2 Neurorobotics: Neuroscience and Robots

Tiffany J. Hwu and Jeffrey L. Krichmar

2.1 Introduction

Neurorobotics is the study of the interaction between neural systems and their physical embodiments on robotic platforms. Since the brain is strongly coupled with the body and situated within the surrounding environment, neurorobots can be a powerful tool for studying the intricate interactions between neural systems and the outside world. Neurorobotics also serves as a way to create autonomous systems that capture the advantages of biology for intelligent behavior. Compared to the general study of cognitive robotics, neurorobotics centers around biological brain functions—for example, the neural circuitry and functional anatomy that support basic cognitive processes. This chapter provides our viewpoints on this field, highlights some of its milestone events, and talks about its future potential.

2.2 Foundational Ideas in Neurorobotics

Many believe that neurorobotics got its beginning with Grey Walter's tortoises, which had simple light sensors and collision detectors attached to a basic analog circuit. His first robots, Elmer and Elsie, were programmed with simple reflexive neural circuits that controlled their movements based on the sensors. Despite the simplicity of these robots, complex and interesting behaviors emerged. For instance, one robot was placed in front of a mirror with a light on its nose. The robot started to react to its own presence in what could be interpreted as narcissistic behavior.

Braitenberg vehicles were another important example of complex behaviors emerging from simple circuitry. First introduced in the book titled *Vehicles* by Valentino Braitenberg (1986), a series of simple robots showed how basic neural circuits could create complex behaviors, some of which could even be attached to abstract human notions, such as emotion, with vehicle names like Fear, Aggression, Love, and Exploration. Each of these vehicles contained a light sensor and a motor on the left and right sides. In the vehicle displaying *fear*, the speed of each motor was directly proportional to the amount of light sensed by the sensor on the equivalent side. This caused the vehicle to speed away from the stimulus source, as if in fear. However, just crossing the wires caused the vehicle to speed toward the stimulus, as if in aggression. This simple robot provided an important

neuroscience lesson on the function of ipsilateral and contralateral connections in the nervous system. By making the motor speeds inversely proportional to the sensors, the vehicle displaying *fear* could turn into *love*, slowing down its movement toward the stimulus. Likewise, aggression then turned into exploration, gently seeking to be away from the stimulus. In this way, Braitenberg demonstrated how changing the balance of excitatory and inhibitory connections can affect behavior. Although the circuits themselves were simple, it was easy to place human interpretations on the resulting behaviors, teaching an important lesson that complex cognitive functions may actually be composed of very simple mechanics.

The Keck Machine Psychology Laboratory at the Neurosciences Institute in La Jolla, California, was also a source of foundational contributions in neurorobotics. Director Gerald Edelman (1987, 1993), whose work in immunology led to the Nobel Prize, advocated his theory of the nervous system in a book titled *Neural Darwinism: The Theory of Neuronal Group Selection*. The theory suggested there was selection of neural circuits during development through synaptic pruning and selection of groups of neurons during adulthood through reentrant connections. Important for neurorobotics was the notion of value systems to tie environmental signals to neuronal groups, which led to the selection of behaviors important for survival. As Edelman would say, "The brain is embodied, and the body is embedded in the environment." Based on this idea, the group developed the Darwin series of Brain-Based Devices (Edelman et al. 1992; Reeke, Sporns, and Edelman 1990). Another phrase that drove this work was "The world is an unlabeled place," which meant that perceptual categories must be selected through experience, rather than supervision. These Brain-Based Devices were robots with large-scale neural networks controlling their behavior (figure 2.1). However, these were not the feedforward-input neural networks that were popular then and became the deep neural networks of today. The Brain-Based Device's neural networks contained anatomical details that resembled biological neural networks. There were sensory streams, top-down connections, and long-range connections between regions that were bidirectional as well as local lateral excitation and inhibition within brain regions. An early Brain-Based Device called Darwin V had an artificial nervous system that could learn preferences and predict the value of objects (Almassy, Edelman, and Sporns 1998). Although the robot was lumbering and did not exactly operate in real time, it did demonstrate operant conditioning and value-based learning.

One of the major venues in the early days of neurorobotics was the annual Simulation of Adaptive Behavior (SAB) conference. For example, SAB 2000 introduced a wide variety of exemplars, which would now be called neurorobots (Meyer et al. 2000). Arleo and Gerstner (2000) presented a model of head direction cells and hippocampal place cells, which was embodied on a Khepera robot, to demonstrate spatial navigation in the rodent. Arsenio (2000) created a neural circuit based on oscillators observed in the brain and showed how these could be used to realize humanoid arm movements and gait patterns. Collins and Wyeth (2000) introduced a cerebellar controller, based on Albus's cerebellar model arithmetic computer (CMAC) neural network, to overcome delays when planning trajectories. Gonzalez and colleagues (2000) constructed a basal ganglia model to show action selection in a mobile robot. The robot would find cylinders, pick them up, and deposit the cylinders outside the wall of the robot arena. At this same meeting, Darwin VII, a Brain-Based Device capable of perceptual categorization, was introduced (Krichmar et al. 2000). For more details on Darwin VII, see the case study below. This is just a sampling of the work going on at this time.

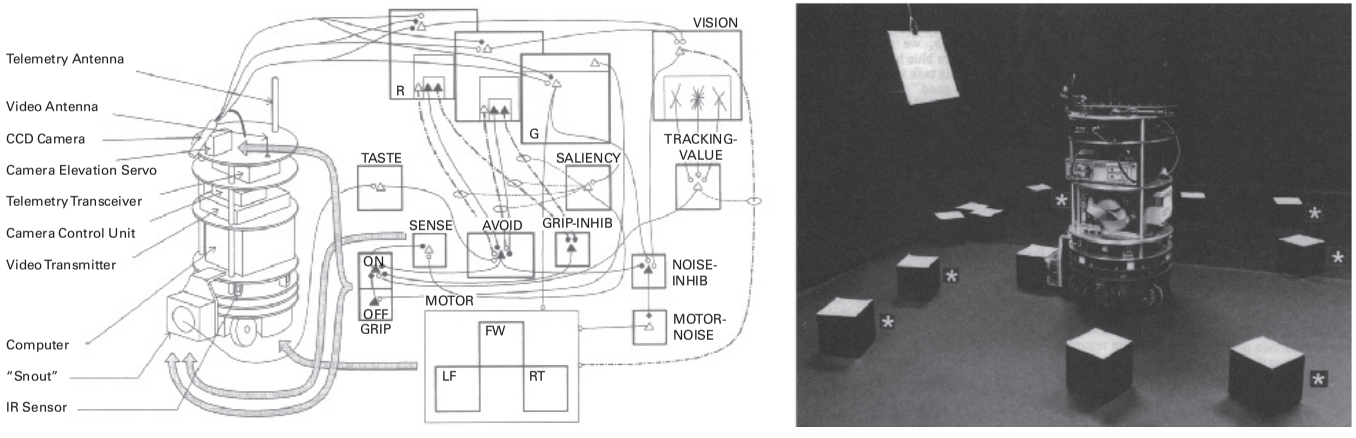

Figure 2.1
Darwin IV Brain-Based Device. *Left:* Neural network model to control Darwin IV's behavior. *Right:* Darwin IV in a conditioning task. *Source:* Adapted with permission from Edelman et al. 1992.

The theme connecting the wide range of methods, robots, and behaviors at SAB 2000 was that neural network models were used to study some aspect of neuroscience by demonstrating behavior in a physical robot. Many of the researchers in these studies were pivotal in establishing the field of neurorobotics as it is known today.

Around this time period, other groups were creating robot designs that could be included within the field of neurorobotics. Rather than building brain circuits, they were investigating how the body and brain interact and how neural networks may develop. For example, Tony Prescott and his group at the University of Sheffield studied whisking in the rodent and developed a robotic sensorimotor circuit with biomimetic whiskers (Pearson et al. 2011). Figure 2.2 shows their Whiskerbot, which was completed around 2005. Dario Floreano helped establish the field of evolutionary robotics (Nolfi and Floreano 2000). Floreano and colleagues used evolutionary algorithms to evolve neural networks that supported a range of behaviors from navigating mazes to developing predator-prey strategies (Floreano and Keller 2010). For more details, the reader should refer to chapter 4. Rolf Pfeifer and Josh Bongard (2006) had the insight that the "body shapes the way we think." They suggested that biological organisms perform morphological computation—that is, the body performs certain processes that would otherwise be performed by the brain.

Even though these biomimetic and evolutionary algorithms were not directly testing brain theories, they were increasing our knowledge of how the brain and body interact, and they were creating novel, biologically inspired algorithms and robot designs that would further the field of robots and AI.

As parallel-computing resources improved, some groups were approaching brain-scale neural simulations. Darwin VII's neural network contained approximately twenty thousand neurons and nearly five hundred thousand synaptic connections, all of which had to be updated in real time to keep up with the active vision and sensors. The Darwin

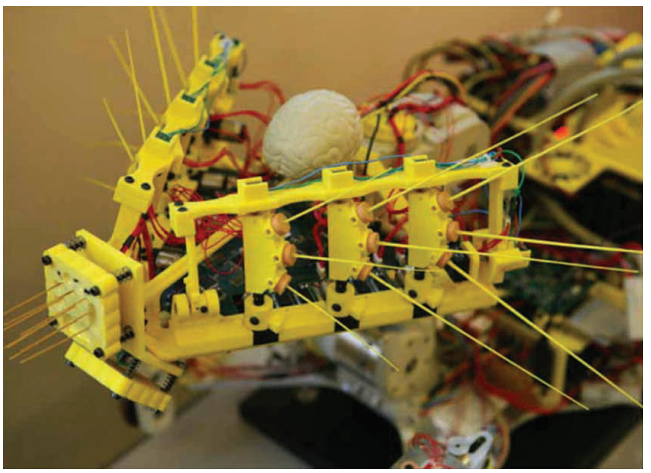

Figure 2.2
Whiskerbot from the University of Sheffield. Whiskerbot had two active whiskers and a detailed neural network model to convert whisker deflection signals into simulated spike trains. *Source:* Adapted with permission from Pearson et al. 2011.

team used a Beowulf cluster with Message Passing Interface (MPI) to achieve real-time performance. Phil Goodman's Virtual Neurobot project had at least one hundred thousand highly detailed neurons on a computer cluster. Although the robot was virtual, it did need to respond in real time to recognize intent and trust in a human actor (Bray et al. 2012).

During this time there was often pushback from the community about the necessity for large-scale modeling. Many interesting results could be achieved with smaller neural networks, often with fewer than one hundred neurons. However, solving a problem in certain domains with small neural networks was unavoidable. For example, a model of the visual cortex that tested theories of feature binding and invariant object recognition (Seth et al. 2004b) required a neuron at every camera pixel (or receptive field) for each feature (two colors and four orientations). Since the network simulated the expansion of visual cortex receptive fields combining primitive features into objects (i.e., $V1 \rightarrow V2 \rightarrow V4 \rightarrow IT$), a large-scale neural network was necessary. However, applying the same modeling detail to a neural network that encoded tactile features with whiskers resulted in an order-of-magnitude-smaller network (Seth et al. 2004a).

In addition to practical reasons, large-scale modeling is often required to realize the neuronal dynamics and anatomical pathways observed in brain responses. Although this fidelity results in highly complex networks, it does allow one to test theories of the brain and make better predictions. Preserving anatomical projections leads to large-scale heterogeneous architectures. Having large groups of neurons with biophysical properties leads to interesting neural dynamics, as was observed in a large-scale model of the hippocampus and surrounding regions (Krichmar, Nitz, et al. 2005). In this model the complex interplay between the entorhinal cortex and the hippocampal subfields resulted in the reliance on different functional pathways at different points in the robot's learning (figure 2.3). Using large-scale neural models does come with a cost beyond computing power. At some point the neural network becomes so complex that it is as difficult to understand as the real brain. Interestingly, the analysis of the large-scale hippocampus model required the development of new tools; one was a recursive backtrace through neural activity (Krichmar, Nitz, et al. 2005), and the other applied Granger causality to the simulated neural network (Krichmar, Seth, et al. 2005).

Nowadays, large-scale neural network models are the norm. Neuromorphic hardware can support brain-scale neural networks at very low power (Indiveri et al. 2011; Merolla et al. 2014; Davies et al. 2018). Deep neural networks with many hidden layers are regularly developed (LeCun, Bengio, and Hinton 2015). With tools such as PyTorch and TensorFlow, graphics processing unit (GPU) clusters, and cloud computing, large-scale neural networks are within the reach of most researchers and students. Moreover, it turns out that size, in the form of many layers, is necessary to solve more challenging problems, such as image recognition (Krizhevsky, Sutskever, and Hinton 2017) or human-level game playing (Mnih et al. 2015).

2.2.1 Case Study: Darwin VII—Perceptual Categorization and Conditioning in a Brain-Based Device

Darwin VII was one of the first neurorobots to demonstrate experience-dependent learning (i.e., learning by sampling the environment without supervisory signals) with a detailed,

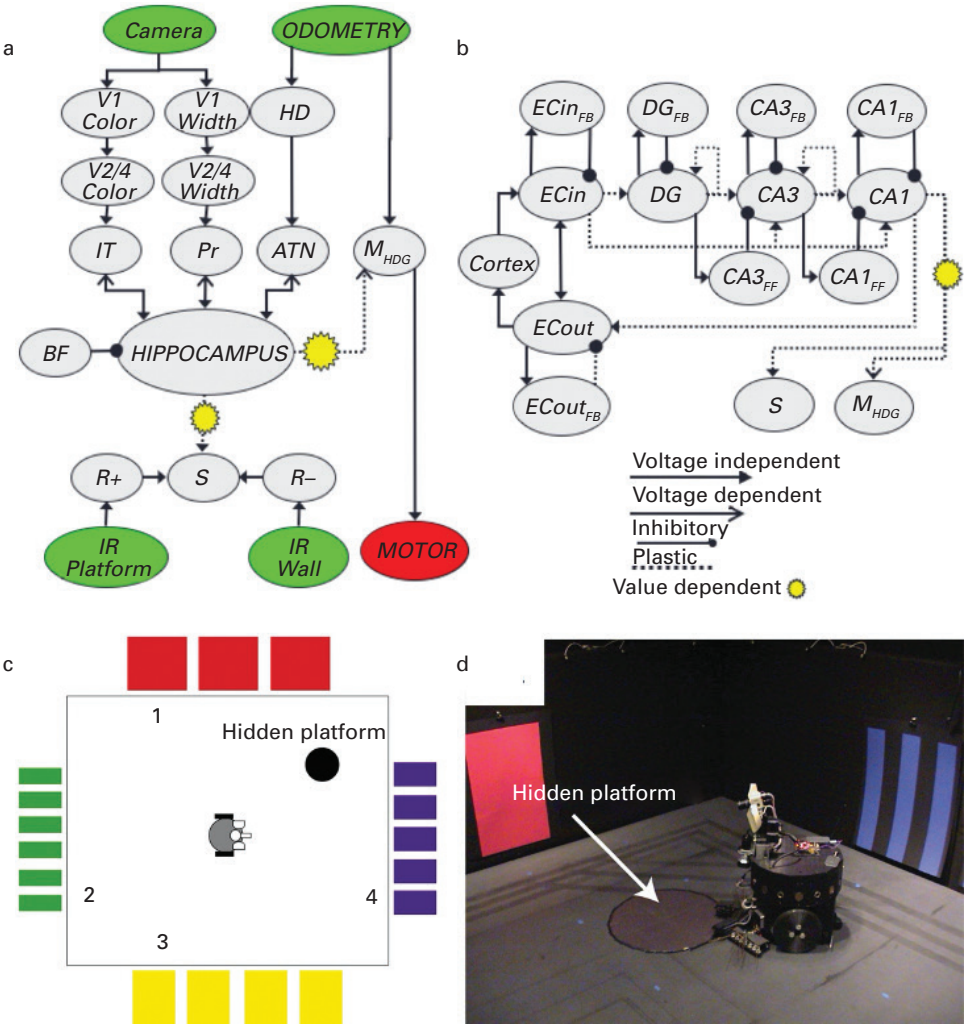

Figure 2.3
Darwin X and a hippocampal model of episodic memory. (a) The overall neural network architecture included neuronal groups for the visual "what" and "where" streams (V1 → V2/4 → IT, V1 → V2/4 → Pr, respectively), head direction system (HD), reward system (R+, R–, S), and hippocampus. (b) Subfields within the hippocampus neural group. Arrows denote synaptic projections between subgroups. (c) Schematic of a dry variant of the Morris water maze. Colors denote landmarks; numbers denote starting positions of trials. (d) Darwin X Brain-Based Device. The hidden platform was a piece of black construction paper that Darwin X could not see with its camera but could detect with a downward-facing IR sensor. Adapted with permission from Krichmar, Nitz, et al. 2005.

neurobiologically plausible neural network (Krichmar and Edelman 2002). Darwin VII autonomously explored its environment and sampled stimuli that contained positive and negative values (figure 2.4). Through its experiences, Darwin VII built up perceptual categories of the objects it sampled. Darwin VII's simulation was based on the anatomy and physiology of vertebrate nervous systems. The simulated nervous system comprised a number of areas labeled according to the analogous cortical and subcortical brain regions for vision, auditory processing, and value. Each area contained different types of neuronal units consisting of simulated local populations of neurons or neuronal groups. The simu-

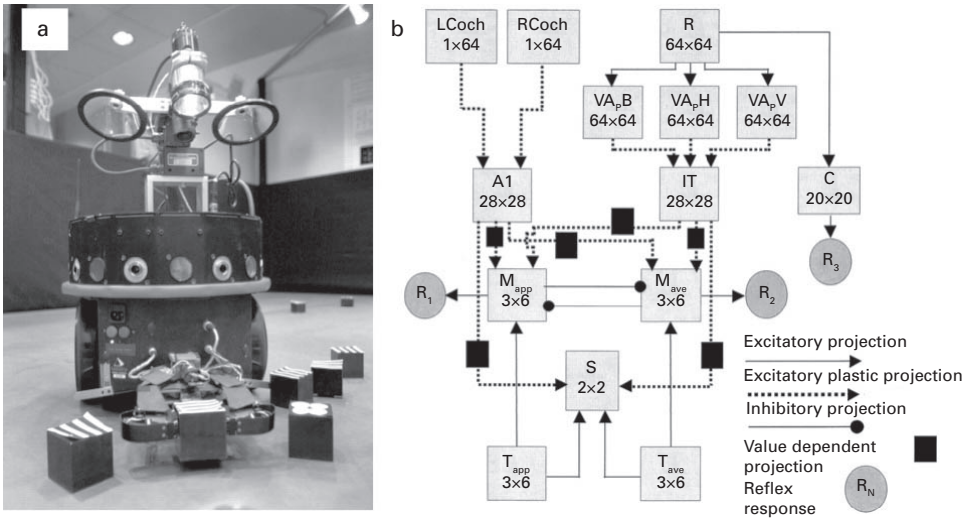

Figure 2.4
Darwin VII robot and neural network. (a) Darwin VII consists of a mobile base equipped with several sensors and effectors. Darwin VII is constructed on a circular platform with wheels that permit independent translational and rotational motion, with pan and tilt movement for its camera and 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. (b) 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. The 64×64 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 VA_P. Three subpartitions in VA_P were selective for blob-like features, short horizontal line segments, or short vertical line segments. Responses within VA_P closely followed stimulus onset and projected nontopographically 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. A1 and IT contained local excitatory and inhibitory interactions producing firing patterns 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. 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. The taste system sent information to the motor areas (M_{app} and M_{ave}) and the value system (S). Area S projected 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"). *Source:* Adapted with permission from Krichmar and Edelman 2002.

lated nervous system contained 18 neuronal areas, 19,556 neuronal units, and approximately 450,000 synaptic connections. Figure 2.4*b* shows a high-level diagram of the different neural areas and the synaptic connections between neural areas in the simulated nervous system. A neuronal unit in Darwin VII was simulated with a mean firing-rate model, and the activity of such a unit corresponded roughly to the firing activity of a group of neurons averaged over a time period of 200 ms. This corresponded to the time needed to process sensory input, compute neuronal unit activities, update the connection strengths of plastic connections, and generate motor output.

The total contribution of synaptic input to unit i was given by

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

where N is the number of connections to unit i, c_{ij} is the weight value of the connection projecting to unit i from unit j, and $s_j(t)$ is the activity of unit j at time step t. Negative values for c_{ij} corresponded to inhibitory connections. The activity level of unit i was given by

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

where

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

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, which differed depending on the neural area.

Connections within and between neuronal areas were subject to activity-dependent modification following a value-independent and a value-dependent synaptic rule. Synaptic modification was determined by both pre- and postsynaptic activity and resulted in either strengthening or weakening of the synaptic efficacy between two neuronal units. The Bienenstock, Cooper, and Munro (BCM) learning rule was used to govern synaptic change because it has a region in which weakly correlated inputs are depressed, and strongly correlated inputs are potentiated (Bienenstock, Cooper, and Munro 1982).

Value-independent synaptic changes in c_{ij} were given by

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

where $s_i(t)$ and $s_j(t)$ are activities of post- and presynaptic 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 is a piecewise linear approximation of the BCM learning rule.

The synaptic change for value-dependent synaptic plasticity was given by

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

where $\overline{S}$ is the average activity of the value system S (see figure 2.4*b*).

Darwin VII's environment consisted of an enclosed area with black walls and a floor covered with opaque black plastic panels, on which metallic cubes were distributed (figure 2.4*a*). The top surfaces of the blocks were covered with black-and-white patterns: blobs and stripes. Stripes on blocks in the gripper could be viewed in either a horizontal or vertical orientation, 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, strongly conductive blocks with a striped pattern and a 3.9 kHz tone were arbitrarily chosen to be positive-value exemplars, whereas weakly conductive blocks with a blob pattern and a 3.3 kHz tone represented negative-value exemplars.

Early during the conditioning trials, Darwin VII picked up and "tasted" blocks that led to either appetitive or aversive responses (see figure 2.5*a*, *left panel*). 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 A1's response to a tone. This shift from value system activity 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, Dayan, and Montague 1997).

After associating visual patterns with taste, Darwin VII continued to pick up and "taste" stripe-patterned blocks but avoided blob-patterned blocks (see figure 2.5*a*, *left panel*). 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 figure 2.5*c*, *left panel*). The right panel of figure 2.5*b* shows the percentage of conditioned responses, which were driven by the auditory or visual stimulus, for seven Darwin VII trials. The increase in conditioned responses showed that Darwin VII learned that auditory or visual cues predicted the value of the object, which resulted in it taking the appropriate behavioral response. These learning curves closely resembled those for similar conditioning experiments in rodents, pigeons, and other organisms.

In Darwin VII, activity in the simulated inferotemporal cortex, IT, provided the basis for visual perceptual categorization. Initially, IT's responses to visual stimuli were weak and diffuse (see IT activity in figure 2.5*a*, *right panel*). After approximately five stimulus encounters, activity-dependent plasticity between primary visual cortex, VA_P, and IT caused IT responses to the different stimuli to become strong, sharp, and separable (see IT activity in figure 2.5*b*, *right panel*). 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 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. These invariant category responses developed as a result of competition among activity-dependent plastic connections between retinotopically mapped VA_P and non-topographically mapped IT.

The behavior of Darwin VII showed that a robot operating on biological principles and without prespecified instructions could carry out perceptual categorization and conditioned responses. 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. Darwin VII

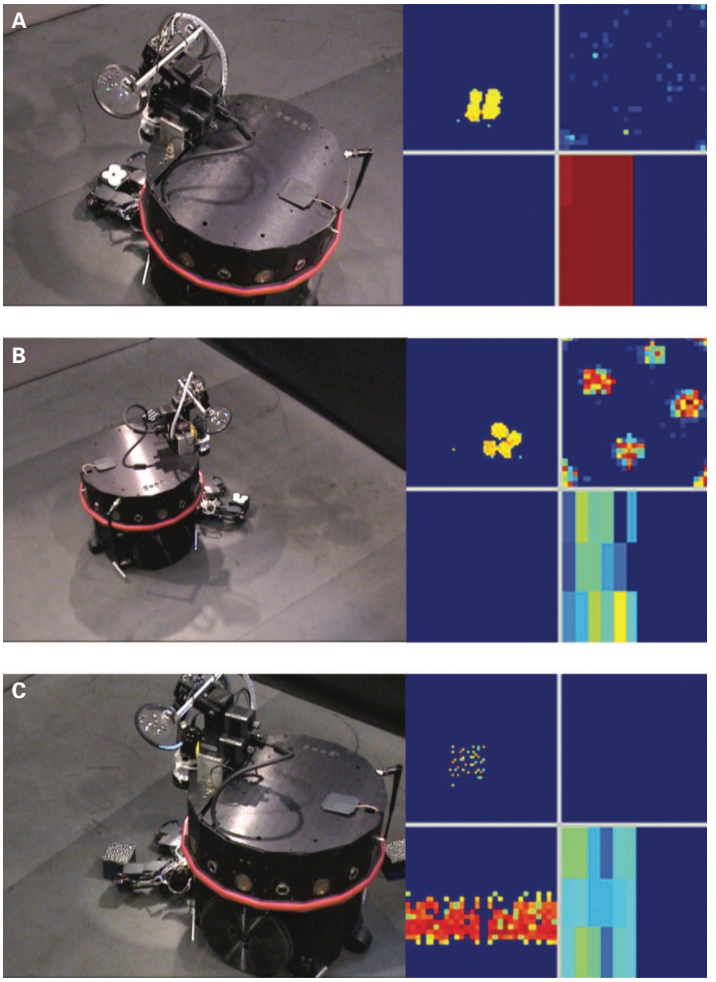

Figure 2.5
Left: Darwin VII during behavioral experiments. The panels to the right of Darwin VII show the activity of selected neural areas in the simulation (R, *top left*; IT, *top right*; A1, *bottom left*; M_{ave}, *bottom right*, *left side*; M_{app}, *bottom right*, *right side*). Each pixel in a selected neural area represents a neuronal unit, and activity is normalized in a range from no activity (*dark blue*) to maximal activity (*bright red*). (a) Darwin VII upon the first encounter with an aversive block. The stimulus block shown in this figure and in (b) had a blob-like visual pattern but did not beep. In this early conditioning trial, Darwin VII is shown picking up and "tasting" an aversive block. Activity in IT is insufficient, but activity in the taste system T_{ave} is sufficient to drive activity in the aversive motor behavior neural area (M_{ave}) above the behavioral threshold. (b) Darwin VII upon the tenth encounter with an aversive block having blob-like visual patterns. After primary conditioning with visual stimuli, activity in area IT is sufficient to drive the M_{ave} neuronal units above the behavioral threshold, triggering a motor response to avoid "tasting" an aversive block. (c) Darwin VII upon the tenth encounter with an aversive block having only auditory cues. After primary conditioning with auditory stimuli, activity in area A1 is sufficient to drive the M_{ave} neuronal units above the threshold to trigger a behavioral response. *Right:* The percentage of conditioned responses (%CR) per stimuli encountered by Darwin VII for auditory and visual stimuli. Each point is the average %CR for seven Darwin VII trials. *Source:* Adapted with permission from Krichmar and Edelman 2002.

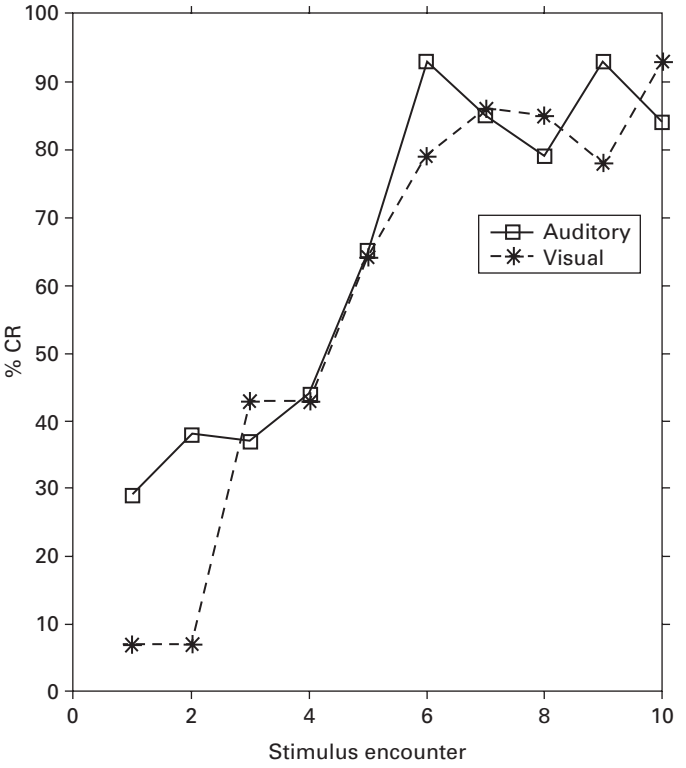

Figure 2.5
(continued)

laid down groundwork for increasingly sophisticated neurorobots with more complex neural circuits and morphologies, which gave further insights into the relationships between brain, body, and behavior.

2.3 Building a Neurorobotics Community

Over the years, a neurorobotics community has emerged in part due to workshops and special journal issues on the topic. The *IEEE Robotics and Automation Magazine* devoted an issue to the topic (Browne et al. 2009). Special sessions were occasionally held on the topic at major IEEE robotics conferences. The European Union's Human Brain Project, a large-scale research project for understanding the nervous system, included a neurorobotics division headed up by Alois Knoll and Florian Rohrbein (Falotico et al. 2017).

In 2004, a special session on "Neurorobotic Models in Neuroscience and Neuroinformatics" took place at the International Conference on the Simulation of Adaptive Behavior (Seth, Sporns, and Krichmar 2005). To introduce the session, it was stated that a neurorobotic device has the following properties: 1) It engages in a behavioral task, 2) it is situated in a structured environment, and 3) its behavior is controlled by a simulated nervous system designed to reflect, at some level, the brain's architecture and dynamics. The session included Auke Ijspeert's research on evolving neural networks for a robotic salamander (Ijspeert, Crespi, and

Cabelguen 2005; Ijspeert et al. 2007). In this research, different motor patterns (i.e., swimming or walking) emerged due to the interaction between brain and body with the specific environment (i.e., water or land). Olaf Sporns and Max Lungarella showed how embodiment can alter and improve information processing in a neural system (Lungarella et al. 2005). In addition, several papers on how the hippocampus contributes to spatial memory were presented (Arleo, Smeraldi, and Gerstner 2004; Banquet et al. 2005; Chavarriaga et al. 2005; Krichmar, Seth, et al. 2005).

Robot models of rodent navigation have made up a number of neurorobotic implementations. One reason for the interest in these models is because robot navigation is a fascinating and complex problem. Another reason is that the neural activity patterns observed in the rat are clear, interesting, and amenable to modeling. For example, a head-direction cell can be modeled with an attractor network and cosine tuning curves (Stringer et al. 2002). A hippocampal place cell can be modeled with a two-dimensional Gaussian (Foster, Morris, and Dayan 2000). The more recent finding of grid cells in the entorhinal cortex has led to a number of proposed neural models (Zilli 2012). Using attractor networks and neural elements that resemble head direction cells, place cells, and grid cells, the Australian RatSLAM team has reported results with neuro-inspired algorithms that are as good as or better than state-of-the-art localization and mapping by conventional robots (Milford et al. 2016). Although great progress has been made in the conventional robotics community with SLAM, or simultaneous localization and mapping (Kohlbrecher et al. 2011; Mur-Artal, Montiel, and Tardos 2015) and path planning (LaValle 2011a, 2011b), a number of open issues still remain when it comes to flexible navigation under dynamic conditions. Under these challenging situations, rodents show superior performance and robustness and still provide inspiration for improved robot navigation algorithms.

2.4 Neurorobotics and Neuromorphic Engineering

An important potential development for the field of neurorobotics is the reemergence of neuromorphic engineering (Indiveri et al. 2011). By reemergence, we mean that the original analog circuits developed by Carver Mead (1990) and his team in the 1980s have led to near-commercially viable computers designed by large companies such as IBM (Merolla et al. 2014) and Intel (Davies et al. 2018). Like neurorobotics, neuromorphic engineering uses inspiration from the brain to build computer architectures and sensors. Because these computers were specifically designed for asynchronous, event-driven processing, spiking neural networks that controlled neurorobots were ideal for these platforms. Moreover, neuromorphic architectures hold great promise for neurorobot applications due to their low power budget and their fast, event-driven responses. For example, the SpiNNaker neuromorphic computer from Manchester has been used in an obstacle avoidance and random exploration task (Stewart et al. 2016). In addition to running neural networks on specialized hardware, very low power neuromorphic vision and auditory sensors are being developed (Liu and Delbruck 2010). Similar to biology, these sensors only respond to change or salient events, and when they do respond, it is with a train of spikes. This allows seamless integration of these sensors with spiking neural networks, and their event-driven nature leads to power efficiency that's ideal for embedded systems, such as robots.

The development of lightweight neuromorphic chips inspired the idea that many computing processes related to outdoor navigation could be implemented on neuromorphic hardware to control ground robots. Neuromorphic hardware is especially beneficial for outdoor navigation, as the robots must rely on battery power for long periods of time and are often used in vital operations such as search and rescue. Spiking implementations of low-level perceptual navigation tasks as well as high-level planning tasks allow for navigation subtasks to run in parallel.

Working with IBM's low-power TrueNorth neuromorphic chip (Esser et al. 2016), we demonstrated that a convolutional neural network (CNN) could be trained to self-drive a robot on a mountain trail (Hwu et al. 2017). Initially, the robot was driven along the trail using remote control. The RGB camera frames, along with the corresponding action controls of steering left, steering right, and driving forward, were recorded for training the CNN. The CNN was first trained with conventional backpropagation techniques, using the RGB images as input and the set of three actions as output. The weights of this neural network were then transferred to weights in a spiking neural network of the same structure as the original CNN. This spiking network was run on the TrueNorth chip, which was powered by the same single hobby-level nickel metal hydride (NiMH) battery used to power the motors of the robot (figure 2.6). The advantage of using this pipeline was that we were able to harness well-developed techniques of CNN training while achieving order-of-magnitude gains in energy efficiency. The circuit diagram and pipeline shown in figure 2.6 could generalize to other hardware and neurorobot applications.

2.4.1 Case Study: Spiking Wavefront Propagation—Brain-Inspired Neuromorphic Path Planning

Navigation is a necessary component of most robots and animals, both of which operate under the constraints of limited time and energy. Using inspiration from brain connectivity, neuron spiking dynamics, and a recent finding that axonal conductance undergoes experience-dependent plasticity (Fields 2015), a model of spiking wavefront propagation was created (Hwu et al. 2018). The model was inspired by the role of the hippocampus in animal navigation. This includes the existence of place cells in the hippocampus, which are active according to the physical location of the animal (O'Keefe and Dostrovsky 1971). These place cells are involved in hippocampal replay, in which the place cells activate in sequence according to potential trajectory routes the animal can take (Dragoi and Tonegawa 2011; Pfeiffer and Foster 2013). Another biological observation behind spiking wavefront propagation is that spreading waves of activity can be found across several areas of the brain including the hippocampus, supporting brain connectivity and memory (Zhang and Jacobs 2015).

Combining these observations, the model of spiking wavefront propagation is able to plan paths through a grid representation of space. Each grid unit corresponds to a discretized area of physical space, and connections between units represent the ability to travel from one area to a neighboring area. Each unit in the grid represents a single neuron with spiking dynamics. The membrane potential of neuron i at time $t+1$ is represented by

$$v_i(t+1) = u_i(t) + I_i(t),$$

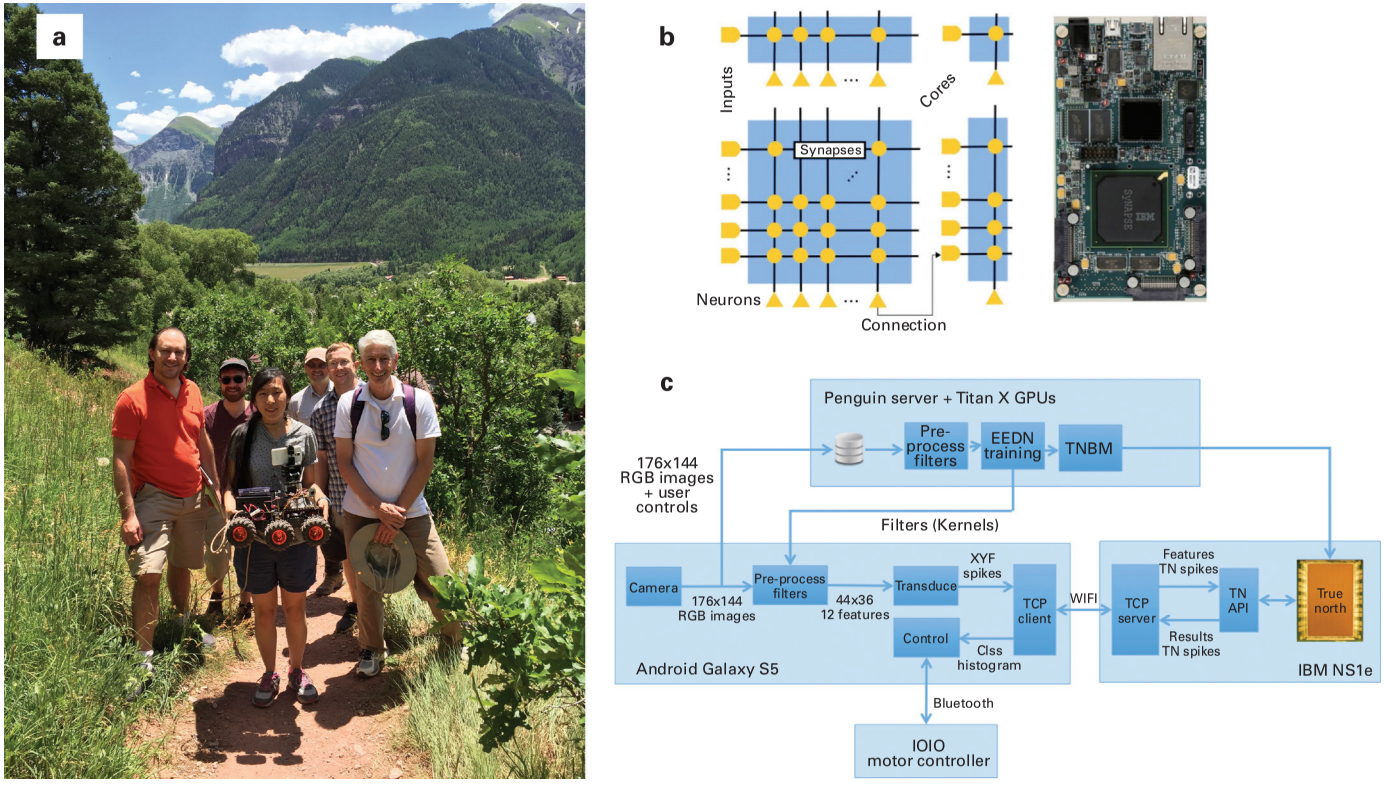

Figure 2.6
A self-driving robot using deep convolutional neural networks on IBM's TrueNorth neuromorphic hardware. (a) The photograph was taken in Telluride, Colorado, where the robot autonomously traversed mountain trails. *Left to right:* Rodrigo Alvarez-Icaza (IBM), Jacob Isbell (University of Maryland), Tiffany Hwu (University of California, Irvine), Will Browne (Victoria University of Wellington), Andrew Cassidy (IBM), and Jeff Krichmar (University of California, Irvine). Missing from the photograph is Nicolas Oros (BrainChip). (b) *Left:* The connectivity of the IBM TrueNorth neuromorphic chip. *Right:* An image of the IBM TrueNorth NS1e board used in the experiments. (c) Data pipeline for running the self-driving robot. Training was done separately with the Eedn MatConvNet package using Titan X GPUs. During testing, a Wi-Fi connection between the Android Galaxy S5 and IBM NS1e transmitted spiking data back and forth, using the TrueNorth (TN) runtime API. *Source:* Adapted with permission from Hwu et al. 2018.

in which $u_i(t)$ is the recovery variable, and $I_i(t)$ is the input current at time t. The recovery variable $u_i(t+1)$ is modeled as

$$u_i(t+1) = (-5 \text{ if } v_i(t) = 1;\ \min(u_i(t)+1,\ 0) \text{ otherwise}),$$

such that if it starts as a negative value, it increases at a steady rate toward a baseline value of 0. The input current $I_i(t+1)$ is represented as

$$I_i(t+1) = \sum_j (1 \text{ if } d_{ij}(t) = 1; 0 \textit{ otherwise}),$$

such that $d_{ij}(t)$ is the delay counter of the signal from neighboring neuron j to neuron i. The delay $d_{ij}(t+1)$ is calculated as

$$d_{ij}(t+1) = (D_{ij}(t) \text{ if } v_j(t) = 1;\ \max(d_{ij}(t)-1,\ 0) \text{ otherwise}),$$

such that it behaves as a timer corresponding to axonal delay with a starting value of $D_{ij}(t)$. This starting value of $D_{ij}(t)$ is a delay value depending on the cost of traversing the spatial area corresponding to the neuron. Taken together, these equations describe the simplified dynamics of a spiking neuron. When a spike from a neighboring neuron occurs, the input current I_i is set to 1, causing a spike. Immediately after, the recovery variable u_i is set to −5, which then counts up by 1 at each successive time step and stops at 0. This mechanism models the refractory period of the neuron. Next, all delay counters d_{ij} for all neighbor neurons j are set to their assigned starting values of D_{ij}.

Multiple possibilities exist for encoding the values D_{ij}. These values should encode the cost of traversing from one area to another. This may be the energy required, the potential risks, or the physical wear. For instance, traveling through rough terrain would be riskier and require more energy for ground robots and therefore have higher costs. A cost map of the same dimensions as the grid can transfer to values of D_{ij} in a one-to-one fashion. The cost map, if known in advance, can be used to populate delay values of the grid prior to running spiking wavefront propagation. They may also be learned on the fly while exploring the terrain. In neuroscience, this would correlate to axonal plasticity, in which the myelin sheath of a neuron consisting of white matter grows in volume with heightened activity and subsequently increases the speed of signals traveling from one neuron to another (Fields 2015). As an agent travels through its environment, either randomly or by intentionally navigating, D_{ij} values are updated each time the agent enters a new grid area using the following equation:

$$D_{ij}(t+1) = D_{ij}(t) + \delta(map_{xy} - D_{ij}(t)),$$

where δ is the learning rate, and map_{xy} is a sample of the cost as the agent traversed location coordinates (x, y) corresponding with grid neuron i. The update rule is applied for each of the j neighbors of neuron i. The advantages of axonal plasticity are that the agent can learn while operating, continuously gaining new information. With a small learning rate, the model accounts for noise in the environment such that if the agent samples a faulty cost value due to sensor error or environmental factors, the effect is averaged across multiple trials. However, learning accurate cost values for an entire grid may require many trials, as each grid area must be traversed several times. It may therefore be preferable to start with an a priori map of costs, updating with sensor-based observations as they occur.

To perform path planning using the grid encoded with costs, an input current is added to the neuron corresponding to the location of the agent to induce a spike. This induces spikes in neighboring neurons, subsequently starting a traveling wave across the entire grid. As the spikes occur, their spikes are recorded using address event representation (AER), which includes pairs of neuron IDs and spike times. Figure 2.7 shows how using AER can be used for path planning.

To plan a path from the start location to any other location, the first spike time of the destination neuron is recorded. The ID of the destination neuron is also recorded on a list. Then the spike times of each neighboring neuron are examined, and the neuron with the most recent spike is appended to the list. The same process is repeated with this neighboring neuron, and so on, until the start neuron is added to the list. The optimal path accounting for length and cost is then returned as the reversed list of neuron IDs.

The present spiking wavefront algorithm was successfully tested on a mobile ground robot traversing over grass, dirt, and asphalt terrains (Hwu et al. 2018). The robot was created from affordable hobbyist parts and an Android phone for computation (figure 2.8, *bottom left*). The robot motors and sensors were powered by a single NiMH battery, making energy savings a priority in its operation. The robot was tested at a large outdoor park in two areas (figure 2.8, *top left*). One area was a grass field surrounded by an asphalt road. Three cost maps were created out of this area (figure 2.8, *top right*): one with a uniform low cost, one with a low cost for the surrounding road, and one with a low cost for the surrounding road and a medium cost for park benches. The other area was grassy with trees, a surrounding outer asphalt road, and a dirt trail cutting straight across. A single cost map was generated from this area, consisting of a low cost for the surrounding road, a high cost for the trees, and a medium cost for the dirt road. Using these different maps, researchers generated a path to navigate between a set of starting and end points with the spiking wavefront algorithm.

Time	Neuron ID (r,c)
1	(1,1)
2	(1,2),(2,1),(2,2)
3	(1,3),(2,3),()3,1
...	...
6	...(3,6),(5,3),...
...	...
9	(3,4),(4,2),(6,6)

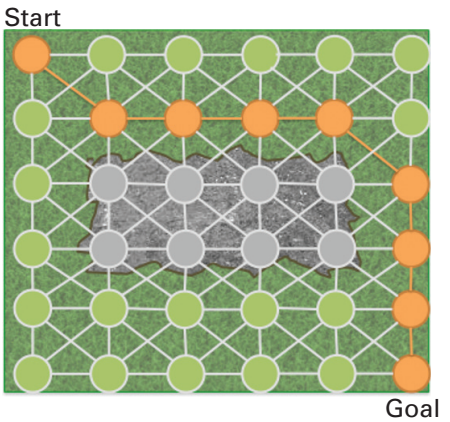

Figure 2.7
Path planning using an address event representation table. *Left:* Spike types and neuron IDs are recorded in this table. In order to plan a path using the trained grid of neurons, the neuron corresponding with the location of the agent receives an impulse spike. This spike triggers a wavefront signal to propagate across the grid surface. Since some neurons have longer axonal delays, the wavefront edge travels at different speeds. Using the table, the neuron corresponding to the goal is identified. Then, stepping back through the time steps, a path of neurons can be traced back to the start neuron (*right*). Since costs are encoded using axonal delays, the planned path avoids costlier terrains with obstacles. *Source:* Adapted with permission from Hwu et al. 2018.

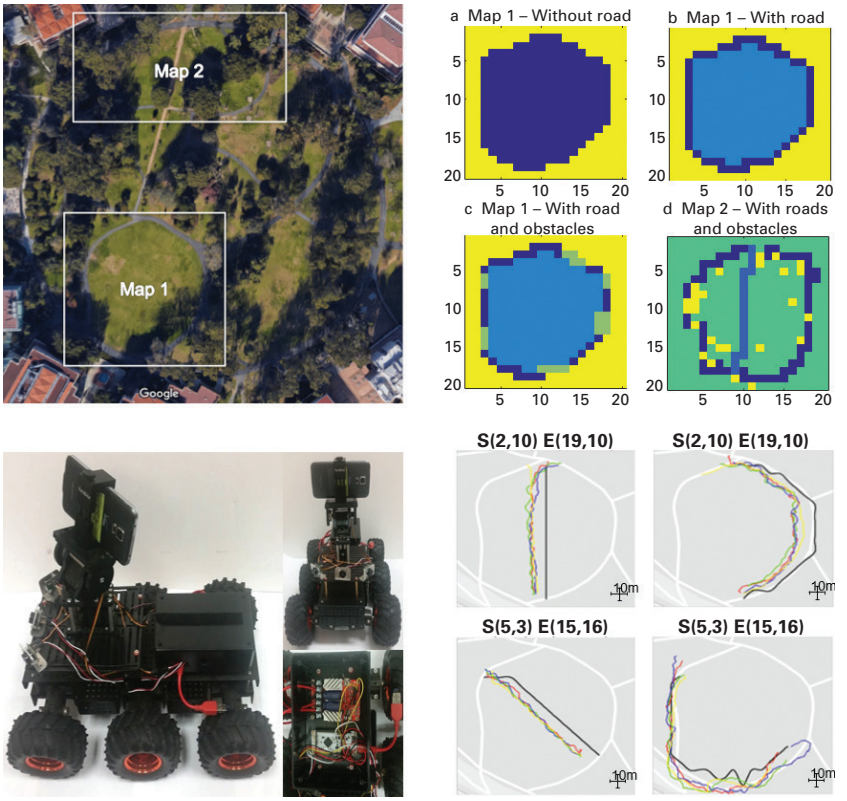

Figure 2.8
Outdoor demonstration of spiking wavefront propagation. *Top left:* A satellite image of an outdoor park where two areas were used to generate cost maps. *Top right:* (a) A uniform cost was given to the grassy area. (b) A low cost was given to the road surrounding the grassy area. (c) A low cost was given to the road, and a medium cost was given to park benches. (d) A low cost was given to the surrounding road, a high cost was given to trees, and a medium cost was given to the dirt path cutting across the area. *Bottom left:* Side, front, and interior views of the Android-based robotics platform. *Bottom Right:* The first row shows two paths planned with the same starting and ending points. The path on the left column was generated using a cost map without the outer road, and the path on the right column was generated using a cost map including the outer road. The bottom row shows the same but with a different set of starting and ending points. When the road is accounted for, the planned path takes the longer, smoother path, as opposed to the shortest path. *Source:* Adapted with permission from Hwu et al. 2018.

Waypoints along the path corresponded to neuronal units representing locations on the map. The robot then used the GPS of the Android phone to drive along the waypoints generated by the algorithm. The paths taken by the robot highlighted trade-offs between finding the shortest path and finding the smoothest path (figure 2.8, *bottom right*). When a uniform cost was used, the shortest path was always chosen. When the road was considered, the robot would occasionally take it, even if it meant traveling a longer distance. For the map containing the dirt road, the robot judged the trade-offs of taking the fastest route versus traveling over bumpy grass. The robot demonstration applied spiking wavefront propagation to cost-aware path planning, showing the possibility of energy savings on an energy-limited mobile platform.

This demonstration combined with the spiking CNN shows the potential for a complete neuromorphic computing solution to outdoor navigation (Hwu, Krichmar, and Zou 2017).

Such a system could enable more computation on mobile platforms and provide more insight into how the brain is able to function with limited energy.

2.5 Future Outlook

Built on a variety of interdisciplinary ideas, neurorobotics has grown into a rich and interesting field. Some of the subtopics of research have remained the same throughout its history, such as navigation, motor planning, mapping, and the development of neural networks. However, the research continues to develop as newer techniques in neurobiology, such as optogenetics, as well as techniques in machine learning and deep neural networks continue to add new tools and insights.

The fields of AI, machine learning, and especially artificial neural networks have enjoyed particular success in recent years. Although deep neural networks have largely been successful, there are a number of new challenges within the field. For the most part, the neural networks work well on specific tasks but have trouble extending knowledge from previously learned tasks to newer but related tasks. Moreover, the neural networks take a large amount of data and training and fail to capture many behaviors that are easy for humans (Larson 2017). This indicates that the study of the brain can contribute much to the field.

According to neuroscientist and entrepreneur Jeff Hawkins (2017), the brain has three key features required for intelligence: 1) learning by rewiring; learning in the brain is both rapid and gradual and can store representations that last over a lifetime; 2) sparse representations; under the constraints of nature, the brain stores information using the fewest metabolic resources possible; 3) embodiment; interaction between the brain and environment together is required for intelligence. We would also argue that the following features are important: 4) value systems; good and bad stimuli from the environment must be learned by detecting saliency and reacting appropriately (Friston et al. 1994; Krichmar 2008) and 5) prediction; we must be able to extrapolate from past experiences to learn how to process future experiences (Clark 2013). Applying these principles, future research in neurorobotics can potentially achieve a more holistic understanding of intelligence, striving for behavior that generalizes across multiple domains and maintains information over long time frames. Neurorobotics is a promising approach to addressing many of the issues the AI community faces today.

2.6 Conclusion

To truly understand intelligence, we believe one must study the brain and body and apply these principles to all applications. Intelligent biological systems are currently our best standard, serving as a model for what AI eventually hopes to achieve. The insights gathered from neurorobotics will ultimately lead to a strong understanding of the essence of intelligence, which will then benefit our understanding of ourselves and lead to applications that improve future technologies.

Additional Reading and Resources

- Edelman, G. M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books. This book introduces an important brain theory that was amenable to testing with neurorobotics.
- Krichmar, J. L., and H. Wagatsuma, eds. 2011. *Neuromorphic and Brain-Based Robots*. Cambridge: Cambridge University Press. This book provides a snapshot of the state of the art in neurorobotics at that time. It covers a range of topics from low-level perception to machine consciousness.
- Tani, Jun. 2016. *Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena*. Oxford: Oxford University Press. Jun Tani has been a pioneer in neurorobotics. His book covers how higher-order cognition might be realized in neurorobots.
- Neurorobotics software and designs:
  - RatSLAM: https://openslam-org.github.io/openratslam.html.
  - Android-based robotics platform: https://www.socsci.uci.edu/~jkrichma/ABR/index.html.

References

Almassy, Nikolaus, Gerald M. Edelman, and Olaf Sporns. 1998. "Behavioral Constraints in the Development of Neuronal Properties: A Cortical Model Embedded in a Real-World Device." *Cerebral Cortex* 8 (4): 346–361. https://doi.org/10.1093/cercor/8.4.346.

Arleo, Angelo, and Wulfram Gerstner. 2000. "Modeling Rodent Head-Direction Cells and Place Cells for Spatial Learning in Bio-mimetic Robotics." In Meyer et al. 2000, 236–245.

Arleo, Angelo, Fabrizio Smeraldi, and Wulfram Gerstner. 2004. "Cognitive Navigation Based on Nonuniform Gabor Space Sampling, Unsupervised Growing Networks, and Reinforcement Learning." *IEEE Transactions on Neural Networks* 15 (3): 639–652. https://doi.org/10.1109/TNN.2004.826221. https://www.ncbi.nlm.nih.gov/pubmed/15384552.

Arsenio, Artur M. 2000. "Neural Oscillator Networks for Rhythmic Control of Animats." In Meyer et al. 2000, 105–114.

Banquet, Jean-Paul, P. H. Gaussier, Mathias Quoy, Arnaud Revel, and Yves Burnod. 2005. "A Hierarchy of Associations in Hippocampo-Cortical Systems: Cognitive Maps and Navigation Strategies." *Neural Computation* 17 (6): 1339–1384. https://doi.org/10.1162/0899766053630369. https://www.ncbi.nlm.nih.gov/pubmed/15901401.

Bienenstock, Elie L., Leon N. Cooper, and Paul W. Munro. 1982. "Theory for the Development of Neuron Selectivity: Orientation Specificity and Binocular Interaction in Visual Cortex." *Journal of Neuroscience* 2 (1): 32–48.

Braitenberg, Valentino. 1986. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.

Bray, Laurence C. Jayet, Sridhar R. Anumandla, Corey M. Thibeault, Roger V. Hoang, Philip H. Goodman, Sergiu M. Dascalu, Bobby D. Bryant, and Frederick C. Harris Jr. 2012. "Real-Time Human-Robot Interaction Underlying Neurorobotic Trust and Intent Recognition." *Neural Networks* 32:130–137. https://doi.org/10.1016/j.neunet.2012.02.029.

Browne, William, Kazuhiko Kawamura, Jeffrey Krichmar, William Harwin, and Hiroaki Wagatsuma. 2009. "Cognitive Robotics: New Insights into Robot and Human Intelligence by Reverse Engineering Brain Functions [from the Guest Editors]." *IEEE Robotics and Automation Magazine* 16 (3): 17–18.

Chavarriaga, Ricardo, Thomas Strösslin, Denis Sheynikhovich, and Wulfram Gerstner. 2005. "A Computational Model of Parallel Navigation Systems in Rodents." *Neuroinformatics* 3 (3): 223–241. https://doi.org/10.1385/NI:3:3:223. https://www.ncbi.nlm.nih.gov/pubmed/16077160.

Clark, Andy. 2013. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences* 36 (3): 181–204. https://doi.org/10.1017/S0140525X12000477. https://www.ncbi.nlm.nih.gov/pubmed/23663408.

Collins, David, and Gordon Wyeth. 2000. "Utilising a Cerebellar Model for Mobile Robot Control in a Delayed Sensory Environment." In *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, edited by Jean-Arcady Meyer, Alain Berthoz, Dario Floreano, Herbert L. Roitblat, and Stewart W. Wilson. Cambridge, MA: MIT Press.

Davies, Mike, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou et al. 2018. "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning." *IEEE Micro* 38 (1): 82–99. https://doi.org/10.1109/MM.2018.112130359.

Dragoi, George, and Susumu Tonegawa. 2011. "Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies." *Nature* 469 (7330): 397–401. https://doi.org/10.1038/nature09633. https://www.ncbi.nlm.nih.gov/pubmed/21179088.

Edelman, Gerald M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.

Edelman, Gerald M. 1993. "Neural Darwinism: Selection and Reentrant Signaling in Higher Brain Function." *Neuron* 10 (2): 115–125.

Edelman, Gerald M., George N. Reeke, W. Einar Gall, Giulio Tononi, Douglas Williams, and Olaf Sporns. 1992. "Synthetic Neural Modeling Applied to a Real-World Artifact." *Proceedings of the National Academy of Sciences* 89 (15): 7267–7271. https://doi.org/10.1073/pnas.89.15.7267.

Esser, Steven K., Paul A. Merolla, John V. Arthur, Andrew S. Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J. Berg et al. 2016. "Convolutional Networks for Fast, Energy-Efficient Neuromorphic Computing." *Proceedings of the National Academy of Sciences* 113 (41): 11441–11446. https://doi.org/10.1073/pnas.1604850113. https://www.ncbi.nlm.nih.gov/pubmed/27651489.

Falotico, Egidio, Lorenzo Vannucci, Alessandro Ambrosano, Ugo Albanese, Stefan Ulbrich, Juan Camilo Vasquez Tieck, Georg Hinkel et al. 2017. "Connecting Artificial Brains to Robots in a Comprehensive Simulation Framework: The Neurorobotics Platform." *Frontiers in Neurorobotics* 11:2. https://doi.org/10.3389/fnbot.2017.00002.

Fields, R. Douglas. 2015. "A New Mechanism of Nervous System Plasticity: Activity-Dependent Myelination." *Nature Reviews Neuroscience* 16 (12): 756–767. https://doi.org/10.1038/nrn4023. https://www.ncbi.nlm.nih.gov/pubmed/26585800.

Floreano, Dario, and Laurent Keller. 2010. "Evolution of Adaptive Behavior in Robots by Means of Darwinian Selection." *PLoS Biol* 8 (1): e1000292. https://doi.org/10.1371/journal.pbio.1000292.

Foster, David J., Richard G. M. Morris, and Peter Dayan. 2000. "A Model of Hippocampally Dependent Navigation, Using the Temporal Difference Learning Rule." *Hippocampus* 10 (1): 1–16. 10.1002/(Sici)1098–1063(2000)10:1.

Friston, K. J., Guilio Tononi, G. N. Reeke Jr., Olaf Sporns, and Gerald M. Edelman. 1994. "Value-Dependent Selection in the Brain: Simulation in a Synthetic Neural Model." *Neuroscience* 59 (2): 229–243. https://www.ncbi.nlm.nih.gov/pubmed/8008189.

Gonzalez, F. Montes, Tony J. Prescott, Kevin Gurney, Mark Humphries, and Peter Redgrave. 2000. "An Embodied Model of Action Selection Mechanisms in the Vertebrate Brain." In Meyer et al. 2000, 157–166.

Hawkins, J. 2017. "What Intelligent Machines Need to Learn from the Neocortex." *IEEE Spectrum* 54 (6): 35–40.

Hwu, Tiffany, Jacob Isbell, Nicolas Oros, and Jeffrey Krichmar. 2017. "A Self-Driving Robot Using Deep Convolutional Neural Networks on Neuromorphic Hardware." In *2017 International Joint Conference on Neural Networks*, 635–641. New York: IEEE.

Hwu, Tiffany, Jeffrey Krichmar, and Xinyun Zou. 2017. "A Complete Neuromorphic Solution to Outdoor Navigation and Path Planning." In *2017 IEEE International Symposium on Circuits and Systems*, 1–4. New York: IEEE.

Hwu, Tiffany, Alexander Y. Wang, Nicolas Oros, and Jeffrey L. Krichmar. 2018. "Adaptive Robot Path Planning Using a Spiking Neuron Algorithm with Axonal Delays." *IEEE Transactions on Cognitive and Developmental Systems* 10 (2): 126–137. https://doi.org/10.1109/TCDS.2017.2655539. https://ieeexplore.ieee.org/document/7827038/.

Ijspeert, Auke Jan, Alessandro Crespi, and Jean-Marie Cabelguen. 2005. "Simulation and Robotics Studies of Salamander Locomotion." *Neuroinformatics* 3 (3): 171–195. https://doi.org/10.1385/NI:3:3:171.

Ijspeert, Auke Jan, Alessandro Crespi, Dimitri Ryczko, and Jean-Marie Cabelguen. 2007. "From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model." *Science* 315 (5817): 1416–1420. https://doi.org/10.1126/science.1138353. https://www.ncbi.nlm.nih.gov/pubmed/17347441.

Indiveri, Giacomo, Bernabé Linares-Barranco, Tara Julia Hamilton, André Van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih-Chii Liu et al. 2011. "Neuromorphic Silicon Neuron Circuits." *Frontiers in Neuroscience* 5:73. https://doi.org/10.3389/fnins.2011.00073. https://www.ncbi.nlm.nih.gov/pubmed/21747754.

Kohlbrecher, Stefan, Oskar Von Stryk, Johannes Meyer, and Uwe Klingauf. 2011. "A Flexible and Scalable Slam System with Full 3D Motion Estimation." In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics*, 155–160. New York: IEEE.

Krichmar, Jeffrey L. 2008. "The Neuromodulatory System: A Framework for Survival and Adaptive Behavior in a Challenging World." *Adaptive Behavior* 16 (6): 385–399.

Krichmar, Jeffrey L., and Gerald M. Edelman. 2002. "Machine Psychology: Autonomous Behavior, Perceptual Categorization and Conditioning in a Brain-Based Device." *Cerebral Cortex* 12 (8): 818–830. https://doi.org/10.1093/cercor/12.8.818.

Krichmar, Jeffrey L., Douglas A. Nitz, Joseph A. Gally, and Gerald M. Edelman. 2005. "Characterizing Functional Hippocampal Pathways in a Brain-Based Device as It Solves a Spatial Memory Task." *Proceedings of the National Academy of Sciences* 102 (6): 2111–2116. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15684078.

Krichmar, Jeffrey L., Anil K. Seth, Douglas A. Nitz, Jason G. Fleischer, and Gerald M. Edelman. 2005. "Spatial Navigation and Causal Analysis in a Brain-Based Device Modeling Cortical-Hippocampal Interactions." *Neuroinformatics* 3 (3): 197–221. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=16077159.

Krichmar, Jeffrey L., James A. Snook, Gerald M. Edelman, and Olaf Sporns. 2000. "Experience-Dependent Perceptual Categorization in a Behaving Real-World Device." In Meyer et al. 2000, 41–50.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. "Imagenet Classification with Deep Convolutional Neural Networks." *Advances in Neural Information Processing Systems* 25:1097–1105. https://doi.org/10.1145/3065386.

Larson, Erik J. 2017. "The Limits of Modern AI: A Story." https://thebestschools.org/magazine/limits-of-modern-ai/.

LaValle, S. M. 2011a. "Motion Planning Part I: The Essentials." *IEEE Robotics and Automation Magazine* 18 (1): 79–89. https://doi.org/10.1109/Mra.2010.940155.

LaValle, S. M. 2011b. "Motion Planning Part II: Wild Frontiers." *IEEE Robotics and Automation Magazine* 18 (2): 108–118. https://doi.org/10.1109/Mra.2011.941635.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–444. https://doi.org/10.1038/nature14539.

Liu, Shih-Chii, and Tobi Delbruck. 2010. "Neuromorphic Sensory Systems." *Current Opinion in Neurobiology* 20 (3): 288–295. https://doi.org/10.1016/j.conb.2010.03.007. https://www.ncbi.nlm.nih.gov/pubmed/20493680.

Lungarella, Max, Teresa Pegors, Daniel Bulwinkle, and Olaf Sporns. 2005. "Methods for Quantifying the Informational Structure of Sensory and Motor Data." *Neuroinformatics* 3 (3): 243–262. https://doi.org/10.1385/NI:3:3:243. https://www.ncbi.nlm.nih.gov/pubmed/16077161.

Mead, Carver. 1990. "Neuromorphic Electronic Systems." *Proceedings of the IEEE* 78 (10): 1629–1636. 10.1109/5.58356.

Merolla, Paul A., John V. Arthur, Rodrigo Alvarez-Icaza, Andrew S. Cassidy, Jun Sawada, Filipp Akopyan, Bryan L. Jackson et al. 2014. "A Million Spiking-Neuron Integrated Circuit with a Scalable Communication Network and Interface." *Science* 345 (6197): 668–673. https://doi.org/10.1126/science.1254642. https://www.ncbi.nlm.nih.gov/pubmed/25104385.

Meyer, Jean-Arcady, Alain Berthoz, Dario Floreano, Herbert L. Roitblat, and Stewart W. Wilson, eds. 2000. *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.

Milford, Michael, Adam Jacobson, Zetao Chen, and Gordon Wyeth. 2016. "RatSLAM: Using Models of Rodent Hippocampus for Robot Navigation and Beyond." In *Robotics Research: The 16th International Symposium ISRR*, 467–485. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-28872-7_27.

Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. 2015. "Human-Level Control through Deep Reinforcement Learning." *Nature* 518 (7540): 529–533. https://doi.org/10.1038/nature14236.

Mur-Artal, Raul, Jose Maria Martinez Montiel, and Juan D. Tardos. 2015. "ORB-SLAM: A Versatile and Accurate Monocular SLAM System." *IEEE Transactions on Robotics* 31 (5): 1147–1163. https://doi.org/10.1109/Tro.2015.2463671.

Nolfi, Stefano, and D. Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.

O'Keefe, John, and Jonathan Dostrovsky. 1971. "The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat." *Brain Research*. https://doi.org/10.1016/0006-8993(71)90358-1. https://www.ncbi.nlm.nih.gov/pubmed/5124915.

Pearson, Martin J., Ben Mitchinson, J. Charles Sullivan, Anthony G. Pipe, and Tony J. Prescott. 2011. "Biomimetic Vibrissal Sensing for Robots." *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1581): 3085–3096. https://doi.org/10.1098/rstb.2011.0164.

Pfeifer, R., and J. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.

Pfeiffer, Brad E., and David J. Foster. 2013. "Hippocampal Place-Cell Sequences Depict Future Paths to Remembered Goals." *Nature* 497 (7447): 74–79. https://doi.org/10.1038/nature12112. https://www.ncbi.nlm.nih.gov/pubmed/23594744.

Reeke, George N., Olaf Sporns, and Gerald M. Edelman. 1990. "Synthetic Neural Modeling: The 'Darwin' Series of Recognition Automata." *Proceedings of the IEEE* 78 (9): 1498–1530. https://doi.org/10.1109/5.58327.

Schultz, Wolfram, Peter Dayan, and P. Read Montague. 1997. "A Neural Substrate of Prediction and Reward." *Science* 275 (5306): 1593–1599. https://doi.org/10.1126/science.275.5306.1593. https://www.ncbi.nlm.nih.gov/pubmed/9054347.

Seth, Anil K., Jeffrey L. McKinstry, Gerald M. Edelman, and Jeffrey L. Krichmar. 2004a. "Texture Discrimination by an Autonomous Mobile Brain-Based Device with Whiskers." In Vol. 5, *Proceedings of the IEEE International Conference on Robotics and Automation*, 4925–4930. New York: IEEE.

Seth, Anil K., Jeffrey L. McKinstry, Gerald M. Edelman, and Jeffrey L. Krichmar. 2004b. "Visual Binding through Reentrant Connectivity and Dynamic Synchronization in a Brain-Based Device." *Cerebral Cortex* 14 (11): 1185–1199.

Seth, Anil K., Olaf Sporns, and Jeffrey L. Krichmar. 2005. "Neurorobotic Models in Neuroscience and Neuroinformatics." *Neuroinformatics* 3:167–170.

Stewart, Terrence C., Ashley Kleinhans, Andrew Mundy, and Jörg Conradt. 2016. "Serendipitous Offline Learning in a Neuromorphic Robot." *Frontiers in Neurorobotics* 10:1. S0954–898x(02)36091–3.

Stringer, S. M., T. P. Trappenberg, E. T. Rolls, and I. E. de Araujo. 2002. "Self-Organizing Continuous Attractor Networks and Path Integration: One-Dimensional Models of Head Direction Cells." *Network* 13 (2): 217–242. PMID: 12061421.

Zhang, Honghui, and Joshua Jacobs. 2015. "Traveling Theta Waves in the Human Hippocampus." *Journal of Neuroscience* 35 (36): 12477–12487. https://doi.org/10.1523/JNEUROSCI.5102-14.2015. https://www.ncbi.nlm.nih.gov/pubmed/26354915.

Zilli, Eric A. 2012. "Models of Grid Cell Spatial Firing Published 2005–2011." *Frontiers in Neural Circuits* 6:16. https://doi.org/ARTN.