Machine Psychology: Autonomous Behavior, Perceptual Categorization and Conditioning in a Brain-based Device

Jeffrey L. Krichmar and Gerald M. Edelman

W.M. Keck Laboratory, The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, USA

In studying brain activity during the behavior of living animals, it is not possible simultaneously to analyze all levels of control from molecular events to motor responses. To provide insights into how levels of control interact, we have carried out synthetic neural modeling using a brain-based real-world device. We describe here the design and performance of such a device, designated Darwin VII, which is guided by computer-simulated analogues of cortical and subcortical structures. All levels of Darwin VII's neural architecture can be examined simultaneously as the device behaves in a real environment. Analysis of its neural activity during perceptual categorization and conditioned behavior suggests neural mechanisms for invariant object recognition, experience-dependent perceptual categorization, first-order and second-order conditioning, and the effects of different learning rates on responses to appetitive and aversive events. While individual Darwin VII exemplars developed similar categorical responses that depended on exploration of the environment and sensorimotor adaptation, each showed highly individual patterns of changes in synaptic strengths. By allowing exhaustive analysis and manipulation of neuroanatomy and largescale neural dynamics, such brain-based devices provide valuable heuristics for understanding cortical interactions. These devices also provide the groundwork for the development of intelligent machines that follow neurobiological rather than computational principles in their construction.

Introduction

A central goal of research in the neurosciences is to understand the relationships between brain structure, function and behavior. Several related factors make this a challenging task. One is the sheer complexity of neuroanatomical networks overlain by the physiological subtleties of brain dynamics. Another is the number of levels of control ranging from molecular events to perception, memory and the coordination of movement. Each level requires the analysis of a number of simultaneous causal factors and chains acting in parallel. The environment itself and interactions between an organism and its econiche add further complexity.

In dealing with this degree of complexity, careful experimental analysis and theory building are obviously essential. However, analytic approaches conducted separately at each level are unlikely alone to provide a full picture of neural patterns in a behaving organism. There are obvious limits on the number of levels simultaneously observable during any given experiment. Moreover, despite the power of mathematical and computational approaches, they have not yet provided a multilevel picture of the non-linear relationships between brain and behavioral events.

To confront these issues and complement these approaches, we have adopted a procedure called synthetic neural modeling (Reeke *et al.*, 1990; Edelman *et al.*, 1992). This consists of building devices capable of behavior, providing them with a computationally simulated nervous system based on known

biological principles of neuroanatomical organization and physiological activity, and then following the behavioral and neuronal responses of such a construct in real time, in a real-world environment. By following behavioral and brain responses completely at all levels of control in a particular environment, one can formulate a synthetic picture that has heuristic value in interpreting data obtained from behaving animals.

A series of such brain-based devices capable of increasingly sophisticated autonomous performance has been tested over the last decade (Edelman *et al.*, 1992; Almassy *et al.*, 1998; Krichmar *et al.*, 2000; Sporns *et al.*, 2000). In these earlier devices, we demonstrated the learning of perceptual responses and emphasized the role of value systems. Value systems are neural structures that are necessary for an organism to modify its behavior based on the salience or value of an environmental cue (Friston *et al.*, 1994). The value system in a brain-based device is analogous to ascending neuromodulatory systems in that its units show uniform phasic responses when activated and its output acts diffusely over multiple pathways by modulating synaptic change (Schultz *et al.*, 1997; Sporns *et al.*, 2000).

In the present report, we describe the construction and performance of Darwin VII, a device capable of perceptual categorization and conditioned behavior. We have extended previous conditioning experiments to include second-order conditioning and have carried out an extensive analysis of the responses of simulated neuronal units. By probing simultaneous brain and behavioral responses at all levels during perceptual and conditioning tasks, we have obtained several new insights into the organization of autonomous behavior. These include a richer understanding of the effects of individual history on learning, of the possible origins of invariant object recognition in an analogue of the inferotemporal cortex, and of the rela-tion of changes in synaptic efficacy to appetitive and aversive conditioned responses.

Materials and Methods

We have developed a heuristic in which a neurally organized mobile adaptive device (NOMAD) explores its environment and through experience-dependent learning develops adaptive behaviors. NOMAD is a part of the Darwin series of automata in which theories of the nervous system are tested by implementing brain-based devices (Reeke *et al.*, 1990; Edelman *et al.*, 1992; Almassy *et al.*, 1998; Krichmar *et al.*, 2000). In the present report, we will refer to the device and the neural simulation together as Darwin VII. The NOMAD portion of Darwin VII consists of a mobile base equipped with a CCD camera for vision (see Appendix, part E), microphones for hearing (see Appendix, part F), conductivity sensors for taste, and effectors for movement of its base, of its head, and of a gripping manipulator having one degree of freedom (Fig. 1).

Darwin VII's behavior was guided by a nervous system simulated on a computer workstation (see Appendix, part A). The simulation was based on the anatomy and physiology of vertebrate nervous systems but obviously with fewer neurons and simpler architecture. The simulated



Figure 1. Darwin VII consists of a mobile base equipped with several sensors and effectors (NOMAD), and a neural simulation running on a remote computer workstation. NOMAD contains a radio modem to transmit status and auditory information to the computer workstation carrying out the neural simulation and to receive motor commands from the simulation. Video output from a CCD camera mounted on Darwin VII is sent to the workstation via RF transmission. RF input and output to NOMAD allows for untethered exploration. NOMAD is constructed on a circular platform (developed by Nomadic Technologies Inc., Mountain View, CA, USA) with wheels that permit independent translational and rotational motion, with pan and tilt movement for its camera and microphones, and with object gripping by a one degree of freedom manipulator or gripper. The CCD camera, two microphones on either side of the camera, and sensors embedded in the gripper that measure the surface conductivity of stimuli provide sensory input to the neuronal simulation. Eight infrared (IR) sensors are mounted at 45° intervals around the mobile platform. The IR sensors are responsive to the boundaries of the environment and were used to trigger reflexes for obstacle avoidance. All behavioral activity other than obstacle avoidance is triggered by signals received from the neural simulation.

nervous system was made up of a number of areas labeled according to the analogous cortical and subcortical brain regions. Each area contains different types of neuronal units consisting of simulated local populations of neurons or neuronal groups (Edelman, 1987). To distinguish modeled areas from real counterparts, the simulated areas are indicated in italics (e.g. *IT*). A neuronal unit in Darwin VII is simulated with a mean firing rate model and the activity of such a unit corresponds roughly to the firing activity of a group of neurons averaged over a time period of 200 ms.

In the present experiments, the simulated nervous system contained 18 neuronal areas, 19 556 neuronal units, and ~450 000 synaptic connections. Figure 2 shows a high-level diagram of the different neural areas and the synaptic connections between neural areas in the simulated nervous system. Further details of the parameters describing neuronal unit activity and neuronal unit connectivity can be found in the Appendix (see Tables A1 and A2, and Appendix, parts B and C). Each simulation cycle took ~200 ms of real time. A simulation cycle is the period during which the current sensory input is processed, the activities of all neuronal units are

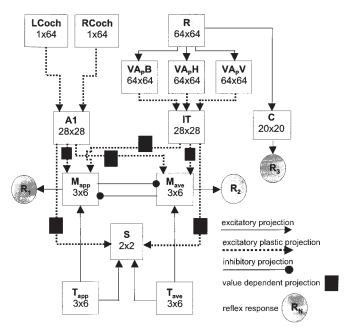


Figure 2. Schematic of the regional and functional neuroanatomy of Darwin VII. There are six major systems that make up the simulated nervous system: an auditory system, a visual system, a taste system, sets of motor neurons capable of triggering behavior, a visual tracking system, and a value system. In the version used in the present experiments, the simulated nervous system contained 18 neuronal areas, 19 556 neuronal units and ~450 000 synaptic connections. A neuronal unit corresponded to the mean activity of a small group of real neurons over ~200 ms. The 64 imes 64 gray level pixel image captured by the CCD camera was relayed to a retinal area R and transmitted via topographic connections to a primary visual area VAP. There were three subpartitions in VAP selective for blob-like features, for short horizontal line segments, or short vertical line segments. Responses within VAP closely followed stimulus onset and projected non-topographically via activity-dependent plastic connections to a secondary visual area, analogous to the inferotemporal cortex (IT). The frequency and amplitude information captured by Darwin VII's microphones was relayed to a simulated cochlear area (LCoch and RCoch) and transmitted via mapped tonotopic and activity-dependent plastic connections to a primary auditory area A1. The activity of each cochlear neuronal unit was broadly tuned to a preferred frequency and scaled according to the signal amplitude (see Appendix, F. Auditory System and its Inputs). A1 and IT contained local excitatory and inhibitory interactions producing firing patterns that were characterized by focal regions of excitation surrounded by inhibition. A1 and IT sent plastic projections to the value system S and to the motor areas M_{app} and M_{ave} . These two neuronal areas were capable of triggering two distinct behaviors, appetitive and aversive. Appetitive or aversive responses were triggered if the difference in instantaneous activity between motor areas M_{app} and M_{ave} exceeded a behavioral threshold ($\beta=0.3$). The taste system (T_{app} and T_{ave}) consisted of two kinds of sensory units responsive to either the presence or absence of conductivity across the surface of stimulus objects as measured by sensors in Darwin VII's gripper. Picking up and sampling the conductivity of a block is innate to Darwin VII's behavior. In all the experiments, strongly conductive blocks activated T_{app} and weakly conductive blocks activated T_{ave} . The taste system sent information to the motor areas (M_{app} and M_{ave}) and the value system (S). Area S projects diffusely with long-lasting value-dependent activity to the auditory, visual and motor behavior neurons. The visual tracking system controlled navigational movements, in particular the approach to objects identified by brightness contrast with respect to the background. To achieve tracking behavior, the retinal area R projected to area C('colliculus') with connection strengths assigned based on learning experiments in a previous study (Edelman et al., 1992). Neural areas $A1_i$, IT_i , S_o and S_i , and their corresponding synaptic projections, are omitted for clarity (see Tables A1 and 2 for complete details).

computed, the connection strengths of all plastic connections are computed, and motor output is generated (see Appendix, parts A and B). Connections between and within neuronal areas were subject to activity-dependent modification following a value-independent (see Appendix, part C) or value-dependent (see Appendix, part D) synaptic rule. Synaptic modification was determined by both pre- and post-synaptic activity and resulted in either strengthening or weakening of the

synaptic efficacy between two neuronal units. We used a modified Bienenstock, Cooper and Munro (BCM) learning rule to govern synaptic change because it has a region in which weakly correlated inputs are depressed and strongly correlated inputs are potentiated (Bienenstock et al., 1982). Simplifying the BCM learning rule by making it piecewise linear and fixing the thresholds, resulted in an efficient biologically based learning rule (see Appendix, part C). Plastic connections that were value-dependent were made between areas involved in responses to salient environmental events $[A1/IT \rightarrow M_{app}/M_{ave}, A1/IT \rightarrow S;$ see (Aston-Jones and Bloom, 1981; Ljungberg et al., 1992)]. Plastic connections that were not value-dependent were made between areas where perceptual categories were to be learned from experience $(VA_P \rightarrow IT, LCoch/RCoch \rightarrow A1)$. Non-plastic connections were between neural areas where there were reflex responses $(T_{app}/T_{ave} \rightarrow M_{app}/M_{ave},$ $R \rightarrow C$), local projections within an area ($IT \rightarrow IT$, $A1 \rightarrow A1$), or between areas where it was assumed that plasticity had already occurred very early in development $[R \rightarrow VA_P$, see Crair et al. (Crair et al., 1998)]. On the assumption that these synaptic changes do not saturate or persist indefinitely, we used a passive synaptic decay term (see ϵ in Appendix, part C) to express a decline in synaptic strength in the absence of activity. Activation of the simulated value system (area S. Fig. 2) signaled the occurrence of salient sensory events and contributed to the modulation of connection strengths of all active synapses in the affected pathways (see value-dependent projections in Fig. 2). For example, 'tasting' a block picked up by Darwin VII's gripper is a salient event affecting subsequent behavior that is reinforced or weakened through synaptic change. Area S is thus analogous to an ascending neuromodulatory value system (Schultz et al., 1997; Sporns et al., 2000).

Experimental evidence suggests that key parameters of neural plasticity may vary over the course of postnatal development (Kato et al., 1991; Fox, 1995; Kirkwood et al., 1995). Since our model was intended to reflect both early and late stages of development, we incorporated such changes in postnatal cortical plasticity by modeling a transition from an earlier critical period to a later adult period. Inasmuch as activitydependent plasticity tends to destabilize the activity of neuronal networks in the absence of homeostatic mechanisms (Turrigiano and Nelson, 2000), connections from VAP to IT, cochlear neuronal units (LCoch and RCoch) to A1, and IT/A1 to M_{app}/M_{ave} (see Fig. 2) underwent several stages of plasticity, characterized by progressive changes in learning rate and synaptic decay as a function of activity in a particular modality (see Table A2 in the Appendix). Without these homeostatic mechanisms, unbounded long-term potentiation could cause network activity to increase and lose its discrimination with the consequence that neuronal units would respond to any input. From earlier to later stages of development, these changes are qualitatively as follows: synaptic plasticity decreases in magnitude; synaptic weights become more resistant to decay back to their original levels; and thresholds for synaptic potentiation

In the experiments in which individual variation was to be examined, each Darwin VII 'subject' shared the same physical device, but had an instantiation in which the simulated nervous system was unique, as a result of different random initializations within the constraints given by Table A2, in both the connectivity between individual neuronal units and the initial connection strengths between those units. Because the connectivity between neuronal units was constrained by a common set of projections, however, large-scale connectivity (i.e. projections between neural areas) was similar between subjects. Details of the neuroanatomical constraint parameters for each synaptic projection, as well as parameters for the synaptic efficacy rules and the projection distributions, can be found in the Appendix (part C and Tables A1 and A2).

Darwin VII's environment consisted of an enclosed area with black walls and a floor covered with opaque black plastic panels, on which we distributed stimulus blocks (6 cm metallic cubes) in various arrangements (Fig. 1). The top surfaces of the blocks were covered with removable black and white patterns; the other surfaces of the blocks were featureless and black. All experiments reported in this paper were carried out with multiple exemplars of two basic designs: blobs (several white patches 2-3 cm in diameter) and stripes (width 0.6 cm, evenly spaced). Stripes on blocks in the gripper can be viewed in either horizontal or vertical orientations, yielding a total of three stimulus classes of visual patterns to be discriminated (blob, horizontal and vertical). A flashlight mounted on Darwin VII and aligned with its gripper caused the blocks, which contained a photodetector, to emit a beeping tone when Darwin VII was in the vicinity. The sides of the stimulus blocks were metallic and could be rendered either strongly conductive ('good taste' or appetitive) or weakly conductive ('bad taste' or aversive). Gripping of stimulus blocks activated the appropriate taste neuronal units (either area T_{app} or area T_{ave}) to a level sufficient to drive the motor areas above a behavioral threshold. In the experiments described in this paper, strongly conductive blocks with a striped pattern and a 3.9 kHz tone were chosen arbitrarily to be positive value exemplars, whereas weakly conductive blocks with a blob pattern and a 3.3 kHz tone represented negative value exemplars.

Basic modes of behavior built into Darwin VII included IR sensordependent obstacle avoidance, visual exploration, visual approach and tracking, gripping and 'tasting', and two main classes of innate behavioral reflex responses (appetitive and aversive). With the exception of obstacle avoidance, selection among the above behaviors was under control of the simulated nervous system. Appetitive and aversive responses were triggered when the difference in activity between motor areas M_{app} and M_{ave} exceeded a threshold (Fig. 2). These responses could be activated by taste (the unconditioned stimulus, US, triggering an unconditioned response, UR) or by auditory or visual stimuli (the conditioned stimulus, CS, triggering a conditioned response, CR). Prior to conditioning, taste triggered the behavioral responses; after conditioning, either a visual pattern or an auditory pattern could evoke behavioral responses. Unconditioned appetitive and aversive behavioral responses consisted of prolonged gripping and 'tasting' of a stimulus block, releasing the block, and then turning counterclockwise. Conditioned appetitive responses, which occurred when motor area activity exceeded the threshold before tasting, differed from unconditioned appetitive responses in that a clockwise turn was executed after 'tasting' a block. In conditioned aversive responses, Darwin VII avoided a stimulus block by backing away without picking it up and then turning clockwise. Thus, during the conditioning experiments, in which many stimuli were encountered over an extended period of time, Darwin VII developed perceptual categories that modified its behavioral responses.

Results

We describe details of two sets of experiments that demonstrate the usefulness of synthetic brain-based devices in testing theories of the nervous system and in understanding how interactions of the nervous system, the body, and the environment affect behavior. The first set focused on visual perceptual categorization and invariance in cortical responses; the second investigated conditioning experiments involving multiple sensory modalities.

Perceptual Categorization

Perceptual categorization is the ability to discriminate and categorize sensory stimuli (Clark et al., 1988; Kilgard and Merzenich, 1998). Development of this ability is obviously necessary for learning and conditioning and, for this reason, was extensively explored in Darwin VII. In primates, the inferotemporal cortex is an area that is believed to be associated with visual object recognition (Tanaka, 1996). In Darwin VII, activity in the simulated inferotemporal cortex, area IT (Fig. 2), provided the basis for visual perceptual categorization. Initially, IT's responses to visual stimuli were weak and diffuse (see IT activity in Fig. 3A). After approximately five stimulus encounters, activitydependent plasticity between VAP and IT caused IT responses to the different stimuli to become strong, sharp and separable (see IT activity in Fig. 3B). It is this strong, discriminative activity of neuronal groups within IT in response to visual stimuli as well as the appropriate behavioral response that we refer to as visual categorization in Darwin VII.

Invariant Object Recognition

In animals, perceptual categorization in the inferotemporal

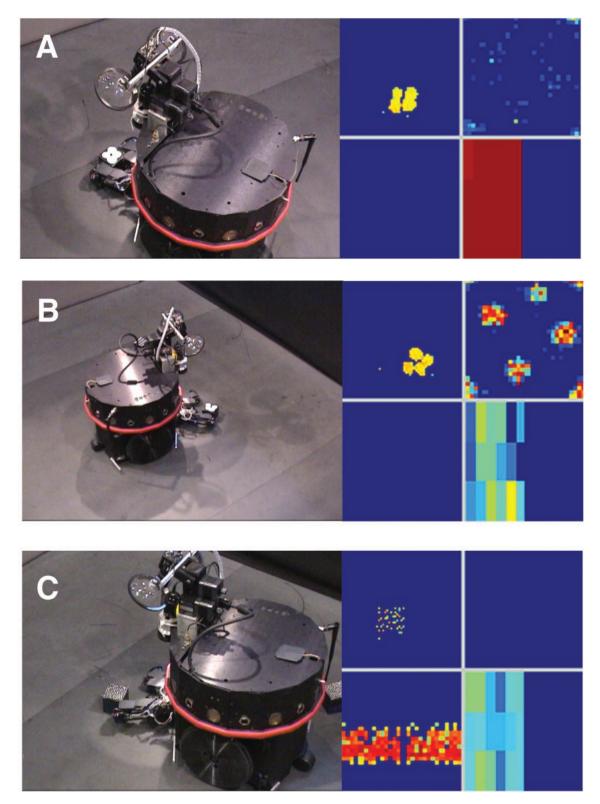


Figure 3. Darwin VII during behavioral experiments. The panels to the right of Darwin VII show activity of selected neural areas in the simulation (R, top left; IT, top right; A1, bottom left; M_{ave} , bottom right, left side; M_{app} , bottom right side in the side side in t

cortex is invariant with respect to differences in position, scale and rotation of an object (Tanaka et al., 1991; Tovee et al., 1994; Ito et al., 1995; Rolls and Tovee, 1995; Tanaka, 1996). Such invariant object recognition has been difficult to achieve in computer vision systems (Mundy and Zisserman, 1992; Mundy et al., 1992; Shashua, 1993; Weinshall, 1993). In the present work, however, Darwin VII's object recognition was observed to be invariant with respect to scale, position and rotation. Visual categorization of a stimulus occurred no matter where an object appeared in Darwin VII's visual field, with the apparent size of the stimulus ranging from a maximum when the object was directly in front of Darwin VII (Fig. 3A) to one-quarter of the maximum size when the object was distal to Darwin VII. Correct categorization of striped blocks in Darwin VII's field of vision, when blocks were not in its gripper, occurred when the stripes on the blocks were rotated over a range of ±30° of a horizontal or vertical reference.

Invariant object recognition required continuous, timevarying sensory input while Darwin VII moved. Invariant responses developed as a result of competition among activitydependent plastic connections between retinotopically mapped VA_P and non-topographically mapped IT. The connections that were potentiated earliest were those with VA_P receptive fields corresponding to regions near Darwin VII's gripper, regions where IT responses to the neural stimulus were first sustained (Fig. 4, top 'First Horizontal Striped Block'). These connections had a competitive advantage; they received not only the earliest but also the longest exposure to the stimulus as a result of the time spent by the block in the gripper. The maintenance of discriminative, persistent patterns of neuronal groups in IT required sustained high activity resulting from strengthening of plastic connections with VA_p neuronal units that received continuously varying images of the block. Upon each approach and withdrawal from the stimulus block, the number of potentiated connections increased, resulting in recruitment of neuronal units with receptive fields that responded to visual stimuli beyond Darwin VII's gripper (Fig. 4, top). An example of the resultant activity in VAp and IT during invariant object recog-

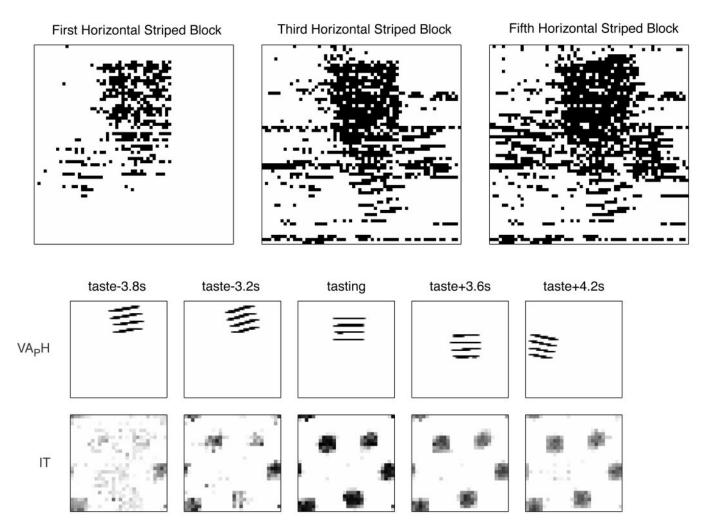


Figure 4. Invariance with respect to position, scale and rotation emerges from a persistent pattern of activity in area IT as the pattern of activity in the IVAP areas moves sequentially across Darwin VII's field of vision. Shown in black (see top row) are the locations of VAPH neuronal units responding to a horizontal visual pattern, whose synaptic weights, going from WAPH to area IT, increased from their initial value. The WAPH neuronal units are pictorially organized such that receptive fields near NOMAD's base are at the top of the figure and receptive fields that are further away from NOMAD's base are at the bottom of the figure. After presentation of the first horizontal block, most of the potentiation of synaptic weights occurred in VAPH neuronal units with receptive fields near the position of Darwin VII's gripper. After each subsequent stimulus presentation with continuous movement, the potentiated receptive fields build up throughout Darwin VII's field of vision. The proportion of potentiated VAPH neuronal units increased from 10% after the first stimulus presentation to 33% after the fifth stimulus presentation. The bottom two rows show sequential activity in VAPH and IT during a stimulus encounter with one striped block. Activity in neuronal area VA_PH is shown in the middle row and that in neuronal area IT is shown in the bottom row as Darwin VII approaches the stimulus, grasps and tastes, and moves away.

nition is shown in the bottom two rows of Figure 4. When the temporal sequence of the images leading to invariance was artificially shuffled (Almassy *et al.*, 1998), invariant object recognition did not occur. As further considered in the Discussion, the invariance arose mainly as a result of an initial strengthening of VA_P to IT synapses that was reinforced and expanded by subsequent inputs from the stimulus block during Darwin VII's movements.

Stimulus History and Individual Variation

Differences in an individual's perceptual history can have an effect on the organization and response of the nervous system. For example, more neurons in the monkey inferotemporal cortex respond to familiar than to unfamiliar stimuli (Kobatake *et al.*, 1998). Using Darwin VII, we performed experiments concerned with experience-dependent effects on categorization during the development of perceptual categories as well as after such development.

We first investigated the effect of variations in presentation frequency of each stimulus class on the development of neuronal unit responses. Darwin VII explored an environment that was partially segregated into two equal sized areas. One area mainly contained blocks with blobs and the other area mainly contained blocks with stripes. In each of 14 separate experiments, Darwin VII started with an identical simulated nervous system that had not sampled any stimuli. The number of neuronal units in IT responding to a given stimulus (whether blob, horizontal stripe or vertical stripe) increased selectively with an increase in the frequency of presentation of that stimulus class. Statistical significance was tested using, r, Pearson's product moment correlation. Stimulus presentation frequency was found to be positively correlated with patterned neural activity in IT that was individually characteristic for each of the visual stimulus classes (blob: r = 0.71, P < 0.005; horizontal: r = 0.75, P < 0.003; vertical: r = 0.61, P < 0.03). These findings are similar to the results of neuronal recordings in the monkey inferotemporal cortex in that more IT neurons responded to familiar than to unfamiliar objects in recognition tasks (Kobatake et al., 1998).

In these experience-dependent responses, competitive and selective interactions among neuronal units from VA_P to IT and within IT governed the changes in the number of those units that responded to a stimulus. An increase in neuronal group size reflected the activity-dependent changes in synaptic connections from neuronal units in VA_P to neuronal units in IT, leading to increased activity in IT. Through intrinsic excitatory connections, this increased activity further recruited neighboring neuronal units in IT. The change in neuronal group size was competitive: a group specific to one stimulus could grow at the expense of another neuronal group associated with another stimulus (Clark et al., 1988).

In the second set of experiments on experience-dependent perceptual categorization, we studied the effect of stimulus presentation frequency on neural mapping in *IT after* visual categories had already been developed. To reach this level of experience, Darwin VII sampled an equal proportion (eight each) of blocks in the three stimulus classes. Darwin VII then sampled either eight additional stimuli containing all three stimulus classes or eight additional stimuli containing any two out of the three stimulus classes. Thus, some stimuli were more frequently sampled than others.

In contrast to the previous experiments on early development, after extensive experience, the number of neuronal units in *IT* responding to more frequently sampled stimuli did not change significantly, suggesting that responses in *IT* had become

Table 1Role of history in perceptual categorization

Test stimuli Effect of stimulus presentation frequency on IT activity after development of visual categorization

	Control	Stripes only	Horizontal and blob Vertical and blob		
Blob	88.8	72.9*	89.9	91.6	
Horizontal	54	55.6	57.7	31.5*	
Vertical	24.2	24.4	17.9*	23.8	

After visual categories had been developed (eight presentations of each stimulus class), the activity in IT stayed relatively constant for stimuli presented in higher proportions, but decreased for stimuli presented in lower proportions. Each row in the table shows the median number of neuronal units in IT responding to stimuli for a control group (eight additional presentations of blob, horizontal and vertical stimuli) and three experimental groups in which eight additional stimuli from two out of the three classes were presented. There were 10 trials for each group with identical initial conditions in the simulated nervous system. The asterisks denote a significant difference (P < 0.05) in medians between the control group and an experimental group using the Wilcoxon Rank Sum test of medians.

'saturated' with respect to the familiar stimuli. However, in the experiments in which Darwin VII responded to the less frequently sampled stimulus, the number of IT neuronal units was significantly less than that in the controls (Table 1). Two factors appear to be responsible for these results. First, the growth of the neuronal groups in IT was limited by intrinsic excitatory and inhibitory connections. Recurrent excitation caused the size of the groups to grow, but lateral inhibition kept that growth in check (see IT activity in Fig. 3B). At a certain size, the different neuronal groups that were active in response to a visual stimulus competed with each other and their growth was halted. In essence, the memory for these perceptual categories is stable. Secondly, the decrease in neuronal units responding to an under-sampled stimulus was governed, in part, by the decay rate, ε, in the activity-dependent synaptic efficacy rule (see Appendix, part C). This caused the efficacy of each synaptic connection that had not been recently updated to decay towards its original value. If, for example, the blob visual stimulus was not encountered for a protracted period of time, synaptic connections from VApB to IT weakened and fewer IT neuronal units responded to that stimulus class. In essence, the perceptual category was forgotten.

In addition to the influences of environmental experience on perceptual categorization, there were noteworthy individual variations in neuronal response patterns related to behavioral differences. Seven Darwin VII subjects, each with nervous systems having different initial conditions in connectivity and connection strengths, were allowed to sample at least 10 aversive and 10 appetitive blocks. The *IT* activity patterns showed significant variations between subjects and between stimulus classes within the same subject (Fig. 5). For an individual subject, many neuronal units were active for more than one stimulus class, but the overall pattern of activity of a given subject's response to each stimulus class was distinguishable. Despite the notable individual differences in patterns of neural activity, all subjects were able to categorize the three stimulus classes as shown by similar behavioral responses.

Response Sampling

In contrast to the limited number of cells whose activity can be monitored in live animals, the design of Darwin VII allowed us to record all such activity in all neuronal units. Neurophysiologists often test whether limited samples from brain areas are robust predictors of responses to input stimuli (Bialek and Zee, 1990; Theunissen and Miller, 1991; Brown *et al.*, 1998). It was therefore of interest to investigate whether a sparse sampling of

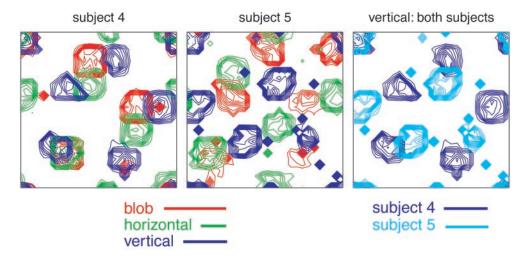


Figure 5. Comparison of patterns of activity in area IT for the three visual stimuli across different Darwin VII subjects. The first two contour plots on the left show, for two representative subjects, the borders of neuronal group activity in response to blobs (red lines), horizontal stripes (green lines) and vertical stripes (blue lines). The contours on the far right show the activity for two different subjects in the same plot, in response to vertical stimuli for subject 4 (dark blue) and for subject 5 (light blue). Across all subjects tested (n=7), the mean overlap of activity within each subject but between stimuli was 26.6% ($\sigma=0.10$). The mean overlap of activity between different subjects was 22.1% ($\sigma=0.09$) in response to blob patterns, 26.9% ($\sigma = 0.11$) in response to horizontal striped patterns, and 20.0% ($\sigma = 0.10$) in response to vertical striped patterns.

neural patterns in the simulated area IT would reliably predict the response to visual stimuli by Darwin VII. We allowed seven individually different Darwin VII subjects to sample at least 20 aversive and 20 appetitive blocks. For each Darwin VII subject, patterns of activity in IT during the development of visual categories (i.e. exposure to the first 10 aversive and appetitive exemplars) were compared with patterns of activity in IT after categorization (i.e. exposure to the last 10 aversive and appetitive exemplars (see Appendix, part G). The accuracy of classification based on IT activity improved with each stimulus exemplar to near perfect performance (Fig. 6). Classifications remained accurate even when relatively small sub-populations (1% of the neuronal units) in IT were sampled; below this range, prediction failed. The relatively small proportion of neuronal units in IT sufficient to classify responses to a given stimulus is in accord with results in live animals, as seen for example, in the limited number of hippocampal neurons needed to reconstruct a rat's position in space (Wilson and McNaughton, 1993) or the limited number of motor cortical neurons needed to predict a monkey's hand position (Georgopoulos et al., 1986).

Conditioning Experiments

In a series of conditioning experiments, Darwin VII was trained to associate the taste value of objects with their visual and auditory characteristics. Weakly conductive objects were assigned innate negative value ('bad taste') and strongly conductive objects were assigned innate positive value ('good taste'). In accord with our prior and arbitrary assignments of block properties, Darwin VII, through experience-dependent learning, associated the blob visual pattern and 3.3 kHz beeping tone with negative value, and the striped visual pattern and 3.9 kHz beeping tone with positive value. Seven individually different Darwin VII subjects participated in the experiments, in which each subject encountered at least 10 appetitive and 10 aversive blocks. In experiments in which only visual stimuli were paired with 'taste', positive conditioned responses occurred in >70% of trials after encountering the sixth exemplar and in >90% after encountering the tenth exemplar. In auditory conditioning trials, conditioned responses occurred in over 80% of trials following exposure to the sixth exemplar. While performance improved

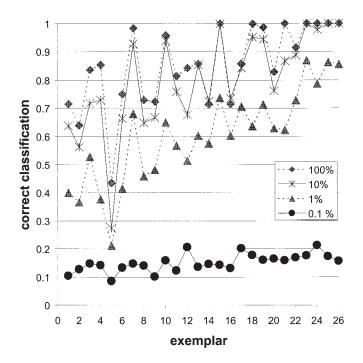


Figure 6. Classification of responses to stimuli by sampling IT activity. Classification of the input stimulus during training is based on IT activity after training was complete. The different data series represent the percentage of IT neuronal units used to classify the input stimuli. Each point on the chart represents the average of seven Darwin VII subjects where the accuracy of each individual subject is based on an average of 10 random sub-samplings of neuronal units in IT (see Appendix, G. Sampling of IT activity for Classification of Responses). The accuracy of classification increases with the numbers of exemplar encounters yielding near total accuracy by the eighteenth encounter when all 784 neuronal units in IT are sampled and also when 10% of the neuronal units in IT are sampled. Accuracy drops to chance (25% or lower) when only 0.1% of the neuronal units in IT is sampled.

with training, it never reached perfection and occasional 'mistakes' were made. This unpredictability is a property of selectionist systems in general. These are systems consisting of a population of variant repertoires which can be differentially

amplified, thus yielding responses to unpredicted or novel events. Such selection has been proposed as being a property of real nervous systems (Edelman, 1987). The unpredictability of behavioral responses in Darwin VII coupled with the variability of a complex environment did not, however, prevent the device from learning after mistakes, from generalizing over sensory inputs, and even from dealing with novel situations.

Early during the conditioning trials, Darwin VII picked up and 'tasted' blocks that led to either appetitive or aversive responses (see Fig. 3A). During this period, it was the output of the taste neuronal units that activated the value system (S) and drove the motor neuronal units (M_{app} and M_{ave}) to cause a behavioral response. After conditioning, however, both the value system and the motor neuronal units were immediately activated upon the onset of IT's response to a visual pattern or AI's response to a tone. This shift following learning, from value system activity that was triggered in early trials by the unconditioned stimulus to value system activity triggered at the onset of the conditioned stimulus, is analogous to the shift in dopaminergic neuronal activity found in the primate ventral tegmental area after conditioning (Schultz et al., 1997).

After associating visual patterns with taste, Darwin VII continued to pick up and 'taste' stripe-patterned blocks, but avoided blob-patterned blocks (see Fig. 3*B*). After associating auditory sounds with taste, Darwin VII continued to pick up the high frequency beeping blocks, but avoided the low frequency beeping blocks (see Fig. 3*C*).

We extended the training paradigm by carrying out secondorder conditioning experiments (Rescorla, 1980). In the first stage of conditioning, a single conditioned stimulus (CS₁; either the tone or the visual pattern) was paired with taste for ~10 encounters with each block type until learning was achieved. In the second stage of conditioning, the two conditioned stimuli (CS₁ and CS₂, tone and visual pattern) were paired together for ~10 encounters with each block type. After the second stage, Darwin VII's performance was tested by presenting CS₂ alone for 10 encounters of each block type. There were four possible behavioral responses for each stimulus encounter: (i) appetitive unconditioned response, (ii) appetitive conditioned response, (iii) aversive unconditioned response, and (iv) aversive conditioned response. When CS1 was visual and CS2 was auditory (see high tone and low tone on the left side of Fig. 7), Darwin VII made the appropriate appetitive and aversive conditioned responses to auditory stimuli. However, when CS₁ was auditory and CS₂ was visual, Darwin VII responded incorrectly to visual aversive stimuli (see blobs on the left side of Fig. 7). As we discuss later, this resulted from the fact that, with this sequence of conditioned stimuli in the aversive learning condition, the blocks were avoided before gripping, and thus taste reinforcement could not occur. Examination of the synaptic weights between area IT and the motor neuronal units in this case showed that the connection strengths from IT to M_{app} were greater than from IT to M_{ave} (Fig. 8A). By altering the synaptic efficacy function (Fig. 8 inset), we were able to assure that aversive stimuli evoked a stronger learning response than appetitive stimuli (Fig. 8B). This change led to more balanced synaptic weights and more appropriate behavioral responses (Fig. 7, right side).

Discussion

Brains in organisms as complex as vertebrates differ in many ways from digital computers or Turing machines. Unlike computer inputs, signals from the environment are not unequivocally coded. Brains of different individuals within a species vary

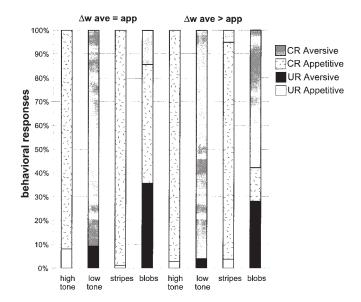


Figure 7. Behavioral responses after second-order conditioning, averaged over 7 Darwin VII subjects. The figure shows the overall percentages of each of the four possible behavioral responses to a stimulus. Left panel: synaptic efficacy rule governing learning rate was the same for aversive and appetitive events. When CS₁ was visual and CS₂ was auditory, 92% of the responses to appetitive auditory stimuli (high tone) were appetitive conditioned responses and 91% of the responses to aversive auditory stimuli (low tone) were aversive conditioned responses. When CS₁ was auditory and CS₂ was visual, 99% of the responses to appetitive visual stimuli (stripes) were appetitive conditioned responses but only 14% of the responses to aversive visual stimuli (blobs) were aversive conditioned responses. Right panel: synaptic efficacy rule governing learning rate was adjusted to be greater for aversive events than appetitive events (see Fig. 8 inset). When CS₁ was visual and CS₂ was auditory, nearly all the responses were conditioned responses (97% for appetitive and 96% for aversive). When CS₁ was auditory and CS2 was visual, the conditioned responses were 91% for appetitive and increased to 57% for aversive as compared to 14% in the previous condition (see left panel). The number of 'incorrect' appetitive conditioned responses to the aversive blob-patterned stimulus dropped from 50% to 14%. The appetitive conditioned responses could all be attributed to the behavior of one of the seven Darwin VII subjects. If this subject was omitted, there were no inappropriate appetitive conditioned responses and the proportion of aversive conditioned responses increased to 67%.

enormously in development, history and fine-scale physiology. Sensorimotor experience is also highly individual despite characteristic species behaviors. Moreover, perceptual categorizations and memories are not simple replicas of world input patterns. Instead, these products of higher brain functions adapt in a species-specific fashion to environmental change. Since many of the functions of individual brains result from complex dynamic interactions at a variety of levels, the elucidation of underlying mechanisms requires simultaneous measurements and sampling across these levels. In living animals, these are difficult to obtain and compare.

These considerations suggest that synthetic modeling of the kind described in this paper may be a useful strategy in attempts to understand higher brain functions. The behavior of Darwin VII shows that a synthetic brain-based device operating on biological principles and without pre-specified instructions can carry out perceptual categorization and conditioned responses. The successful performance of the device rests on the selectional modulation of its neuronal activity by behavior as well as on the existence of constraints provided by its value system. In both the perceptual categorization and conditioning experiments, the development of categorical responses required exploration of the environment and sensorimotor adaptation through specific and highly individual changes in connection strengths. We

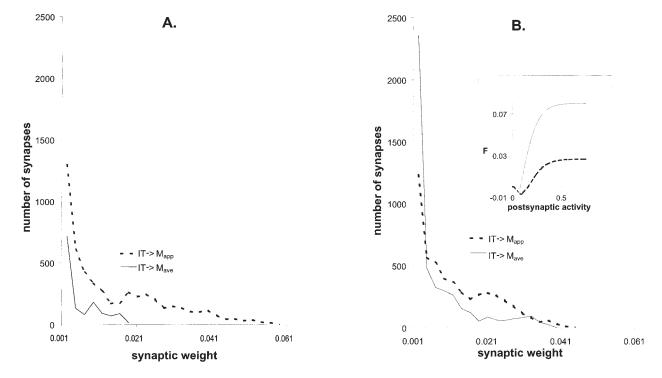


Figure 8. Distribution of synaptic weights for the connections between neuronal units in areas IT and motor behavior units (Mapp and Mave) for all seven Darwin VII subjects. The distributions show the number of synapses with different synaptic weight values. Summing across all subjects, each projection, $I\vec{T} \rightarrow M_{ave}$ and $IT \rightarrow M_{app}$, has ~15 000 synapses, but only those weights that changed from their initial value are shown in the figures. (A) Appetitive and aversive learning rates for the synaptic efficacy rules are equal (lower trace in inset; see function F in Appendix, C. Activity-dependent Synaptic Plasticity). The synaptic weights for the projection from IT to M_{app} are higher than those for the projection from IT to M_{app}. Approximately 5000 $IT \rightarrow M_{app}$ synapses increased from their initial weight, whereas ~1300 $IT \rightarrow M_{ave}$ synapses increased from their initial weight. (B) Synaptic efficacy is higher for the IT to M_{ave} connection (upper trace in inset) than the IT to M_{app} connections (lower trace in inset). Synaptic weight distributions for the projection from IT to M_{ave} are similar to those from the projection from IT to M_{app} . In both $IT \rightarrow M_{app}$ and $IT \rightarrow M_{app}$, ~5000 synapses increased from their initial weight. The inset shows the learning functions for the rule for synaptic efficacy change. For the definition of the synaptic efficacy function F, see the Appendix, part C. Parameters for the different projections are given in Table A2.

observed Darwin VII's overall behavior while at the same time recording the state of every neuronal unit and synaptic connection in its simulated nervous system. By collecting these neuronal data, we were able to demonstrate the development of neuronal groups during categorization and recognition by individual subjects (Fig. 5), to show that reliable classification of responses to visual stimuli could be based on the sampling of a small sub-population of neuronal units (Fig. 6), and to relate learning responses to functional changes in synaptic efficacy (Figs 7 and 8).

Darwin VII's nervous system has three features that are critical for understanding the mechanisms underlying perceptual categorization: (i) Connectivity from a topographically mapped primary area with transient activity to a non-topographically mapped higher area with more persistent activity. (ii) Sensory input that is continuous and temporally correlated with selfgenerated movement. (iii) Activity-dependent learning in which competitive mechanisms categorize sensory information and select for appropriate behavioral repertoires.

All of these features were necessary for Darwin VII to achieve invariant object recognition. Because a given visual stimulus spent more of the time in Darwin VII's gripper, VA_P neuronal units with receptive fields near the gripper were initially selected for and their corresponding connections to neuronal units in IT were potentiated (see Fig. 4, 'First Horizontal Shaped Block'). Once localized and patterned activity began in IT, it tended to sustain itself via local recurrent excitation combined with lateral inhibition. Continual overlapping input from VA_P as Darwin VII moved toward and away from the stimulus block (Almassy et al.,

1998) led to the reinforcement of the specific pattern of changes in synaptic strength from the retinotopically mapped VA_P neuronal units to the non-topographic populations of neuronal units in IT. By these means, the activity of VA_P neuronal units that drove IT neuronal units into a stimulus-dependent pattern of activity expanded from those with receptive fields near Darwin VII's gripper to those involving an almost complete coverage of the visual field (see Fig. 4, 'Fifth Horizontal Striped Block'). As a result, IT neuronal units were primed to respond to stimuli over a wider range of Darwin VII's visual field. Invariant object recognition is thus a system property that emerges dynamically from competitive neuronal group interactions within and between neural areas. These interactions differ from those of other models in which images are typically static and invariant object recognition is achieved by arranging features to line up across multiple views (Mundy and Zisserman, 1992; Mundy et al., 1992; Shashua, 1993; Weinshall, 1993), by deriving a learning rule that utilizes the temporal trace of neural activity (Wallis and Rolls, 1997; Rolls and Stringer, 2001), or by placing the main responsibility for invariance on neuronal properties alone (Tovee et al., 1994; Rolls and Tovee, 1995).

One striking characteristic of Darwin VII observed under all circumstances was the individuality of the patterns displayed by each subject's neural responses even for repetitions of the same behavior (see Fig. 5). This is consistent with the observation that adaptive behaviors tend to remain similar despite changes in context and variance in system properties resulting from multiple interactions across circuitry, plastic synaptic connections, fluctuating value systems, and variable object encounters. Thus, Darwin VII is structurally and behaviorally degenerate: different circuits and dynamics can yield similar behavior (Tononi *et al.*, 1999; Edelman and Gally, 2001). The developmental experiments comparing responses to strongly biased samples of appetitive or aversive stimuli indicate, however, that even with identical starting architectures, changes in experiential sequences can have profound effects. While this has been documented phenomenologically with living organisms, the experiments reported here may suggest possible mechanisms underlying such epigenetic biases.

The ability to modify various levels of control in Darwin VII provides insights into the neural mechanisms of learning during conditioning. For example, when CS1 was an auditory cue and CS₂ was a visual cue, our second-order conditioning experiments revealed an asymmetry that was initially unexpected: a predominance of appetitive conditioned responses over aversive responses that is analogous to the psychological phenomenon of overshadowing. Overshadowing occurs when an intense, salient stimulus gains control of responses at the expense of another less salient stimulus (Pavlov, 1927; Staddon, 1983). In the secondorder conditioning experiments where the CS₁ was an auditory cue and CS2 was a visual cue, behavior similar to that of overshadowing occurred in Darwin VII for two reasons. First, because of the simple tonotopic mapping in A1, responses to auditory stimuli were stronger and easier to categorize than visual stimuli. No overshadowing occurred when CS₁ was visual and CS₂ was auditory, since visual categories in IT and the appropriate behavioral response developed during primary conditioning when visual stimuli (CS₁) were paired with taste (US). Secondly, during the second stage of conditioning when both CS₁ and CS₂ were present, responses to the reinforcement (i.e. taste) of appetitive stimuli overshadowed aversive learning. This is attributable to the fact that after aversive learning the blocks were avoided before gripping and therefore taste reinforcement did not take place. Thus, in this sequence, Darwin VII generalized incorrectly that all visual stimuli were predictive of positive value. In the appetitive learning condition, this avoidance did not occur and reinforcement came from the US (taste) and CS₁ (auditory cue). Conditioning performance more consistent with animal models was obtained by altering the synaptic efficacy so that changes in plasticity were on average larger for aversive events than for appetitive events (see Figs 7 and 8). These results are consistent with the observation that the brain uses different learning rates for punishment and reward and that, in some cases, this difference may be crucial for an organism's survival (Garcia et al., 1955; Siucinska et al., 1999; O'Doherty et al., 2001).

The design of brain-based devices such as Darwin VII that possess neuroanatomical structure and large-scale neural dynamics differs fundamentally from that of robots. Unlike Darwin VII, robotic approaches using classical artificial intelligence are based on data representation, rule-driven algorithms, and the manipulation of formal symbol systems (Moravec, 1983; Nilsson, 1984). Artificial intelligence has been somewhat successful in emulating logical aspects of human behavior, but has been less successful in dealing with perception, categorization and movement in the world, which is a strength of synthetic neural models and brain-based devices (Reeke and Edelman, 1988; Pfeifer and Scheier, 1997). Purely reactive or behaviorbased robots carry out actions that are controlled through arbitration of several primitive behavioral repertoires without neural architectures (Brooks, 1986; Arkin, 1993). Evolutionary robotics, in which control systems are selected after each trial or lifetime according to a fitness function (Nolfi and Floreano,

2000), can evolve complex behaviors with very simple systems, but also do not emphasize neuronal responses. A recent hybrid between evolutionary algorithms and artificial neural network learning rules was designed to mutate learning rules between trials, allowing learning during the lifetime of the robot (Floreano and Mondada, 1998). Typically, however, the artificial neural networks controlling the evolutionary robot's behavior were small (on the order of tens of artificial 'neural units') and they did not reflect neuroanatomical organization.

In its present form, Darwin VII has several limitations. In comparison to organisms that its behavior mimics, it has an extremely simple nervous system. Re-entrant connections (Edelman, 1987) within a neural area are present in the model, but re-entrant connections between neural areas, such as *A1* and *IT*, are currently not implemented. This limits intra-modal and cross-modal interactions, making its behavior purely stimulus driven. Moreover, the activity in a simulation cycle is the average of a relatively small population of neurons over 100–200 ms, and the spiking dynamics of individual neurons cannot presently be explored with this model. Despite these limitations, Darwin VII's performance shows that, regardless of the existence of individual variance, neurally based devices acting in the real world can carry out consistent behaviors.

One might ask why the simulation must include behaviors in the real world. Why not simulate the environment as well as the brain? The answer rests in the constructive nature of the brain and behavioral responses to real-world inputs (Chiel and Beer, 1997; Clark, 1997). For example, to specify the outlines of an environmental object in a pure computer simulation of the environment would contribute an *a priori* bias in the form of a detailed albeit implicit instruction. In contrast, by acting in the real world, Darwin VII decides 'for itself' on object properties and then constructs appropriate responses. By using a real-world environment, not only is the risk of introducing biases into the model reduced, but also the experimenter is freed from the substantial burden of constructing a highly complex simulated environment (Edelman *et al.*, 1992).

Although the world of Darwin VII is much simpler than a real econiche, there does not seem to be a fundamental restriction on constructing a more complex phenotype to deal with a richer environments. Experiments exploring the effects of different neuroanatomical arrangements, the effects of lesions, and of altered synaptic responses are also now possible. As in the present experiments, the behaviors of the resulting brain-based devices would emerge solely as a result of internally generated activity of their nervous systems rather than of responses to any programmed instructions from computer software. Devices of this kind might prove useful in situations of novelty where computation is not possible or in cases of great local complexity where programming proves infeasible. In the near future, such devices are not likely to include behaviors as rich as those of higher vertebrates, and therefore their greatest practical use may at present be to complement computers in a hybrid arrangement, i.e. a brain-based device linked to a conventional digital computer. Since the fundamental operation of such devices includes random fluctuations and unpredictable behaviors, they are not in any strict sense Turing machines. Although the phrase 'machine psychology' may thus appear to be a misnomer, it may be nevertheless be usefully applied to the behavior of non-living things that learn. In any case, providing such synthetic constructions with increasingly sophisticated neural circuits and body forms should give further valuable insights into the relationships between brain, body and behavior.

Notes

This work was supported by the W.M. Keck Foundation and the Neurosciences Research Foundation, which supports The Neurosciences Institute. We thank J. A. Snook for his contribution to the design of Darwin VII. We are grateful to Drs J. A. Gally and G. N. Reeke for detailed criticism.

Address correspondence to Jeffrey L. Krichmar, W.M. Keck Laboratory, The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, USA. Email: krichmar@nsi.edu.

References

- Almassy N, Edelman GM, Sporns O (1998) Behavioral constraints in the development of neuronal properties: a cortical model embedded in a real-world device. Cereb Cortex 8:346–361.
- Arkin RC (1993) Modeling neural function at the schema level: implications and results for robotic control. In: Biological neural networks in invertebrate neuroethology and robotics (Beer RD, Ritzmann RE, McKenna T, eds), pp. 383-409. Boston, MA: Academic Press.
- Aston-Jones G, Bloom FE (1981) Norepinephrine-containing locus coeruleus neurons in behaving rats exhibit pronounced responses to non-noxious environmental stimuli. J Neurosci 1:887–900.
- Bialek W, Zee A (1990) Coding and computation with neural spike trains. J Stat Phys 59:103–115.
- Bienenstock EL, Cooper LN, Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci 2:32-48.
- Brooks RA (1986) A robust layered control system for a mobile robot. IEEE I Robotics Automation 2:14–23.
- Brown EN, Frank L, M, Tang D, Quirk MC, Wilson MA (1998) A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. J Neurosci 18:7411–7425.
- Chiel HJ, Beer RD (1997) The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. Trends Neurosci 20:553–557.
- Clark A (1997) Being there. Putting brain, body, and world together again. Cambridge, MA: MIT Press.
- Clark SA, Allard T, Jenkins WM, Merzenich MM (1988) Receptive fields in the body-surface map in adult cortex defined by temporally correlated inputs. Nature 332:444-445.
- Crair MC, Gillespie DC, Stryker MP (1998) The role of visual experience in the development of columns in cat visual cortex. Science 279: 566-570.
- Edelman GM (1987) Neural Darwinism: the theory of neuronal group selection. New York: Basic Books.
- Edelman GM, Gally JA (2001) Degeneracy and complexity in biological systems. Proc Natl Acad Sci USA 98:13763-13768.
- Edelman GM, Reeke GN, Gall WE, Tononi G, Williams D, Sporns O (1992) Synthetic neural modeling applied to a real-world artifact. Proc Natl Acad Sci USA 89:7267–7271.
- Floreano D, Mondada F (1998) Evolutionary neurocontrollers for autonomous mobile robots. Neural Netw 11:1461–1478.
- Fox K (1995) The critical period for long-term potentiation in primary sensory cortex. Neuron 15:485-488.
- Friston KJ, Tononi G, Reeke GN, Sporns O, Edelman GM (1994) Value-dependent selection in the brain: simulation in a synthetic neural model. Neuroscience 59:229-243.
- Garcia J, Kimeldorf DJ, Koelling RA (1955) A conditioned aversion towards saccharin resulting from exposure to gamma radiation. Science 122:157-159.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233: 1416-1419.
- Ito M, Tamura H, Fujita I, Tanaka K (1995) Size and position invariance of neuronal responses in monkey inferotemporal cortex. J Neurophysiol 73:218–226.
- Kato N, Artola A, Singer W (1991) Developmental changes in the susceptibility to long-term potentiation of neurones in rat visual cortex slices. Brain Res Dev Brain Res 60:43–50.
- Kilgard MP, Merzenich MM (1998) Cortical map reorganization enabled by nucleus basalis activity. Science 279:1714–1718.
- Kirkwood A, Lee HK, Bear MF (1995) Co-regulation of long-term

- potentiation and experience-dependent synaptic plasticity in visual cortex by age and experience. Nature 375:328-331.
- Kobatake E, Wang G, Tanaka K (1998) Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. J Neurophysiol 80:324–330.
- Krichmar JL, Snook JA, Edelman GM, Sporns O (2000) Experience-dependent perceptual categorization in a behaving real-world device.
 In: Animals to animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior (Meyer J-A, Berthoz A, Floreano D, Roitblat H, Wilson SW, eds), pp. 41-50.
 Cambridge, MA: MIT Press.
- Ljungberg T, Apicella P, Schultz W (1992) Responses of monkey dopamine neurons during learning of behavioral reactions. J Neurophysiol 67:145-163.
- Moravec HP (1983) The Stanford cart and the CMU rover. Proc IEEE 71:872-884.
- Mundy J, Zisserman A, eds (1992) Geometric invariance in computer vision. Cambridge, MA: MIT Press.
- Mundy JL, Welty RP, Brill MH, Payton PM, Barrett EB (1992) 3-D model alignment without computing pose. Image understanding workshop, Morgan Kaufmann, San Mateo, CA.
- Nilsson N (1984) Shakey the robot. Menlo Park, CA: SRI International.
- Nolfi S, Floreano D (2000) Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines. Cambridge, MA: MIT Press.
- O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) Abstract reward and punishment representations in the human orbitofrontal cortex. Nat Neurosci 4:95–102.
- Pavlov IP (1927) Conditioned reflexes. London: Oxford University Press.Pfeifer R, Scheier C (1997) Sensory-motor coordination: the metaphor and beyond. Robotics Autonomous Syst 20:157-178.
- Reeke GN, Edelman GM (1988) Real brains and artificial intelligence. Daedalus Proc Am Acad Arts Sci 117:143-173.
- Reeke GN, Sporns O, Edelman GM (1990) Synthetic neural modeling: the 'Darwin' series of recognition automata. Proc IEEE 78:1498–1530.
- Rescorla RA (1980) Pavlovian second-order conditioning: studies in associative learning. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rolls ET, Stringer SM (2001) Invariant object recognition in the visual system with error correction and temporal difference learning. Netw Comput Neural Syst 12:111–129.
- Rolls ET, Tovee MJ (1995) The responses of single neurons in the temporal visual cortical areas of the macaque when more than one stimulus is present in the receptive field. Exp Brain Res 103:409–420.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599.
- Shashua A (1993) Projective depth: a geometric invariant for 3D reconstruction from two perspective/orthographic views and for visual recognition. International Conference on Computer Vision,
- Siucinska E, Kossut M, Stewart MG (1999) GABA immunoreactivity in mouse barrel field after aversive and appetitive classical conditioning training involving facial vibrissae. Brain Res 843:62-70.
- Sporns O, Almassy N, Edelman GM (2000) Plasticity in value systems and its role in adaptive behavior. Adaptive Behav 8:129–148.
- Staddon JER (1983) Adaptive behavior and learning. Cambridge: Cambridge University Press.
- Tanaka K (1996) Inferotemporal cortex and object vision. Annu Rev Neurosci 19:109–139.
- Tanaka K, Saito H, Fukada Y, Moriya M (1991) Coding visual images of objects in the inferotemporal cortex of the macaque monkey. J Neurophysiol 66:170-189.
- Theunissen FE, Miller JP (1991) Representation of sensory information in the cricket cercal sensory system. II. Information theoretic calculation of system accuracy and optimal tuning-curve widths of four primary interneurons. J Neurophysiol 66:1690–703.
- Tononi G, Sporns O, Edelman GM (1999) Measures of degeneracy and redundancy in biological networks. Proc Natl Acad Sci USA 96: 3257–3262.
- Tovee MJ, Rolls ET, Azzopardi P (1994) Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. J Neurophysiol 72:1049–1060.
- Turrigiano GG, Nelson SB (2000) Hebb and homeostasis in neuronal plasticity. Curr Opin Neurobiol 10:358–364.
- Wallis G, Rolls ET (1997) Invariant face and object recognition in the visual system. Prog Neurobiol 51:167-194.

Weinshall D (1993) Model based invariants for 3-D vision. Int J Comput Vis 10:27-42.

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. Science 261:1055-1058.

Appendix: Specifics of Neuronal Responses, Input and Output in Darwin VII

A. Computation

The simulated nervous system was run on a Quad Pentium III Xeon Linux workstation capable of communicating with the mobile base. The workstation received visual input via radio frequency (RF) video transmission from a CCD camera mounted on the mobile base (see Appendix, part E). The workstation received auditory and gripper information, and transmitted motor and actuator commands via an RF modem (see Appendix, part F).

B. Neuronal Unit Activity

The total contribution of input to unit i is given by

$$A_i(t) = \sum_{l=1}^{M} \sum_{j=1}^{N_l} c_{ij} s_j(t)$$

where M is the number of different anatomically defined connection types (see Table A2) and N_l is the number of connections per type M projecting to unit i. Negative values for c_{ij} corresponded to inhibitory connections. The activity level of unit i is given by:

$$s_i(t+1) = \phi \left(\tanh \left(g_i \left(A_i(t) + \omega s_i(t) \right) \right) \right)$$

where

$$\phi(x) = \begin{cases} 0; & x < \sigma_i \\ x; & \text{otherwise} \end{cases}$$

and ω determined the persistence of unit activity from one cycle to the next, σ_i is a unit specific firing threshold, and g_i is a scale factor. Specific parameter values for neuronal units are given in Table A1.

C. Activity-dependent Synaptic Plasticity

Activity-dependent synaptic changes in c_{ij} were given by

$$\Delta c_{ij}(t+1) = \varepsilon \left(c_{ij}(0) - c_{ij}(t)\right) + \eta s_j(t) F(s_i(t))$$

where $s_i(t)$ and $s_j(t)$ are activities of post- and pre-synaptic units, respectively, η is a fixed learning rate, ε is a decay constant, and $c_{ij}(0)$ is the initial (t=0) weight of connection c_{ij} . The decay constant ε governed a passive, uniform decay of synaptic weights to their original starting values. The function F, similar to the BCM learning rule (Bienenstock et al., 1982), determines limits on potentiation and depression that depend upon post-synaptic activity. F was implemented as a piecewise linear function defined by two thresholds $(0 < \theta_1 < \theta_2 < 1)$, two inclinations (k_1, k_2) and a saturation parameter ρ (ρ = 6 throughout).

$$F(s) = \begin{cases} 0; & s < \theta_1 \\ k_1(\theta_1 - s) & \theta_1 \le s < (\theta_1 + \theta_2)/2 \\ k_1(s - \theta_2) & (\theta_1 + \theta_2)/2 \le s < \theta_1 \\ k_2 \tanh(\rho(s - \theta_2))/\rho & \text{otherwise} \end{cases}$$

See Figure 8 inset for a representative chart of the function. Specific parameter settings for fine-scale synaptic connections are given in Table A2.

Table A1
Values of parameters defining properties of neuronal units in Darwin VII

Area	Size	g	σ	ω
R	64 × 64	1.50	0.20	0.00
VA _P -B	64×64	1.33	0.65	0.50
VA _P -H, VA _P -V	64×64	1.50	0.50	0.50
IT	28×28	1.10	0.04	0.03
IT_i	14×14	1.35	0.02	0.15
LCoch, RCoch	1×64	2.00	0.10	0.30
A1	28×28	1.50	0.50	0.50
$A1_i$	28×28	1.35	0.02	0.15
M_{app} , M_{ave}	3×6	2.00	0.10	0.30
T_{app} , T_{ave}	3×6	2.00	0.10	0.30
S	2×2	2.00	0.05	0.15
S_o	4×4	3.00	0.15	0.22
S_i	2×2	2.00	0.10	0.22
C	20 × 20	1.33	0.65	0.50

D. Value-dependent Synaptic Plasticity

A value term was computed as

$$V(d) = 1 + f(d) \frac{\overline{S} + v(d-1)(d-1)}{d}$$

where d is the delay of the value term and is incremented every simulation cycle after the onset of area S activity with a range of 1-9, \overline{S} is the average activity in area S, and v(d-1) is the value of V at time d-1. f is a convolution function that scales the activity over the delay period with values of 0.1, 0.1, 0.3, 0.7, 1.0, 1.0, 0.7, 0.3 and 0.1 for the nine delay increments. The effect of this convolution is to delay onset of the value system activity and spread the activity over time. The synaptic change for value-dependent synaptic plasticity was given by:

$$\Delta c_{ij}(t+1) = \varepsilon \left(c_{ij}(0) - c_{ij}(t)\right) + \eta s_j(t) F(s_i(t)) V(d)$$

E. Visual System and its Input

The CCD camera sent 320 × 240 pixel monochrome video images, via an RF transmitter, to an ImageNation PXC200 frame grabber attached to the computer running the neural simulation. The image was clipped, such that only the center square of the image remained, and it was then spatially averaged to produce a 64×64 pixel image. Each pixel was normalized between 0 (black) and 1 (white) and mapped directly to neuronal units of area R in the neural simulation. R neuronal units projected retinotopically to neuronal units in neural area VA_P , which in turn projected to neural area IT non-topographically (see Fig. 2 and Table A2).

F. Auditory System and its Inputs

Microphone input was amplified and filtered in hardware. An RMS (root mean square) chip measured the amplitude of the signal and a comparator chip produced a square waveform which allowed frequency to be measured. Every millisecond, the microcontroller on NOMAD calculated the overall microphone amplitude by averaging the current signal amplitude measurement with the previous three signal amplitude measurements. The microcontroller calculated the frequency of the microphone signal at each time point by inverting the average period of the last eight square waves. LCoch and RCoch each had 64 neuronal units. Their response was based on the frequency and amplitude information received from the microcontroller via the RF modem. Each cochlear neuronal unit had a cosine tuning curve with a tuning width of 1 kHz and a preferred frequency, which ranged over the ensemble of units from 2.9 to 4.2 kHz. Activity of a cochlear neuronal unit was obtained by multiplying the value from its cosine tuning curve by the amplitude of the microphone signal. Cochlear neuronal units projected tonotopically to neuronal units in neural area A1 (see Fig. 2 and Table A2).

Table A2 Properties of anatomical projections and connection types in Darwin VII

Projection	Р	Arbor	$c_{ij}(0)$	η	ε	θ_1	θ_2	<i>k</i> ₁	k ₂
R→VA _P -B	1.00	03×3	0.03, 0.03	0.0	0.0	0.0	0.0	0.0	0.0
$R \rightarrow VA_{p}-H$	1.00	$[] 0 \times 4$	0.04	0.0	0.0	0.0	0.0	0.0	0.0
$R \rightarrow VA_P - V$	1.00	$[]4 \times 0$	0.04	0.0	0.0	0.0	0.0	0.0	0.0
Early: $VA_PB \rightarrow IT$	0.01	non-topo	0.04, 0.08	0.125	0.00125	0.10	0.25	0.45	0.45
Mid: $VA_PB \rightarrow IT$	0.01	non-topo	0.04, 0.08	0.05	0.0005	0.10	0.25	0.45	0.45
Late: $VA_PB \rightarrow IT$	0.01	non-topo	0.04, 0.08	0.03	0.0003	0.10	0.25	0.45	0.45
Early: VA_PH , $VA_PV \rightarrow IT$	0.0175	non-topo	0.04, 0.08	0.125	0.00 125	0.10	0.25	0.45	0.45
Mid: VA_PH , $VA_PV \rightarrow IT$	0.0175	non-topo	0.04, 0.08	0.05	0.0005	0.10	0.25	0.45	0.45
Late: VA_PH , $VA_PV \rightarrow IT$	0.0175	non-topo	0.04, 0.08	0.03	0.0003	0.10	0.25	0.45	0.45
$ T\rightarrow T_i $	1.00	Θ2,3	0.06, -0.06	0.0	0.0	0.0	0.0	0.0	0.0
$IT_i \rightarrow IT$	1.00	[] 1 × 1	-0.36, -0.50	0.0	0.0	0.0	0.0	0.0	0.0
<i>IT</i> → <i>IT</i>	1.00	0 1	0.10, 0.14	0.0	0.0	0.0	0.0	0.0	0.0
$IT \rightarrow M_{app}, M_{ave} \dagger$	0.15	non-topo	0.0006, 0.0010	0.006	0.00006	0.01	0.16	0.1	0.16
$IT \rightarrow M_{ave} \uparrow$ (see note)	0.15	non-topo	0.0006, 0.0010	0.006	0.00006	0.005	0.08	0.1	0.48
$IT \rightarrow S_{o} \uparrow$	0.05	non-topo	0.0005, 0.0015	0.02	0.002	0.01	0.18	0.05	0.05
$R \rightarrow C$	0.50	[] 2 × 2	0.10	0.10	0.0	0.0	0.0	0.0	0.0
Early: LCoch, RCoch→A1	0.65	[] 13×2	0.10, 0.13	0.01	0.0001	0.14	0.29	0.45	0.15
Mid: LCoch, RCoch→A1	0.65	[] 13×2	0.10, 0.13	0.005	0.00005	0.14	0.29	0.45	0.15
Late: LCoch, RCoch→A1	0.65	[] 13×2	0.10, 0.13	0.003	0.00003	0.14	0.29	0.45	0.15
$A1 \rightarrow A1_i$	1.00	[] 8×18	0.06, 0.06	0.0	0.0	0.0	0.0	0.0	0.0
$A1_i \rightarrow A1$	1.00	0 1	-0.30, -0.46	0.0	0.0	0.0	0.0	0.0	0.0
$A1 \rightarrow A1$	0.25	0 1	0.05, 0.075	0.0	0.0	0.0	0.0	0.0	0.0
$A1 \rightarrow M_{app}, M_{ave}^{\dagger}$	0.15	non-topo	0.0006, 0.0010	0.006	0.00006	0.01	0.16	0.1	0.16
$A1 \rightarrow M_{ave} \uparrow$ (see note)	0.15	non-topo	0.0006, 0.0010	0.006	0.00006	0.005	0.08	0.1	0.48
$A1 \rightarrow S_0 \dagger$	0.05	non-topo	0.0005, 0.0015	0.02	0.002	0.01	0.18	0.05	0.05
$T_{app}, T_{ave} \rightarrow S_o, M_{app}, M_{ave}$	1.00	0 1	0.12, 0.12	0.0	0.0	0.0	0.0	0.0	0.0
$M_{app} \longleftrightarrow M_{ave}$	1.00	non-topo	-0.12, 0.12	0.0	0.0	0.0	0.0	0.0	0.0
$S_i \rightarrow S$	1.00	$[] 2 \times 2$	-0.27, -0.30	0.0	0.0	0.0	0.0	0.0	0.0
$S_o \rightarrow S$	1.00	[] 2 × 2	0.09, 0.11	0.0	0.0	0.0	0.0	0.0	0.0
$S_0 \rightarrow S_i$	0.80	[] 2 × 2	0.05, 0.06	0.0	0.0	0.0	0.0	0.0	0.0

A pre-synaptic neuronal unit is connected with a post-synaptic neuronal unit with a given probability (P) and given projection shape (Arbor). This arborization shape can be rectangular '[]' with a height and width (h × w), circular '0' with a radius (r), doughnut shaped '\text{\Theta}' with the shape constrained by an inner and outer radius (r_1, r_2), or non-topographical 'non-topog' where any pairs of pre-synaptic and post-synaptic neuronal units have a given probability of being connected. The initial connection strengths, $c_{ii}(0)$, are set randomly within the range given by a minimum and maximum value (min, max). A negative value for $c_{ij}(0)$ indicates inhibitory connections. Projections marked † are value-dependent. Non-zero values for η, ε, θ1, θ2, k_1 and k_2 signify activity-dependent plastic connections. In the perceptual categorization experiments, the properties of the projections were not altered over time. In the conditioning experiments, we modeled a transition from a critical learning period to a later adult period. Specifically, the learning rate η and the decay rate ε , for primary sensory (VA_P or LCoch/RCoch) to higher sensory (IT or A1) projections decreased based on stimulus exposure. 'Early' corresponded to the 1st through 20th exemplars for a given modality (i.e. auditory: LCoch/RCoch->A1; visual: VAp->/T), 'Mid' corresponded to the 21st through 40th exemplars, and 'Late' corresponded to approximately the 41st through 60th exemplars. Note that higher learning rates for aversive events were used in some of the conditioning experiments (see text and Fig. 8 for details).

G. Sampling of IT Activity for Classification of Responses

In order to test the classification of responses to visual stimuli based on ITactivity, the patterns of activity in IT during the development of visual categories were compared with templates consisting of individual patterns of IT activity in response to visual stimuli after categorization had developed. Since patterns of activity varied for each Darwin VII subject, a separate template and comparison needed to be made for each individual. A template was created for each Darwin VII subject by taking the average activity of neuronal units in area IT in response to the last 10 presentations of a particular stimulus class. Templates were made for each of the visual stimulus classes (blob, horizontal, vertical) as well as a template with random activity which was achieved by shuffling the blob template. The metric used to compare activity of IT with the templates

$$\hat{s} = \min_{j} \left\{ \sum_{i} \sqrt{\left(s_{j}(i) - IT(i)\right)^{2}} \right\}$$

where j is an index over the stimulus classes, and s_j is the template for stimulus j, i is an index into the neuronal units, IT is the current activity of neuronal area IT, and s is the predicted stimulus. A classification was regarded as correct if the predicted stimulus, s, was equal to the observed stimulus. In results reporting a sub-sample of IT activity, a randomly chosen percentage of IT neuronal units was selected for each Darwin VII subject. To ensure that the accuracy did not depend on a specific subset of neuronal units, 10 such random samples were averaged together for each Darwin VII subject. Thus, the accuracy of classification reported in Figure 6 is the average of seven Darwin VII subjects where the classification accuracy of each individual subject is based on an average of 10 random samples (see Fig. 6).