through synchronous activity across cortical areas brought about by reentrant signalling. Second, we apply the principles of Darwin VIII to a novel platform, the Segway Robotic Mobility Platform (RMP), in a dynamic environment, namely a soccer game.

## 2. Visual binding in a laboratory setting.

Darwin VIII, a BBD incorporating an extensive visual cortical neural simulation, demonstrated the ability to parse a scene composed of ambiguous visual shapes into separate and coherent perceptual categories. It solved the so-called binding problem, that is, it linked responses in different brain areas and modalities to yield selective responses to percepts in the absence of any superordinate control from a master or executive brain area [1]. The behavior of Darwin VIII exploited interaction between neural areas, and revealed that reentrant activity (i.e. ongoing reciprocal excitatory activity brought about by connections between neuronal units in different neural areas) is sufficient for recognizing and distinguishing among multiple objects in a scene.

Darwin VIII was designed to demonstrate visual categorization and selective conditioning in a rich environ-

ment [18]. This BBD had a camera for vision, microphones to pick up auditory cues from the environment, and infrared (IR) sensors to detect the boundaries of the environment.

2.1. Experimental Setup

Figure 1 shows Darwin VIII in its environment, which consists of an enclosed area with various shapes hung on two walls. Near the walls with visual shapes, infrared beams were set up that control speakers (see Figure 1A). When Darwin VIII's movement broke an IR beam, a tone was emitted. Darwin VIII reflexively oriented towards the sound source and gradually came to associate the sound with the object it saw near the sound source. After conditioning, the sound is no longer necessary; Darwin VIII approaches visual objects that have become associated with preferred sounds.
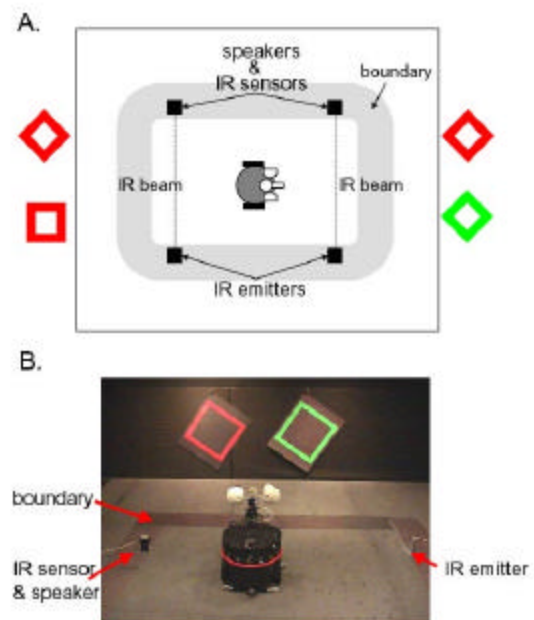
Figure 1. Experimental setup for Darwin VIII. A. Darwin VIII views objects on two of the walls of an arena. The area Darwin VIII explores is constrained by a boundary of reflective construction paper. Detection of this boundary by infrared sensor triggers a reflexive turn. When Darwin VIII breaks the beam from the IR emitter to the IR sensor, a tone is emitted from the speaker. B. Photograph of the experimental environment.

2.2. Simulated Nervous System

Darwin VIII's simulated nervous system contains areas corresponding to cortical and sub-cortical areas in the vertebrate nervous system (see Figure 2). Specifically, Darwin VIII's brain includes simulated cortical areas of

the visual system that respond to shape and color (*V1→V2→V4→IT*), a motor system (*C*), an auditory system (*Mic-left→Aleft, Mic-right→Aright*), and a value system (*S*). Activity in *S* is analogous to that of ascending neuromodulatory systems in that it is triggered by salient events, influences large regions of the simulated nervous system, and persists for several cycles [19]. Due to its projection to the tracking area *C*, area *S* has a direct influence on behavior.

Neuronal units in Darwin VIII roughly correspond to the activity of 100 real neurons over 100 ms. The neuronal units have a firing phase parameter, which specifies the relative timing of this activity within each simulation cycle (for details, see [8, 18]). This modeling feature provides temporal specificity without incurring the computational costs associated with modeling spiking neurons in real time. Simulated synaptic connections follow known vertebrate neuroanatomical projections (arrows in Figure 2) and include extensive reentrant connectivity within and among neural areas. In Darwin VIII, reentrant connections among neuronal units encourage phase coherence and therefore lead to the emergence of neural synchrony.

Synaptic strengths are subject to modification according to a synaptic rule that depends on the phase and activities of the pre- and postsynaptic neuronal units. Plastic synaptic connections are either value-independent (see *IT→IT* in Figure 2) or value-dependent (see *IT→S*, *IT→C* in Figure 2). Both of these rules are based on a modified BCM learning rule [20] in which thresholds defining the regions of depression and potentiation are a function of the phase difference between the presynaptic and postsynaptic neuronal units (for details, see [18]). Synapses between neuronal units with strongly correlated firing phases are potentiated and synapses between neuronal units with weakly correlated phases are depressed; the magnitude of change is determined as well by pre- and postsynaptic activities. This learning rule is similar to a spike-time dependent plasticity rule [21] applied to jittered spike trains where the region of potentiation has a high peak and a thin tail, and the region of depression has a comparatively small peak and fat tail [22].

Value-dependent synaptic plasticity differs from the value-independent rule in that an additional term, based on the activity and phase of the value system, modulates the synaptic strength changes. Synaptic connections terminating on neuronal units that are in phase with the value system are potentiated, and connections terminating on units out of phase with the value system are depressed.

2.3. Image Processing and Computation

The CCD camera sent 320x240 pixel RGB video images, via an RF transmitter, to a frame grabber attached to one of the workstations running the neural simulation.

The image was spatially averaged to an 80x60 pixel image. Gabor filters were used to detect edges of different orientations (45, 90, 135, 180 degrees). The output of the Gabor function mapped directly onto the neuronal units of the corresponding *V1* sub-area. Color filters (red positive center with a green negative surround, red negative center with a green positive surround) were applied to the image. The outputs of the color filters were mapped directly onto the neuronal units of *V1-Red* and *V1-Green*. *V1* neuronal units projected retinotopically to neuronal units in *V2* (see Figure 2).

Computation in the Darwin VIII simulation was carried out on a Beowulf cluster with 12 1.4 GHz Pentium IV processors using MPI. A simulation cycle, in which all the neuronal units and plastic synaptic connections were updated, took approximately 100 ms.
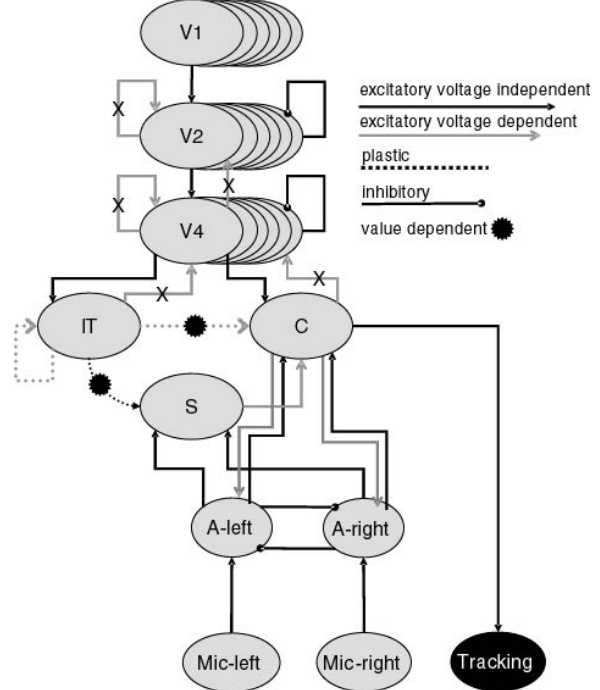
Figure 2. Global schematic of the regional and functional neuroanatomy of Darwin VIII. In the version used in the present experiments, the simulated nervous system contained 28 neuronal areas, 53,450 neuronal units, and approximately 1.7 million synaptic connections. The gray ellipses denote different neural areas. Arrows between the areas denote projections from one area to another. Projections marked with an 'X' are removed during lesion experiments. Tracking commands were issued to NOMAD's wheels based on activity in area *C*.

When the BBD triggers a speaker as it approaches a visual object, the tone emitted by the speaker activates its value system. At this time all of the value-dependent connections between neural areas (see value-dependent

projections in Figure 2) are subject to value-dependent modification. Specifically, the changes dictated by the BCM synaptic change rule are further modulated by the average activity of the value system (area *S* in Figure 2).

As a consequence of these anatomical and dynamical characteristics, Darwin VIII autonomously approaches and views multiple visual shapes containing overlapping features (e.g. red squares, red diamonds, green squares and green diamonds) and can be trained to prefer one of these shapes by associating that shape with a positive-value tone (see Figure 1). It demonstrates this preference by orienting toward the preferred object.

2.4. Experimental Results

When confronted by a pair of these shapes, Darwin VIII learns successfully to track towards the preferred object, designated the target, and to avoid the other objects, designated the distracters. At first, this orientation is in response to the tone, but after approximately 10-15 minutes of viewing pairs of objects, the visual pattern alone is enough to elicit this preference.

All subjects successfully track the four different targets over 80% of the time (Figure 3A). Successful perform-ance on this task is not trivial. Targets and distracters appear in the visual field at many different scales and at many different positions as Darwin VIII explores its envi-ronment. Moreover, because of shared properties, targets cannot be reliably distinguished from distracters on the basis of color or shape alone. Thus, the behavior of Dar-win VIII demonstrates visual categorization and selective conditioning in a rich visual environment.

To investigate the importance of the presence of reen-trant connections in the model, certain interareal reentrant connections were lesioned at different stages of the ex-perimental paradigm In one case, previously trained subjects were retested after lesioning. In a second, reen-trant connections were lesioned in both training and tes t-ing stages. Lesions were applied to a subset of interareal excitatory reentrant connections (projections marked with an 'X' in Figure 2), which had the effect of transforming the simulated nervous system into a 'feed-forward' model of visual proces sing. To compensate for the reduction in activity due to these lesions, neuronal unit outputs in areas *V2* and *V4* were amplified. Figure 3B shows that subjects with intact reentrant connections performed significantly better than either lesioned group. The decrease in per-formance observed in the absence of reentry indicates that reentrant connections are essential for behavior above chance in the discrimination task.

These observations indicate that reentrant connectivity is necessary for the reliable discrimination of targets from visually similar distracters. In contrast to previous mo dels of target selection, which required external intervention or an artificial environment [23, 24], Darwin VIII autono-

mously solved the binding problem in a rich environment even in the face of self-movement that generated changes in the size and location of visual stimuli.
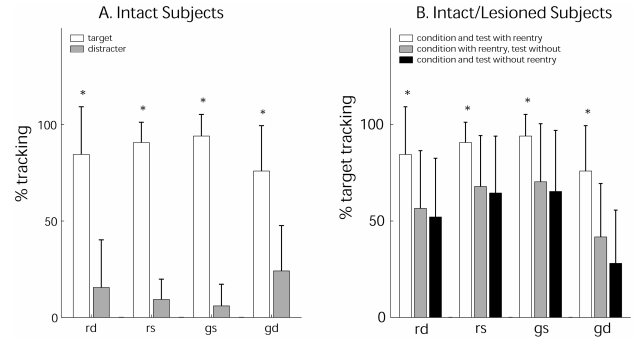
Figure 3. Darwin VIII behavior following condi-tioning. Three separate Darwin VIII subjects were conditioned to prefer one of 4 target shapes ('rd' = red diamond, 'rs' = red square, 'gs' = green square, 'gd' = green diamond). Activity in *V2* areas was used to assess the percentage of time for which NOMAD's visual field was centered on a particular visual shape. Bars in both graphs represent the mean percentage tracking time with error bars denoting the standard deviation. A. Darwin VIII subjects with intact reentrant con-nections tracked the targets (white bars) signifi-cantly more than the distracters (gray bars) for each target shape, averaging over all ap-proaches (asterisks denote p < 0.01 using a paired sample nonparametric sign test). B. Com-parison of intact subjects with lesioned subjects. White bars indicate target tracking performance of subjects with reentrant connections intact, light gray bars indicate subjects with lesions only during testing, and black bars indicate sub-jects with lesions during both training and test-ing. Intact subjects tracked significantly better than both lesion groups (Asterisks denote p < 0.01 using the Wilcoxon ranksum test).

During the behavior of an intact Darwin VIII subject, we observed circuits comprising synchronously active neuronal groups distributed throughout different areas in the simulated nervous system. Multiple objects were dis-tinguishable by the differences in phase between the cor-responding active circuits. A snapshot of Darwin VIII's neural responses is given in Figure 4, in which the device is approaching a red diamond target and a green diamond distracter towards the end of a training session. Each pixel in each neural area represents the activity (brightness) and phase (color) of a single neuronal unit. The figure shows two dynamic neural circuits differentiated by their distinct phases which were elicited respectively by the red dia-mond and the green diamond. As shown in the figure, Darwin VIII had not yet reached the beam that triggers the speaker to emit a tone. The activity of area *S* was

nonetheless in phase with the activity in areas *V2* and *V4* corresponding to the target, and is therefore predictive of the target's saliency or value. Area *IT* has two patterns of activity, indicated by the two different phase colors, which reflect two perceptual categories. The increased activity in area *C* on the side of the target is causing Darwin VIII to orient towards the target (i.e. the red diamond).
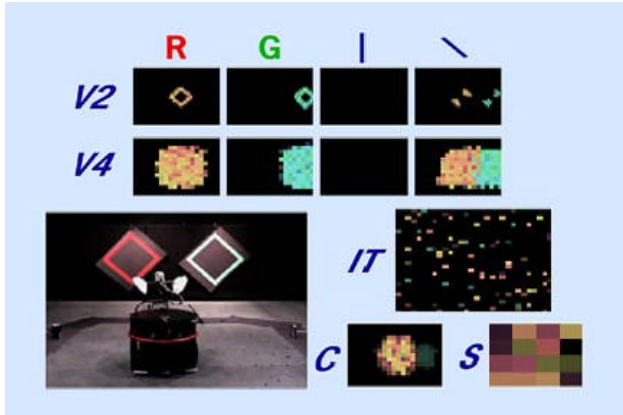
Figure 4. Snapshot of Darwin VIII's neuronal unit activity after approximately 10 minutes of conditioning. Darwin VIII is approaching a red diamond target (left) and a green diamond distracter (right) towards the end of a training session. Darwin VIII has not yet broken the beam that triggers the sound from the speakers located on the left side of the floor. The panels next to Darwin VIII show the activity and phase of selected neural areas (top row; *V2-red, V2-green, V2-vertical, V2-diagonal,* **second row;** *V4-red, V4-green, V4-vertical, V4-diagonal,* **third row (to the right of Darwin VIII);** *IT,* **fourth row (to the right of Darwin VIII);** *C* **and** *S***). Each pixel in the selected neural area represents a neuronal unit; the activity is normalized from no activity (black) to maximum activity (bright colors), and the phase is indicated by the color of the pixel (colors were chosen from a pseudocolor map, there is no connection between the color of the stimulus object and the color representing the phases of neuronal responses). The neuronal units responding to the attributes of the red diamond share a common phase (red-orange color), whereas the neuronal units responding to the green diamond share a different phase (blue-green color).**

Object recognition and perceptual categorization were unsupervised. The simulated nervous system of a given subject developed distinct patterns of activity for each object it observed based on its own experience. Because images of the visual objects varied considerably in size and position as Darwin VIII explored its enclosure, suc-

cessful discrimination required invariant object recognition. Darwin VIII's response to objects was position and scale invariant; it responded reliably to target images which appeared within ±20-degrees of the center field of view (the range of the visual field was approximately ±35-degrees) and as the apparent target size ranged from 8-degrees to 27-degrees of visual angle. This invariance was achieved due to generalization of a continuous stream of input due to self-movement.

The key mechanisms incorporated into Darwin VIII are reentrant connections within and among areas, neuronal units with a mean firing rate and a relative firing phase, and a value system modulating synaptic plasticity. The operation of these mechanisms, in conjunction with the sensorimotor correlations generated by self-motion, enable Darwin VIII to categorize visual objects, bind the features of visual objects, segment a scene, and demonstrate selective behavior in a rich real-world environment.

Our results are consistent with the hypothesis that visual binding results from the dynamic synchronization of neural activity mediated by reentrant connections among many dispersed neural areas. The performance of Darwin VIII suggests that specific timing relations and firing rates can act in a complementary mode to regulate behavior, and that synchrony among groups of neurons, as distinct from synchrony between pairs of individual neurons, may play a significant role in adaptive neural function.

3. Brain-Based Device Playing Soccer

Recently, we applied the visual binding and scene segmentation model of Darwin VIII to a BBD that can play soccer in both indoor and outdoor environments under varying lighting conditions and surfaces. The platform for this device is based on a modification of the Segway balancing technology and allows people on Segway Human Transporters (HT) to interact with Segway RMP robots (http://www.segway.com/segway/rmp/). The rules for this game are currently under development (see [25]) and a new league based on the Segways will be proposed for RoboCup 2005 (http://www.robocup.org). The rules will dictate that the BBD have the ability to both catch and kick a ball and also that that all sensing, actuating, and computing are local to the device (see Figure 5).

The neural simulation constructed for segmenting a soccer scene used the same principles of image processing, visual categorization, reentrant signalling and value-dependent learning as our previous Darwin automata. A high-level schematic of the neural architecture is given in Figure 6.
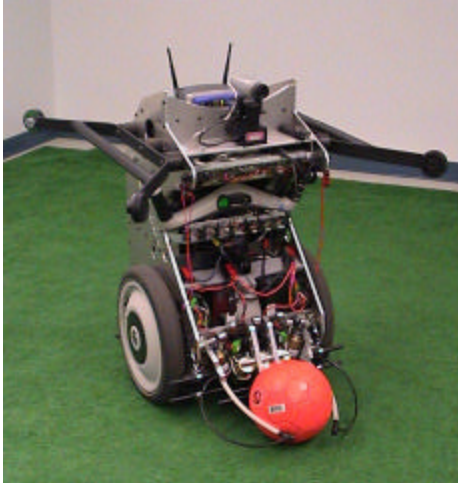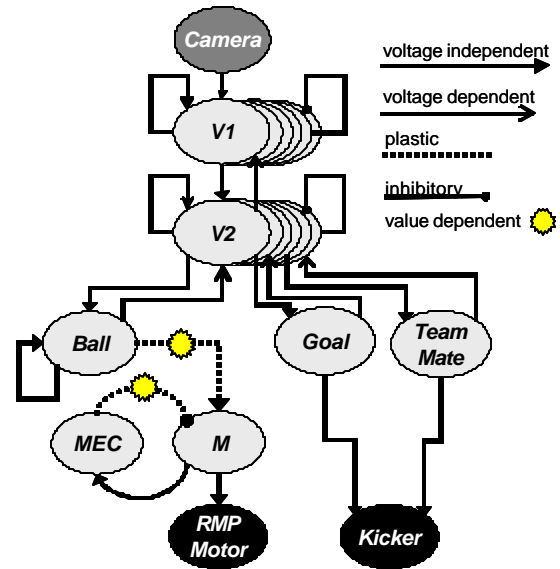
Figure 5. Brain-based soccer playing device based on the Segway RMP balancing platform. The BBD neural simulation receives sensory input from a CCD camera, IR sensors used for ball detection and obstacle avoidance, and odometry from the RMP. The simulation outputs to a camera pan-tilt unit, solenoids to capture the ball, solenoids to kick the ball, and motor commands to the RMP wheels.

Since many of the key elements of the soccer scene were pre-specified by the rules, certain neural areas were dedicated to responding to objects such as the ball, goal post, and teammate. Each of these areas had neuronal units that responded selectively to attributes of the object (e.g. *Goal* neuronal units were active when pre-synaptic neuronal units with similar receptive fields in areas *V2-Green*, *V2-Yellow*, and *V2-Horizontal* neuronal units were active). Activity in these object detection areas triggered motor actions (e.g. activity in Goal triggered kicking the soccer ball).

Because the mapping from visual ball recognition to the Segway RMP wheel motions was non-linear and complex, value-dependent plasticity was used to learn this mapping (see value-dependent learning projections in Figure 6). A form of temporal difference learning was developed in which value was increased when the number of active neuronal units in the *Ball* neural area increased and the activity of *Ball* neuronal units that respond to the center of device's field of view increased. This learning rule had the effects of potentiating *Ball*→*M* and *MEC*→*M* connections when a motor movement brought the ball closer to the device or of depressing connections due to erroneous movements away from the ball (see Figure 6).

Training the BBD to effectively track a soccer ball through value-dependent plasticity took approximately three minutes. In the first twenty seconds, movements were slow and the device's camera did not stay centered on the ball. By the final twenty seconds, ball tracking was fast and the device's camera stayed fixated on the ball.


Figure 6. Schematic of the neural simulation architecture used for the Brain-based soccer playing device. Visual areas and their connectivity are similar to that of Darwin VIII. Neural areas *Ball*, *Goal*, and *TeamMate* respond specifically to those key objects on a soccer field and cause reflexive motor actions such as moving the Segway RMP or capturing and kicking the soccer ball. Ball tracking was achieved via plastic connections from the *Ball* and motor efference copy (*MEC*) areas to the motorneurons (*M*).
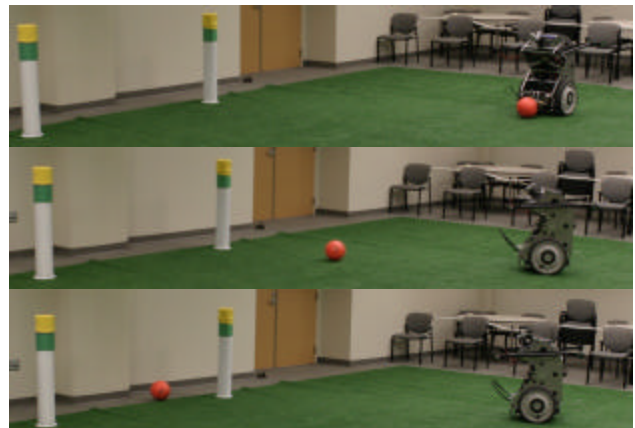
Figure 7. Goal shooting sequence. In the top panel, the BBD recognizes and acquires the soccer ball and then centers its gaze between the two goal posts. In the bottom two panels, the BBD kicks the ball between the goal posts.

Kicking to a teammate or the goal was achieved by the device recognizing the appropriate object, centering the object on its camera, and then kicking the ball (see Figure 7). Reentrant connections between neural areas facilitated dynamic synchronization, and color constancy [26]; al-

lowing the BBD to recognize objects in a noisy environment under non-uniform lighting conditions (see Figure 8).
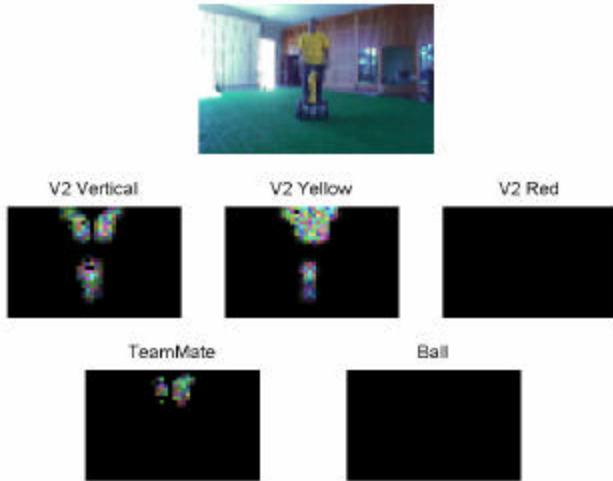
Figure 8. Snapshot of the BBD's camera input and selected areas in the neural simulation just prior to passing to a teammate. The camera input is sub-sampled to 60x80 pixels before processing. The activity and phase representation in the neural areas is the same as depicted in Figure 4.

The soccer playing BBD has successfully performed key elements of soccer playing; ball chasing, passing back and forth between itself and a human teammate on a Segway HT, and goal kicking. Video clips of the devices soccer playing capabilities can be found at: http://www.nsi.edu/nomad/segway.

4. DISCUSSION

Higher brain functions depend on the cooperative activity of an entire nervous system, reflecting its morphology, its dynamics, and its interaction with the phenotype and the environment. BBDs are designed to incorporate these attributes to allow tests of the self-sufficiency of such theories of brain function. We have demonstrated that BBDs can address many difficult tasks, without instruction or intervention, such as invariant object recognition [14], visual binding of objects in a scene [18], and texture discrimination using whiskers [17]. Like the brain, these BBDs operate according to selectional principles through which they form categorical memory, associate categories with innate value, and adapt to novel environments. These devices may provide the groundwork for the development of intelligent machines that follow neurobiological rather than computational principles in their construction.

We designed BBDs to simultaneously test parallel brain functions that could not presently be examined in any single animal in the laboratory. The BBDs were designed to yield data, in the form of neuronal activities in different brain regions that could be directly compared with experimental data. Equally important in the design is that a BBD must demonstrate adaptive behavior and this behavior must be measurable by an observer. The BBD's neural model, by necessity, is developed at a systems level, in which the structure of the brain and its different regions gives rise to adaptive behavior. Although the devices are still too simple to make direct comparisons to neurophysiology, they can make predictions about the neuroanatomical and dynamical constraints that subserve adaptive behavior.

Any model of brain function must not only take into consideration the structure of different brain regions, but must also pay attention to the connectivity within and between these brain areas. Brain function is more than the activity of disparate regions; it is the interaction between these areas that is crucial, as we have shown in Darwin VIII and the soccer playing BBD using the Segway. Thus, brains are defined by a distinct neuroanatomy in which there are areas of special function, which are defined by their connectivity to sensory input, motor output, and to each other.

Brains do not function in isolation; rather are tightly coupled to the organism's morphology, history, and environment. Therefore, our brain models are embodied in a physical device and explore a real as opposed to a simulated environment. The real environment is required for two reasons. First, simulating an environment can introduce unwanted and unintentional biases in a model. For example, a computer generated object presented to a vision model already has its shape and segmentation defined by the modeler and is directly presented to the model, whereas a device that views an object hanging on a wall has to discern the shape and figure from ground by segmentation based on its active vision. Second, real environments are rich, multimodal, and noisy. An artificial design of such an environment would be computationally intensive and difficult to simulate. All these interesting features of the environment come for "free" when we place the BBD in the real world. The modeler is freed from simulating an environment and can concentrate on the development of a device that can actively explore the real world.

The advantage of a synthetic model is that these measurements can be carried out in every neuron and synapse of the BBD's nervous system during the acquisition and recall of a behavior. To be effective, researchers using synthetic models need to analyze their data in such a way that they can compare their results to empirical data. By analyzing the neural dynamics of the model (i.e. spike rates, correlations between areas, neural dynamics and prediction), and choosing a behavioral paradigm similar to those used when studying behaving animals (i.e. mazes, conditioning paradigms, decision-making tasks, etc.), the modeler can directly compare the BBD's behavior with

the results of psychological and neurophysiological experiments. This places the burden on modelers to include sufficient complexity in their models so that these psychological and physiological metrics can be compared.

We have obtained a number of insights and made several predictions based on the results of experiments with BBDs. In Darwin VIII, the model suggests that synchrony between widely separated neural areas may play a key role in solving the binding problem and demonstrates the importance of reentrant connections in facilitating binding through synchrony. The observed behavior demonstrates that binding through synchrony is feasible in an unlabeled real-world environment in which objects are constantly changing in size and position.

The development of adaptive and autonomous behavior by BBDs is novel in its neurally based approach and has implications for the construction of intelligent machines. The design and construction of such behaving devices based on principles of nervous systems may have much to offer to basic understanding and practical applications.

References

[1]   A. Treisman, "The binding problem," *Curr Opin Neurobiol*, vol. 6, pp. 171-8, 1996.

[2]   K. M. Shafritz, J. C. Gore, and R. Marois, "The role of the parietal cortex in feature binding.," *Proc Natl Acad Sci USA*, vol. 99, pp. 10917-10922, 2002.

[3]   M. N. Shadlen and J. A. Movshon, "Synchrony unbound: a critical evaluation of the temporal binding hypothesis.," *Neuron*, vol. 24, pp. 67-77, 1999.

[4]   G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books, Inc., 1987.

[5]   C. M. Gray, "The temporal correlation hypothesis of visual feature integration: still alive and well.," *Neuron*, vol. 24, pp. 31-47, 1999.

[6]   W. Singer, "Neuronal synchrony: a versatile code for the definition of relations?," *Neuron*, vol. 24, pp. 49-65, 1999.

[7]   G. M. Edelman, "Neural Darwinism: selection and reentrant signaling in higher brain function," *Neuron*, vol. 10, pp. 115-25., 1993.

[8]   G. Tononi, O. Sporns, and G. M. Edelman, "Reentry and the problem of integrating multiple cortical areas: simulation of dynamic integration in the visual system," *Cereb Cortex*, vol. 2, pp. 310-35, 1992.

[9]   C. von der Malsburg and J. Buhmann, "Sensory segmentation with coupled neural oscillators," *Biol Cybern*, vol. 67, pp. 233-42., 1992.

[10]   A. Knoblauch and G. Palm, "Scene segmentation by spike synchronization in reciprocally connected visual areas. II. Global assemblies and synchronization on larger space and time scales," *Biol Cybern*, vol. 87, pp. 168-84., 2002.

[11]   A. Knoblauch and G. Palm, "Scene segmentation by spike synchronization in reciprocally connected visual areas. I. Local effects of cortical feedback," *Biol Cybern*, vol. 87, pp. 151-67., 2002.

[12]   G. M. Edelman, G. N. Reeke, W. E. Gall, G. Tononi, D. Williams, and O. Sporns, "Synthetic neural modeling applied to a real-world artifact," *Proc Natl Acad Sci U S A*, vol. 89, pp. 7267-71, 1992.

[13]   N. Almassy, G. M. Edelman, and O. Sporns, "Behavioral constraints in the development of neuronal properties: a cortical model embedded in a real-world device," *Cereb Cortex*, vol. 8, pp. 346-61, 1998.

[14]   J. L. Krichmar and G. M. Edelman, "Machine psychology: autonomous behavior, perceptual categorization and conditioning in a brain-based device," *Cereb Cortex*, vol. 12, pp. 818-30., 2002.

[15]   H. J. Chiel and R. D. Beer, "The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment," *Trends Neurosci*, vol. 20, pp. 553-7, 1997.

[16]   A. Clark, *Being there. Putting brain, body, and world together again.* Cambridge, MA: MIT Press, 1997.

[17]   A. K. Seth, J. L. McKinstry, G. M. Edelman, and J. L. Krichmar, "Spatiotemporal processing of whisker input supports texture discrimination by a brain-based device," in *Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior*, S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J. A. Meyer, Eds. Cambridge, MA: The MIT Press, 2004.

[18]   A. K. Seth, J. L. McKinstry, G. M. Edelman, and J. L. Krichmar, "Visual binding through reentrant connectivity and dynamic synchronization in a brain-based device," *Cereb Cortex*, 2004, In press (http://www.nsi.edu/nomad/aks_jlk_cc2004_preprint.pdf)

[19]   W. Schultz, P. Dayan, and P. R. Montague, "A neural substrate of prediction and reward," *Science*, vol. 275, pp. 1593-9., 1997.

[20]   E. L. Bienenstock, L. N. Cooper, and P. W. Munro, "Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex," *J Neurosci*, vol. 2, pp. 32-48., 1982.

[21]   W. Senn, H. Markram, and M. Tsodyks, "An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing," *Neural Comput*, vol. 13, pp. 35-67, 2001.

[22]   E. M. Izhikevich and N. S. Desai, "Relating STDP to BCM," *Neural Comput*, vol. 15, pp. 1511-23, 2003.

[23]   L. Itti and C. Koch, "A saliency-based search mechanism for overt and covert shifts in attention.," *Vision Res*, vol. 40, 2000.

[24]   B. A. Olshausen, C. H. Anderson, and D. C. Van Essen, "A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information.," *Journal of Neuroscience*, vol. 13, pp. 4700-4719, 1993.

[25]   B. Browning, P. E. Rybski, J. Searock, and M. M. Veloso, "Development of a Soccer-Playing Dynamically-Balancing Mobile Robot," presented at IEEE International Conference on Robotics & Automation, New Orleans, 2004.

[26]   J. Wray and G. M. Edelman, "A model of color vision based on cortical reentry," *Cereb Cortex*, vol. 6, pp. 701-716, 1996.