Journal of Integrative Neuroscience, Vol. 6, No. 3 (2007) 403–431

Imperial College Press
www.icpress.co.uk

Special Report

SENSORY INTEGRATION AND REMAPPING IN A MODEL OF THE MEDIAL TEMPORAL LOBE DURING MAZE NAVIGATION BY A BRAIN-BASED DEVICE

JASON G. FLEISCHER* and JEFFREY L. KRICHMAR†

The Neurosciences Institute
10640 John Jay Hopkins Dr., San Diego
California, 92121, USA
**fleischer@nsi.edu*
†*krichmar@nsi.edu*

Received 25 March 2007
Accepted 29 July 2007

Information from many different sensory modalities converges on the medial temporal lobe in the mammalian brain, an area that is known to be involved in the formation of episodic memories. Neurons in this region, called place cells, display location-correlated activity. Because it is not feasible to record all neurons using current electrophysiological techniques, it is difficult to address the mechanisms by which different sensory modalities are combined to form place field activity. To address this limitation, this paper presents an embodied neural simulation of the medial temporal lobe and other cortical structures, in which all aspects of the model can be examined during a maze navigation task. The neural simulation has realistic neuroanatomical connectivity. It uses a rate code model where a single neuronal unit represents the local field potential of a pool of neurons. The dynamics of these neuronal units are based on measured neurophysiological parameters. The model is embodied in a mobile device with multiple sensory modalities. Neural activity and behavior are analyzed both in the normal condition and after sensory lesions. Place field activity arose in the model through plasticity, and it continued even when one or more sensory modalities were lesioned. An analysis that traced through all neural circuits in the model revealed that many different pathways led to the same place activity, i.e., these pathways were degenerate. After sensory lesions, the pathways leading to place activity had even greater degeneracy, but more of this variance occurred in entorhinal cortex and sensory areas than in hippocampus. This model predicts that when examining neurons causing place activity in rodents, hippocampal neurons are more likely than entorhinal or sensory neurons to maintain involvement in the circuit after sensory deprivation.

Keywords: Hippocampus; lesion; multimodal; place field remapping; backtrace; degenerate neural pathways.

*Corresponding author.

1. Introduction

The medial temporal lobe, including the hippocampus, is necessary for the acquisition of multimodal memories in vertebrates. Multimodal memory responses have been shown in rodent hippocampus, where pyramidal cells display place-correlated firing fields [38]. These place cells can have firing fields that are related to visual cues, olfactory cues, tactile cues, auditory cues, or various combinations thereof [37]. Many place cells maintain their firing fields even when the animals are deprived of some sensory modalities [41, 43]. However, some cells do change the location and firing rate of their place fields when sensory cues are modified, i.e. they remap. They do so in a fashion that suggests the remapped firing fields are independent of the original ones [40, 17]. These findings indicate that place field activity, like episodic memory, is multimodal, associative, and capable of pattern completion even when some input is missing.

It is likely that the unique anatomy and connectivity of the hippocampal region and its surrounding areas are critical for forming multimodal memories capable of pattern completion. Highly processed neocortical information from many sensory modalities converges onto the medial temporal lobe [28]. After several levels of further processing within the medial temporal lobe, and specifically the hippocampus, information diverges in broad projections back to the neocortex [28, 51]. Within the hippocampus itself, there are several levels of looping over different time-scales [2, 8, 50, 52]. The trisynaptic loop runs from entorhinal cortex to dentate gyrus and then in turn to the CA3 and CA1 subfields of hippocampus before returning to entorhinal cortex. The perforant path projects to dentate gyrus but also to CA3 and CA1, allowing various shortcuts of the trisynaptic loop. A possible function of this unique anatomy is that the convergence of sensory projections onto this area provides the multisensory information needed to form reliable associative memories. Moreover, the looping of information and relatively sparse connectivity within the hippocampus may allow it to integrate sensory input over time [26].

It is presently not feasible to simultaneously trace all neural circuits that contribute to place field activity using current recording techniques. To address this limitation, we have constructed brain-based devices (BBDs) that allow us to examine fully a simulated nervous system at all levels while the device carries out a behavioral task [23–26, 30]. A BBD consists of a robotic platform which is controlled by a large-scale neural simulation that is based on features of vertebrate neuroanatomy and neurophysiology. By studying a BBD behaving in a real environment, we are able to investigate all of the detailed dynamic interactions of its nervous system with the body and the environment.

Recently, we have used brain-based devices (Darwin X and XI) to investigate the formation of place activity in the hippocampus, both in an open field task similar to the Morris water maze [26], and during plus-maze navigation where we observed the formation of prospective and retrospective place fields [14]. In these studies, we observed place field activity and followed the detailed dynamics of the pathways that

gave rise to such activity using a technique called backtrace analysis [25]. Backtrace analyses have revealed that hippocampal place fields were driven mainly by the trisynaptic loop in early learning trials but increasingly by perforant path connections in late trials, and that the trisynaptic loop is more involved in the formation of the prospective and retrospective fields than in the emergence of context-independent fields.

In the present paper, we analyze the behavioral performance and neuronal activity of Darwin XI during lesion studies while navigating a plus-maze. We found that selectively removing sensory input streams to the simulated medial temporal lobe during navigation had no effect on the majority of place fields. Because the computational design of a BBD allowed us to determine the detailed microcircuitry and synaptic activity of all neuronal units, we were able to trace structural and functional connections in complete detail using backtrace analysis. This analysis revealed that many degenerate, i.e. functionally equivalent but morphologically separate pathways led to the same hippocampal place activity. Moreover, the degeneracy was more pronounced in the simulated entorhinal and sensory cortical areas than in the hippocampus. Lesions of sensory inputs to the model resulted in place unit backtraces that had a greater difference in entorhinal cortex units than hippocampus units when compared to backtraces of place activity without lesions; a result which is consistent with pattern completion theories of hippocampal function [33, 39]. These results lead us to predict that, when examining neurons causing place activity in rodents, hippocampal neurons are more likely than entorhinal or sensory neurons to maintain involvement in the circuit after sensory deprivation.

2. Methods

2.1. *Device, task and environment*

Darwin XI's task was to navigate the plus-maze shown in Fig. 1. The maze is similar to one used in studies of rodent hippocampal place activity [13]. Darwin XI explored the maze autonomously and made a choice of direction at the intersection of the maze based on the activity in its simulated nervous system. On each trial the device would start, alternately, on the East or West arm of the maze and would, with a correct behavioral choice, find the rewarding platform at the end of the South maze arm.

The plus-maze was five meters on the East-West axis and four meters on the North-South axis. Textures formed from pegs arranged in vertical and horizontal patterns were present on the arms of the maze at whisker height. All walls of the maze were black. Visible cue cards hanging on the perimeter of the maze were colored red, yellow, green, and blue. Different colored cue cards had different stripe widths. At the end of the goal arm was a platform that could not be detected by the camera, but had a distinctive reflectivity that could be detected by a downward pointing infrared (IR) transceiver. This platform was used to trigger Darwin XI's

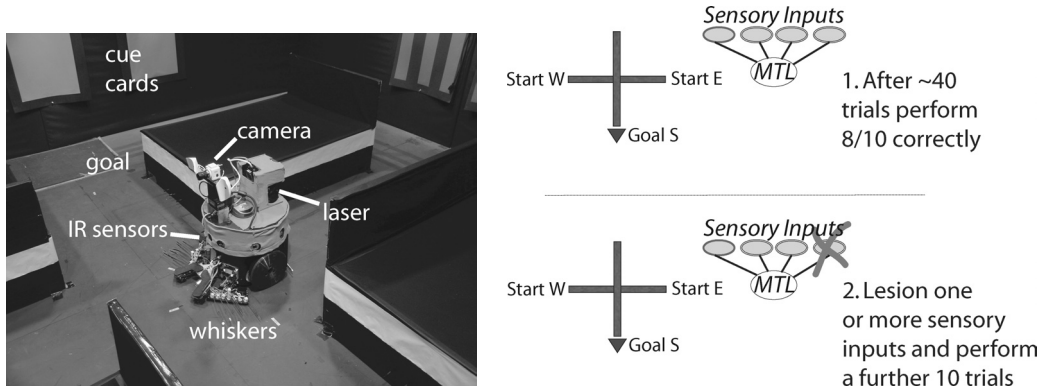

(A) (B)

Fig. 1. Darwin XI's environment and experimental protocol. (A) The device at the choice point of its plus-maze environment. Darwin XI began a trial alternately at the east or west start arm, and used its artificial whiskers to follow the maze arm until it reached the choice point. As it followed the maze wall, its whiskers sensed patterns of pegs, its camera sensed color cue cards on the perimeter, its compass provided heading, and its laser provided range information. (B) Darwin XI was first trained to find a rewarding stimulus at the end of the South goal arm (B1). After it successfully learned to choose the South arm, one or more sensory streams entering the simulated entorhinal cortex were lesioned and the device was tested for a further 10 trials (B2). Adapted from Fleischer *et al.* [14].

reward system which drove learning in value-dependent synapses affecting motor action choice at the intersection (see Sec. 2.2 and App. A for details).

The robotic device performing this task was 0.35 m in diameter and height and moved at approximately five cm/s. A CCD color camera mounted on a pan-tilt unit provided visual input to the neural simulation. Deflection-sensing artificial whiskers were arranged on each side of the device in two groups. Whiskers arranged in a row parallel to the ground were used for wall-following, and whiskers stacked vertically were used to detect texture patterns of pegs set at different heights along the wall. A magnetic compass provided heading input to the neural simulation. A rear-facing laser rangefinder provided distance-to-object information. A set of three IR transceivers, mounted front left, front middle, and front right, were used to detect the maze intersection by the absence of walls.

2.2. *Neural simulation and behavior*

Darwin XI's simulated nervous system incorporated a model of the entorhinal cortex, the hippocampus, and sensory cortical regions. The gross architecture (Fig. 2) was based on known neuroanatomical and neurophysiological parameters obtained from the literature. The simulation contained 50 neural areas (see Table A.1), and approximately 80 000 neuronal units with 1.2 million synapses. Simulated areas are denoted by italics (e.g., *CA1*). Each neuronal unit was described by a mean firing

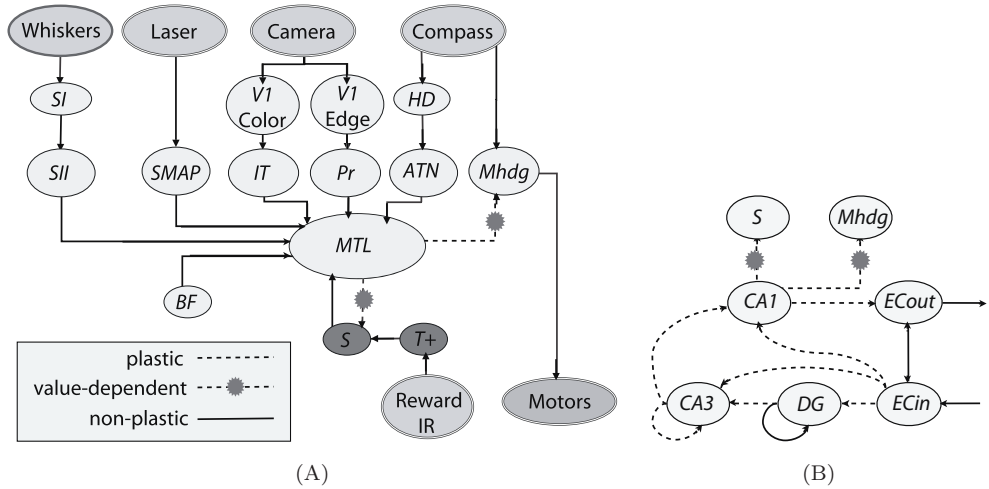

Fig. 2. Simplified schematic of Darwin XI's neural simulation. (A) The schematic shows the organization of the simulation. Simulated neural areas: *SI,SII* somatosensory (thalamic whisker areas not shown), *SMAP* population-coded laser localization, *V1* early visual (several sub areas not shown), *IT* inferotemporal, *PR* parietal, *HD* head direction, *ATN* anterior thalamic nucleus, M_{HDG} motor, *BF* basal forebrain, *S* value system, *T+* rewarding stimulus, *MTL* medial temporal lobe, including hippocampus. (B) The schematic shows the details of the *MTL*. Projections from sensory areas converge on the entorhinal cortex input layer, *ECin*, and *ECin* projects to the *DG*, *CA3* and *CA1* fields of the hippocampus. *CA1* projects to *ECout* which in turn projects to cortical sensory areas. Learning in the hippocampus occurs at synapses that have experience-dependent plasticity (dotted lines in). Value-dependent synapses from *CA1* influence behavioral choices (dotted lines with markers). Inhibitory neural areas and connections are omitted from figure for clarity. Adapted from Fleischer *et al.* [14].

rate model in which the activity of each unit corresponded to the firing rate of 100 neurons over 200 ms.

The overall model, shown in Fig. 2, included dorsal and ventral visual pathways from camera input (*V1-Color*→*IT* and *V1-Edge*→*PR*); head direction input from an internal compass (*HD*→*ATN*), whisker texture patterns on the walls which created whisker barrel activity (*SI*→*SII*), and a pseudo-cortical area (*SMAP*) which contained an estimate of location in the environment obtained through a probabilistic localization algorithm that integrates laser rangefinder and self-motion information over time. The *SMAP* area is the only sensory input which is not directly analogous to an area of the vertebrate brain, but such information can be seen as similar to what a path integration system, long hypothesized to be present in animals [6] would generate.

Inputs from all these sensory streams converged sparsely on the input layers of entorhinal cortex, such that on average, each neuronal unit in *ATN, IT, PR, SII*, and *SMAP* projected to only one *ECin* unit, and that each neuronal unit in *ECin* received an average of three projections from various sensory input stream units, which in turn projected via the perforant path to the dentate gyrus and the

CA1→*CA3* subfields of hippocampus (*ECin*→*DG*, *ECin*→*CA3*, and *ECin*→*CA1*). The model also contained the trisynaptic loop through hippocampus and back to the output layers of entorhinal cortex (*ECin*→*DG*→*CA*3→*CA1*→*ECout*).

Strengths of synapses within the hippocampus were activity-dependent, and synapses from *CA1* to a motor cortical area, M_{HDG}, were modified by a temporal difference reinforcement learning rule based on the activity of a simulated dopaminergic value system, *S*. Connections from *CA1* to *S* were also modified by the same temporal difference rule. When the device detected the reflective floor plate at the end of the goal arm, *T+* was activated, which in turn activated the value system *S* producing unconditional reward.

The device had two built-in behaviors. When in a maze arm, it used its whiskers to follow the wall as described in Ref. [44]. When the three front-mounted IRs detected the maze intersection, the device chose which arm to enter based on the neural activity of a motor cortical area M_{HDG}. In the intersection, the device stopped and panned its camera 90° to each side. During this time, the neuronal activity in M_{HDG} was summed for each direction, and the new direction was chosen with a softmax probability, $\beta = 40$, proportional to the M_{HDG} activity in that direction.

A full algorithmic description of the simulation and additional details of the sensory processing and action selection can be found in App. A.

This study examines three Darwin XI subjects, each of which consisted of the same robotic device and neural simulation containing the same connectivity on the macro level. Each subject differed, however, in its microconnectivity, as determined by a random draw from synaptic probability distributions (see App. A). Each subject also differed by its individual history of experience, which altered those connections. We ran multiple subjects, as would be done with animal experiments, to account for these differences in experience and microconnectivity in the results.

A trial consisted of a trip from one of the start arms to either the North (incorrect choice) or South (correct, rewarded choice) maze arms. During a trial, a Darwin subject's neural simulation took in sensory input and modified plastic synapses in the simulated hippocampus and motor areas according to synaptic plasticity rules (see App. A). After a trial, the new synaptic values would be saved to disk and then reloaded at the beginning of the next trial. A subject was trained until it reached a behavioral criterion of 80% correct performance in a block of 10 trials. The three subjects reached that criterion at 50, 50, and 30 trials respectively, which is comparable to the number of trials required by rats to learn the same task [13]. At this point, each subject's simulated nervous system was copied and the BBD was tested for behavioral performance on a further 10 trials after the removal of one or more sensory input streams. The sensory streams were lesioned by removing the connections from the relevant sensory input area (*ATN, IT, PR, SII, ATN*, or combinations thereof) to *ECin*.

Because Darwin XI is a computational simulation rather than an animal, any number of copies of the same trained subject can be made and tested under

different conditions. In this way, the connectivity and connection strengths were identical for each subject at the beginning of the lesion trials. The three subjects were tested with lesions of the visual input ($PR \rightarrow ECin$ and $IT \rightarrow ECin$), the whisker input ($SII \rightarrow ECin$), the head direction input ($ATN \rightarrow ECin$), the laser input ($SMAP \rightarrow ECin$), and various combinations of these lesions.

2.3. *Place field analysis*

Mean firing rate maps of neuronal activity were obtained by partitioning the maze arms into a place grid consisting of 10 cm wide slices across the maze arms. The location of the device at each time step of the simulation was recorded by two overhead cameras. A rate map for a particular neuronal unit is the summation over ten maze navigation trials of the unit's activity at each location in the grid, divided by the amount of time the device spent in each grid location. A neuronal unit was said to have a place field in a rate map if there was any grid location where the mean activity was greater than 10% of the maximum possible, a threshold equivalent to a 4 Hz spiking rate for a rodent *CA1* pyramidal cell.

The correlation between two rate maps was obtained as Pearson's R, calculated on the mean firing rates of the two maps over each location in the map. Two mean firing rate maps were considered to be correlated, and included in our analysis, only if a t-test determined that $R > 0$ with $p < 0.05$.

2.4. *Backtrace analysis*

Functional pathways causing place field activity were examined using an iterative method called backtrace analysis [25]. The backtrace starts with a particular reference neuronal unit at a specific time and recursively examines all neuronal units that caused the observed activity in this reference unit, going up to six simulation cycles back in time. A neuronal unit was added to the backtrace if that presynaptic unit was active one time step back from when a postsynaptic unit, already present in the backtrace, was active. Reference units in the *CA1* area that had place activity in the normal condition (no lesion) and under multiple lesion conditions were selected for analysis.

3. Results

3.1. *Behavioral performance*

No single sensory stream was necessary for Darwin XI to successfully navigate the maze, and even most combinations of two simultaneous lesions did not produce a deficit in behavior. The results of the behavioral testing are shown in Fig. 3, which depicts the fraction of correct choices (those leading to the South goal arm) in the 10 trials of the normal condition immediately before behavioral criterion, compared to the performance after the lesioning of one or more sensory pathways in the next 10

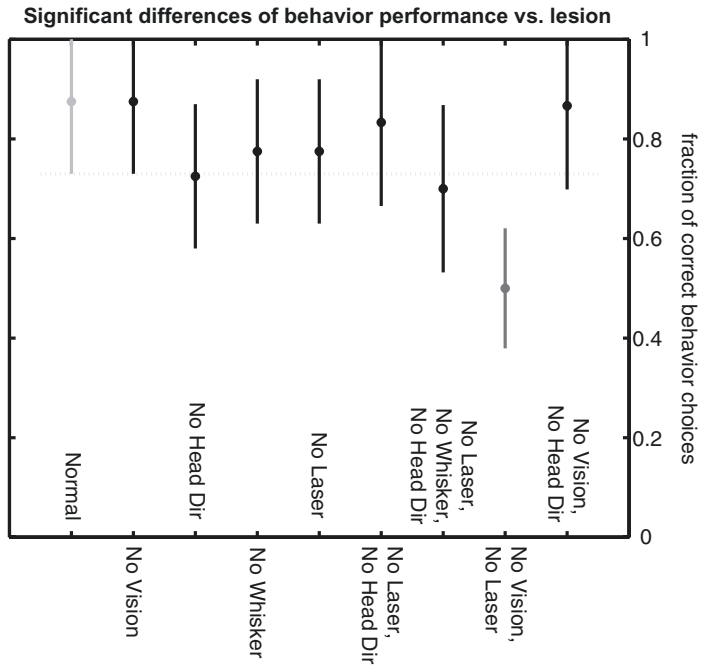

Fig. 3. Behavioral performance of Darwin XI in the plus-maze under different conditions. The graph shows the fraction of correct arm choice for three subjects under normal (no lesion) and lesion conditions. In the normal case, the ten trials before each subject reached behavioral criterion were used. In the lesion cases, the ten trials after reaching behavioral criterion were used. A Tukey-Kramer multiple comparison of means (confidence intervals on the graph, $\alpha = 0.05$) shows that only "No Vision and No Laser" case is significantly worse than the normal case.

trials. A conservative test (ANOVA, $p < 0.05$, Tukey-Kramer comparison, $\alpha = 0.05$) revealed that only when both the visual (PR and IT) and laser rangefinder ($SMAP$) inputs were lesioned was behavioral performance significantly worse than the normal (no lesion) condition. This result illustrates the robustness of the device's behavior to sensory loss, and also demonstrates that the spatial information present in *CA1* activity arises from the integration of multiple sensory streams.

3.2. *Lesion effects on place activity*

In order to understand how Darwin XI's neural place activity enabled successful behavior in spite of sensory lesions, we conducted a thorough study of place activity and the neural circuits that led to such activity. Place activity was calculated using rate maps as described in Sec. 2.3. Rate maps were calculated over the last 10 trials before behavioral criterion (normal case) or over the 10 trials after behavioral criterion (lesion cases). In the normal condition, approximately 13% of all *CA1* units (see Table 1) qualified as having place fields. These place units had a mean (standard deviation) of 1.10 (0.65) bits of spatial information.

Table 1. The effect of sensory lesions on place fields. Without any lesion, 238 neuronal units in *CA1* displayed place field activity out of a possible 1728 *CA1* neuronal units from three subjects. The table shows, for each sensory lesion, how many of those place units had a field in the same maze arm as in the normal condition (*Stable*); how many of those place units had fields that switched to a different maze arm (*Remapped*); the number of new place units appearing that had not been present in the normal condition (*Gained*); and the number of place units in the normal condition that no longer had place field activity in the lesion condition (*Lost*).

Lesion Condition	Stable	Remapped	Gained	Lost
Whisker	203	5	112	30
Laser	191	7	97	40
Head direction	172	15	101	51
Vision	102	13	69	123
Laser + head dir	154	16	108	68
Vision + whisker	73	15	74	150
Vision + laser	62	12	50	164
Vision + head dir	42	10	32	186

Table 1 shows that when sensory streams were lesioned, most place units were stable, but some lost their place field, others remapped, and some new place units emerged. Stable place units are the ones whose field is in the same maze arm for both the normal condition and the lesion condition. Units that gain or lose fields are those that have changes in the maximum mean firing rate of their rate maps such that they now reach or no longer reach the place field criterion in Sec. 2.3. No single sensory lesion caused a loss of more than 52% of the place units from the normal case. Several lesions caused the emergence of up to 43% new place units that had not been in the normal condition.

To quantify stability at a finer grain, we calculated how much or little stable place fields are shifting on the same maze arm. As described in Sec. 2.3, we calculated the correlation between rate maps in the normal and lesion conditions for each case of a place unit remaining stable after lesioning (1072 comparisons). The mean correlation was 0.764 with a standard deviation of 0.144. After inspection by eye of approximately 200 such comparisons, we can say that a place unit with a correlation score of greater than 0.7 is rarely judged to have shifted its location along the maze arm.

Examples of typical changes induced on place fields by lesions can be seen in Fig. 4. Approximately 25% of neuronal units with place fields were stable during every single sensory lesion. Figure 4(A) shows an example of such a unit. Most place units ($\approx 70\%$) lost their field during at least one sensory lesion (see for example Fig. 4(B)) but were stable during other conditions. Around 5% of units with place fields remapped during lesions, to the extent that the field switched from one maze arm to another. An example of a remapping field can be seen in Fig. 4(C).

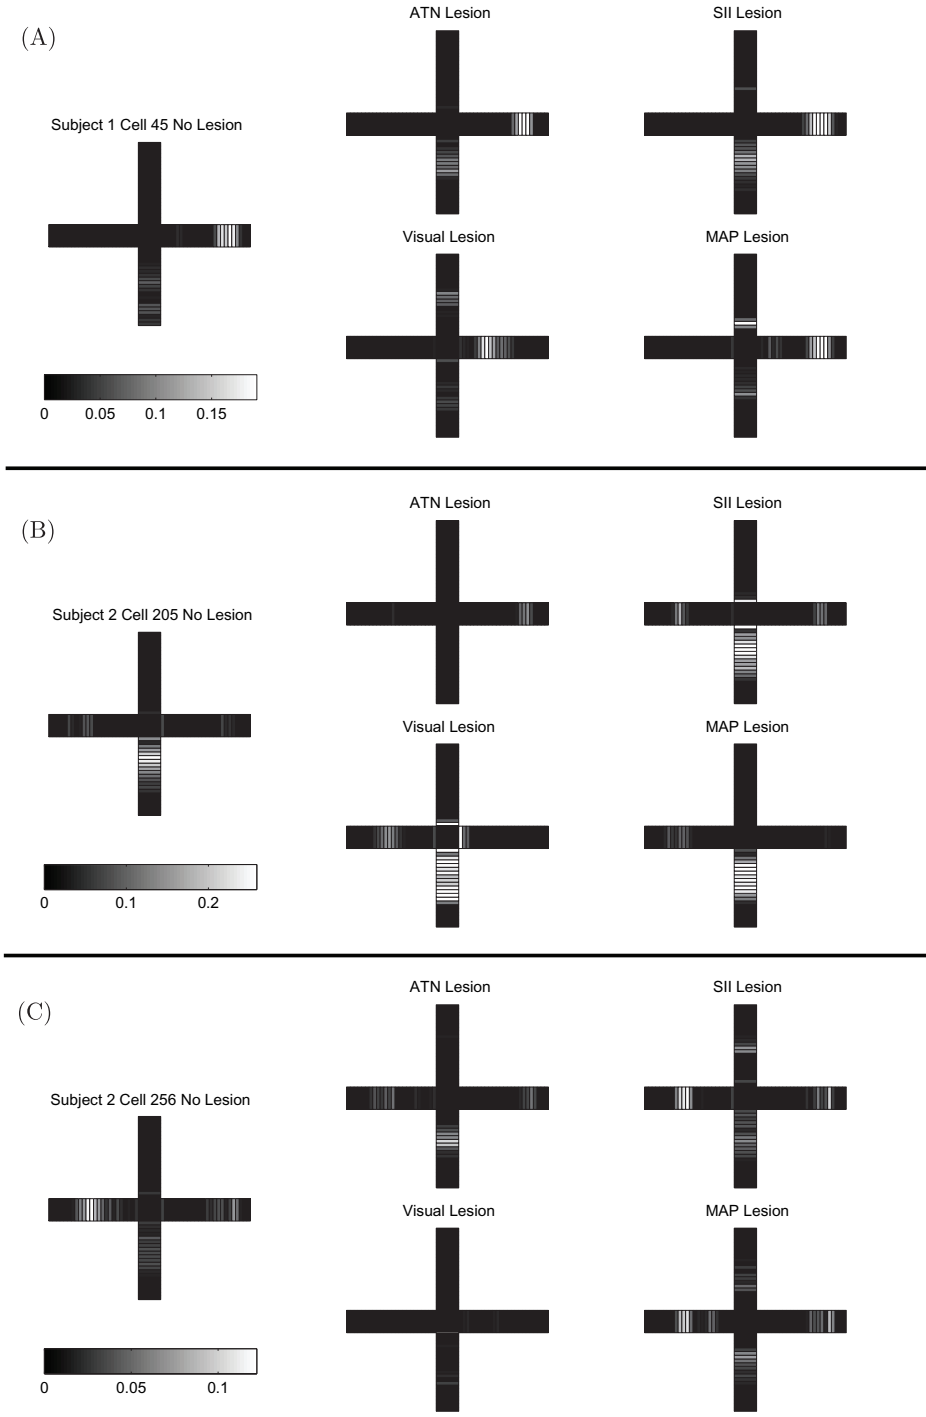

Fig. 4. The effects of lesions on the rate maps of three representative *CA1* place units. In (A), the place field was robust to lesions, changing very little for each individual lesion. In (B), the loss of head direction information (lesion of *ATN*) resulted in the loss of the place field. In (C), the place field remapped during the head direction lesion and lost its field during the visual lesion (lesion of *PR* and *IT*).

3.3. *Lesion effects on backtraces*

In this study, we used the backtrace analysis to investigate two questions: (1) When conditions are ostensibly the same, do multiple neural pathways generate a neuronal unit's place-correlated firing? (2) Under lesion conditions, are neural pathways leading to place-correlated firing quantitatively different?

To address these questions, backtraces were performed from *CA1* neuronal units that displayed place fields in both the normal condition and in at least two other lesion conditions. The neuronal unit from which a backtrace begins is called the reference unit. For each such reference unit, three to four trials with place activity and correct behavior choice were manually selected from the ten trials *before* behavioral criterion was reached (for the normal case) or the ten trials *after* behavioral criterion was reached (for lesioned cases). For each trial and reference unit combination, the first cycle on which the unit began to fire above 10% of the maximum firing rate became the start cycle for the backtrace. The backtrace analysis, described in the methods section, revealed the excitatory neuronal units in the six simulation time steps previous to the start cycle that caused the firing of the reference unit. Examples of backtraces can be seen in Fig. 5.

Backtraces from the same unit on different trials quickly become very different from each other, even though they led to the same place activity. These backtraces are degenerate, i.e. they are functionally equivalent but contain different neural pathways. In order to quantify the degeneracy, Fig. 6 depicts the amount of shared neuronal units when comparing backtraces. All comparisons were carried out on backtraces that started from the same reference unit, but differed in trial or condition. When comparing two normal backtraces of the same unit, there was nearly a complete overlap at the beginning of the backtraces. However, the overlap declined significantly (ANOVA, $p < 0.05$) by three time steps back ($t - 3$) reached an asymptote, where approximately half of all units being common between backtraces. Comparing elements of two backtraces from a reference unit under the same lesion condition produced similar results (not shown in this figure). This variance of causal pathways leading to the "same" neuronal event implies that a multitude of quite different activity patterns can lead to the same neuronal activity, i.e. the simulated nervous system is degenerate in the sense described in Ref. [12]. The overlap of normal backtraces ("Normal vs normal", Fig. 6) also provides a control condition with which to compare the same reference units' backtraces after a lesion.

Lesions produced backtraces that were quantitatively different from the normal condition even when the same cell had a similar place field. Figure 6 shows that by two or three times steps back, most of the lesion comparisons had less overlap than the control comparison. A conservative statistical test (ANOVA, $p < 0.01$, Tukey-Kramer $\alpha = 0.05$) revealed that there was a significant difference for all time steps between the vision lesion comparison ("Normal vs visual lesion", Fig. 6) and the control ("Normal vs normal", Fig. 6) comparison. For the head direction and laser lesion comparisons, there was a significant difference to the control comparison by

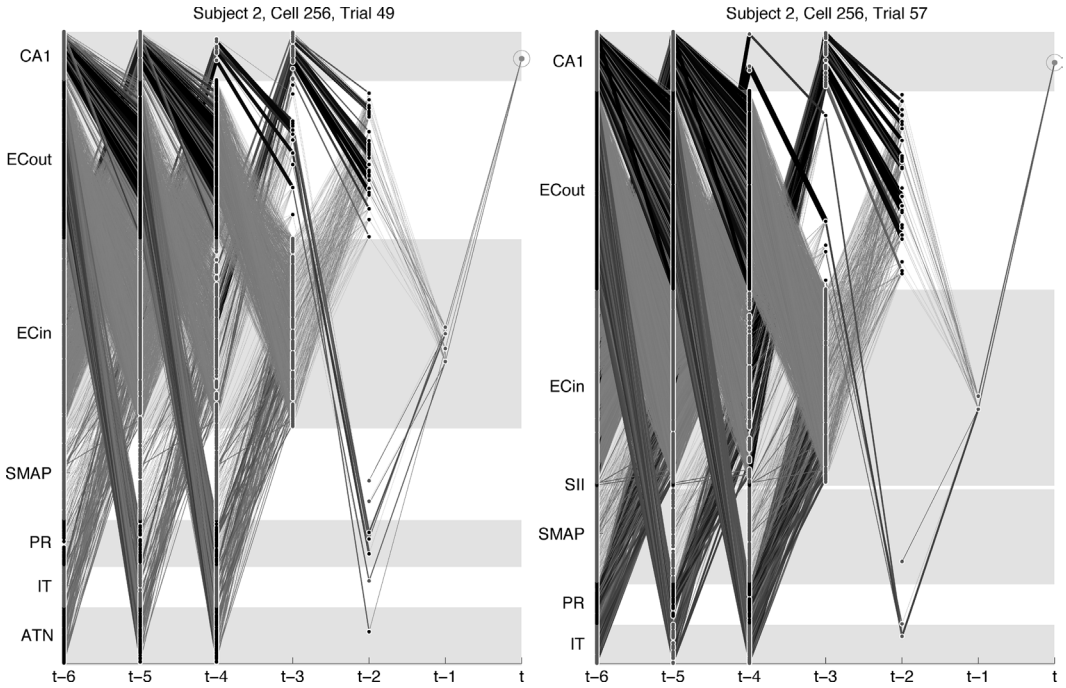

Fig. 5. Examples of backtraces starting from the reference neuronal unit in Fig. 4(A) (shown in the charts as a bulls-eye). The x-axis depicts simulation time ranging from six time cycles to one time cycle before the onset of reference unit activity. Each dot represents a neuronal unit in the backtrace. Lines between neuronal units denote the synaptic influence (presynaptic activity times synaptic weight). Thicker darker lines correspond to higher influence. On the left, the backtrace of the unit at the onset of place activity during a normal condition. On the right, the same unit at the onset of place activity after a lesion of the *ATN* head direction area, where the place field remaps to a different maze arm. At one time step back the backtraces share two common *ECin* units (50% overlap), and at two time steps back they share one *IT* unit, one *SMAP* unit, and 12 *ECout* units (46% overlap). The rest of the time steps have 38%, 38%, 43%, and 44% overlap at three through six times steps back respectively. The largest change in units occurs in entorhinal cortex, which is typical of lesion backtraces.

three time steps back. No significant differences were observed between control and whisker lesion comparisons. The increased degeneracy in backtraces after lesioning implies that place activity could be created through an entirely different set of pathways when sensory streams were lesioned, and that these pathways were not active when all sensory information was intact. Knowing that lesions produce a measurable difference in backtraces, we wondered whether lesion backtraces might differ from the normal condition in a systematic way.

To address this question, we examined the differences in neuronal unit composition between the backtraces in the normal and lesion conditions (see Fig. 7). The difference between conditions is characterized by two numbers, the neuronal units that existed in the backtrace during the normal condition but not during the lesion

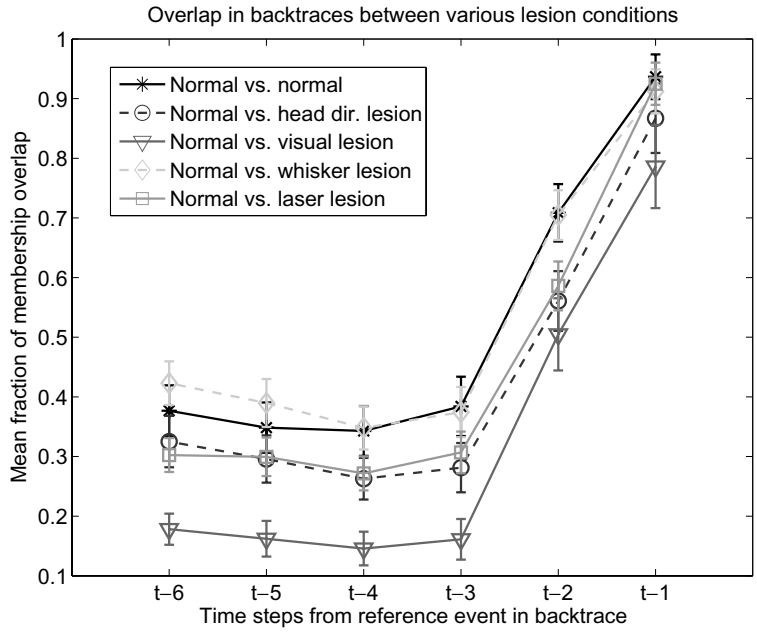

Fig. 6. Comparisons of neuronal elements among different backtraces starting from the same reference unit. The chart shows the overlap of neuronal units between backtraces from the normal condition (no lesion) and backtraces from various lesion conditions. The overlap is the number of neuronal units shared between backtraces divided by the number of neuronal units in the smaller of the two backtraces. In the lesion (normal) comparison, all backtraces of a particular reference unit in the normal condition are compared to all (other) backtraces of that same reference unit in the lesion (normal) condition. The figure displays the mean and standard error bars of particular comparisons, calculated from 39 reference units with place fields both in the normal case and in more than one lesion condition in the same location. Comparisons used 132 normal, 94 head direction lesion, 96 visual lesion, 93 laser lesion, and 107 whisker lesion backtraces.

condition (bars on the negative side of the abscissa of Fig. 7), and neuronal units that existed in the backtrace during the lesion condition but not during the normal condition (bars on the positive side of the abscissa of Fig. 7). We further divided these measured differences by which broad neural region they occurred in: sensory areas (*ATN*, *IT*, *PR*, *SII*, *SMAP*), entorhinal cortex (*ECin*, *ECout*), or hippocampus (*DG*, *CA3*, *CA1*). The raw number of differences were normalized to produce measurements of percentage change from the normal condition for each region, and for each time step of the backtrace.

The backtraces of place units showed different compositions after lesioning. Changes occurred primarily in sensory areas and entorhinal cortex, with the hippocampus changing significantly less (ANOVA, $p < 0.05$) than sensory or entorhinal areas (see Fig. 7). This result is true for backtraces from both place units that were stable after lesioning and for place units that remapped after lesioning. This implies that for any major change in the place field activity of a particular *CA1* neuronal

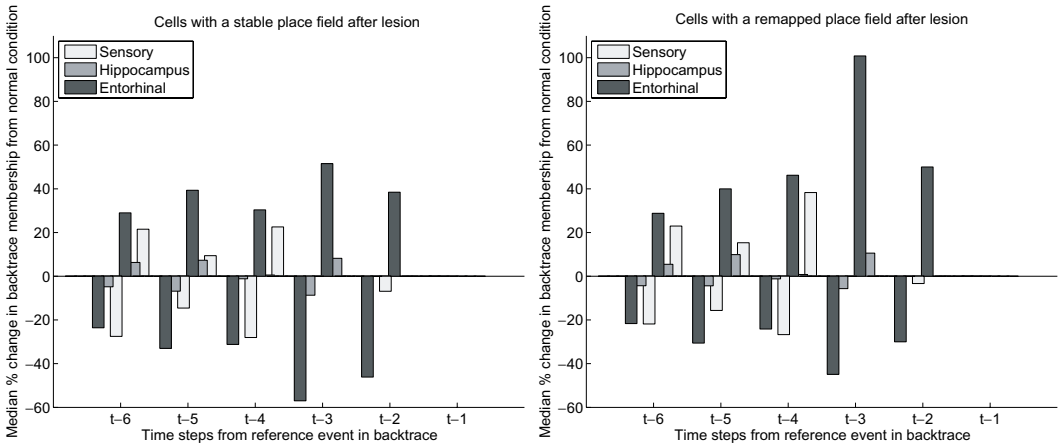

Fig. 7. Changes in backtrace membership occurring between the normal case and all lesion cases. The changes are broken down by time step of the backtrace (abscissa) and the ordinate shows the median percentage change from the normal case over all units. Each time step is further broken down into changes occurring in sensory areas, hippocampus, and entorhinal cortex. Neuronal units that were not in the normal condition backtrace but were present in the lesion condition backtrace are shown by the bars on the positive side of the abscissa. Neuronal units that were present in the normal condition but were not in the lesion condition are shown by the bars on the negative side of the abscissa. *Left:* changes in the backtraces from reference units that are stable after lesion. *Right:* changes in the backtraces from reference units that remap after lesion.

unit, we would expect to see a larger change in the population dynamics of *EC* than in hippocampal areas.

There was also a difference in backtrace composition when comparing place units that were stable after a lesion to those that remapped (left vs right subfigures in Fig. 7). Remapping place units also displayed more of an increase in new units added to the backtrace as compared to place units with stable fields. The most prominent feature of the remapped units' backtrace differences was a large addition of entorhinal units involved in the backtrace at time step $t-3$. An increased number of sensory units involved in the backtrace was also evident at time step $t-4$, which could conceivably have induced the entorhinal changes at time step $t-3$. It is possible that the entorhinal differences observed at $t-3$ created the remapping, and those changes were in turn caused by the sensory differences at $t-4$. A new study, experimentally designed to specifically address the sensitivity of place activity to perturbation of inputs could help elucidate why place fields remap.

4. Conclusion

The medial temporal lobe is thought to be necessary for the formation and maintenance of associative memories. In this study, we showed that a BBD having a simulated medial temporal lobe and receiving inputs from multiple sensors acquired multimodal memories that were sufficient for solving a standard plus-maze task. The

BBD's behavioral performance was robust to sensory lesions (see Fig. 3). No single lesion caused a deficit in maze navigation; only a combined lesion of both the laser localization input, which includes path integration information, and the visual input produced a significantly worse performance. Interestingly, a recent study of rodent navigation has very similar results [11]. Cooper and Mizumori induced multiple sensory deficits by selectively inactivating retrosplenial cortex, which they hypothesize to be involved in the path integration input to hippocampus, and running animals in the dark. Only combined inactivation and darkness produced navigation deficits in these animals, whereas the same animals were able to navigate in the light.

The BBD approach allowed us to look at the time course of all neural activity during behavior under a variety of conditions, something which is currently impossible in animal experiments. We observed the formation of a large number of place fields in Darwin XI's simulated hippocampus. These place fields were formed by integrating multiple streams of sensory input. Firing fields were often maintained no matter which single sensory stream was lesioned (see Table 1 and Fig. 4). Although some units lost their field when a particular lesion was made, the whole population of units maintained enough place activity so that correct behavioral choices were maintained until at least three of five sensory inputs were removed (see Fig. 3). There was also a strong relationship between behavioral performance after a lesion and the extent of changes in *CA1* place fields, which is a result also found in rodents [11].

An important contribution of the BBD approach is the ability to trace back in time the circuits responsible for creating a particular neural event. This backtrace analysis was performed on neuronal units that formed place fields to determine how consistently the same pathways were involved in forming place activity, and also whether quantitatively different pathways produced place activity after a lesion was made.

As expected from an examination of the behavioral results above, the backtrace analysis indicated that there were many degenerate pathways by which the same pattern of place field activity could be generated (see Figs. 5 and 6). Even when normal condition backtraces were compared with each other, at three time steps ($\sim$ 600 ms) back, there was less the 50% overlap between the pathways leading to the same place activity on successive trials. Lesions in general produced even less overlap, demonstrating that there are degenerate pathways that produce place field activity in lesion conditions that are not active before the lesion is made. This model predicts that in behaving animals similar degeneracy of pathways leading to place activity should be occurring, and that this degeneracy should be more pronounced during sensory deprivation conditions.

To quantify the kinds of changes produced by lesions, we looked at the difference in types of neuronal units (sensory, entorhinal, or hippocampal) that were present in backtraces before and after lesions occurred. The analysis of backtrace differences showed that changes in the causal chains leading to place field activity were more pronounced in sensory and entorhinal cortices than in the hippocampus itself (see Fig. 7). This implies that there may be less degeneracy of pathways once information

is flowing through hippocampus. Such a finding is consistent with theories that highlight the pattern completion abilities of hippocampus [33, 39]. However, this model does not perform pattern completion in the same way as the classical attractor network model of hippocampus. In attractor networks, all place fields would remap simultaneously during manipulation of sensory input [22]. Instead, the Darwin XI model dynamically switches neural pathways that lead to place activity, that in turn leads to behavioral responses. This is similar to rodent place fields, where there is evidence that place cells may remap independently of each other and may also display remapping without subsequent changes in behavior [22, 40, 17]. Therefore, the Darwin XI model may provide a more realistic mechanism to investigate remapping, although a recent theoretical extension of attractor networks to multiple simultaneous maps displays a similar ability to partially remap [49].

The utilization of a brain-based device had an advantage over electrophysiological recordings, in that the mechanisms of the simulated nervous system were fully observable. The emphasis on realistic neuroanatomy and neurophysiology allowed us to examine complete neuronal circuit dynamics during behavior, as well as to make predictions of how neural activity gave rise to such behavior. Of course, real progress in understanding how the hippocampus integrates sensory information will depend on further work in behaving animals. Our predictions could conceivably be tested as the technology for simultaneous multi-unit recording improves.

In this and previous studies with BBDs, activity in the simulated hippocampus has been shown to be place specific, journey dependent and multimodal in its responses. Our work with Darwin X and XI [14, 26, 25] has shown that BBDs simulating anatomical and physiological details of the medial temporal lobe and surrounding regions can support the formation of spatial memory, episodic memory, and associative memory. The results using these models may have heuristic value in analyzing findings from studies of behaving animals.

Acknowledgments

This work was supported by the Neurosciences Research Foundation, the Defense Advanced Research Agency (DARPA), and the Office of Naval Research. We thank Dr Gerald M. Edelman for his suggestions and valuable criticism. We would also like to thank Brian Cox, Donald Hutson, Doug Moore and Jim Snook for their contribution to the design of Darwin XI, and our anonymous reviewers for their useful comments.

Appendix A

This appendix provides algorithmic details of the neural simulation. Please refer throughout this section to Fig. 2 for a schematic of the simulation's connectivity and to Tables A.1 and A.2 for parameters describing the artificial anatomy and physiology.

Table A.1. Information about the neural areas in Darwin XI's simulated nervous system. The number of neuronal units and topology is shown in the size column. The other columns are physiological parameters for each neural area as defined in the simulation equations above. Roughly described, the parameters are scaling factor, g, firing rate thresholds, σ^{fire} and σ^{vdep}, and persistence factor, ω. Neural areas are described in both App. A and the legend of Fig. 2. This table includes inhibitory areas not described previously. Areas whose name has the suffix -iff, indicating they are inhibitory feedforward, receive projections from the excitatory area of the same name and project to another area. Areas whose name has the suffix -ifb, indicating they are inhibitory feedback, receive projections from the excitatory area of the same name and project back to that excitatory area.

Area	Size	g	σ^{fire}	σ^{vdep}	ω
4 × *Color* (Red, Green, Blue, Yellow)	60 × 80	1.0	0.0	0.1	0.0
4 × *Width* (*Wid2, 4, 8, 16*)	60 × 80	1.0	0.0	0.1	0.0
4 × $V1_{color}$	6 × 8	1.0	0.2	0.1	0.0
4 × $V1_{width}$	15 × 20	1.0	0.2	0.1	0.0
4 × *WL* (*LT, LM, RT, RM*)	1 × 20	1.0	0.3	0.0	0.8
4 × *SI* (*LT, LM, RT, RM*)	1 × 20	1.0	0.2	0.0	0.8
HD	1 × 360	1.0	0.2	0.1	0.0
M_{Hdg}	1 × 60	1.0	0.0	0.0	0.0
M_{Hdgi}	1 × 60	1.0	0.0	0.1	0.0
SMAP	30 × 30	1.0	0.2	0.1	0.0
SII	30 × 30	1.0	0.2	0.0	0.8
PR	30 × 30	1.0	0.0	0.1	0.0
IT	30 × 30	1.0	0.0	0.1	0.0
ITi	15 × 15	1.0	0.0	0.1	0.15
ATN	30 × 30	1.0	0.0	0.1	0.5
ATNi	30 × 30	1.0	0.0	0.1	0.15
ECIN	37 × 37	1.0	0.1	0.1	0.5
ECINifb	18 × 18	1.0	0.02	0.1	0.0
ECOUT	37 × 37	1.0	0.1	0.1	0.5
ECOUTifb	18 × 18	1.0	0.02	0.1	0.0
DG	37 × 37	0.75	0.1	0.1	0.5
DGifb	18 × 18	1.0	0.02	0.1	0.0
DGiff	18 × 18	1.0	0.02	0.1	0.0
CA3	18 × 18	1.25	0.05	0.1	0.5
CA3ifb	10 × 10	1.0	0.02	0.1	0.0
CA3iff	18 × 18	1.0	0.02	0.1	0.0
CA1	24 × 24	1.0	0.02	0.1	0.5
CA1ifb	12 × 12	1.0	0.02	0.1	0.0
CA1iff	12 × 12	1.0	0.02	0.1	0.0
BF	1 × 1	1.0	0.01	0.1	0.15
S	4 × 4	1.0	0.0	0.0	0.0
T+	1 × 1	1.0	0.0	0.0	0.0

Table A.2. Non-plastic connections between neural areas in Darwin XI's simulated nervous system. A presynaptic neuronal unit connects to a postsynaptic neuronal unit with a given probability (P) and a given projection topology (Arbor). This arborization shape can be rectangular with a particular height and width "$h \times w$", doughnut-shaped with an inner and outer radius "$\Theta r1 \times r2$", nontopological "nontopo" in which any pair of presynaptic and postsynaptic neurons have an equal probability of being connected, or "S2Special" indicating that the post-synaptic neuronal unit took input from three neuronal units, each of which was in a different subarea of S1. The initial connection strengths $c_{ij}(0)$ are set with uniform random probability within the range given by (minimum, maximum) in that column. A negative value of $c_{ij}(0)$ denotes inhibitory connections. Connections can be either voltage-independent (VI) or voltage-dependent (VD). ϕ denotes the persistence of the synapse. A nonzero value of η, the learning rate parameter, signals a plastic connection that changes according to a modified BCM rule with parameters k_1 and k_2. An explanation of neural area abbreviations can be found in both App. A and in Fig. 2.

Projection	Arbor	P	$c_{ij}(0)$	Type	ϕ	η	k_1	k_2
$4 \times Color \rightarrow V1_{color}$	[] 1×1	1.0	0.012, 0.014	VI	1.0	0.0	0.0	0.0
$4 \times V_{color} \rightarrow V1_{color}$	[] 2×2	0.4	0.5, 0.6	VD	1.0	0.0	0.0	0.0
$4 \times Width \rightarrow V1_{width}$	[] 1×1	1.0	0.008, 0.009	VI	1.0	0.0	0.0	0.0
$4 \times V1_{width} \rightarrow$ itself	[] 1×1	0.4	0.5, 0.6	VD	1.0	0.0	0.0	0.0
$4 \times V1_{width} \rightarrow$ all other $V1_{width}$	[] 2×2	1.0	-0.012, -0.014	VI	1.0	0.0	0.0	0.0
$4 \times V1_{color} \rightarrow IT$	nontopo	0.05	0.01, 0.04	VI	1.0	0.0	0.0	0.0
$IT \rightarrow IT$	[] 1×1	1.0	0.08, 0.14	VI	1.0	0.0	0.0	0.0
$IT \rightarrow ITi$	Θ 2×3	1.0	0.06, 0.08	VI	1.0	0.0	0.0	0.0
$ITi \rightarrow IT$	[] 1×1	1.0	-0.36, -0.5	VI	1.0	0.0	0.0	0.0
$4 \times V1_{width} \rightarrow PR$	[] 1×1	0.25	0.25, 0.3	VI	1.0	0.0	0.0	0.0
$PR \rightarrow PR$	Θ 4×6	1.0	-0.06, -0.08	VI	1.0	0.0	0.0	0.0
$HD \rightarrow M_{Hdg}$	[] 1×1	1.0	0.01, 0.01	VI	1.0	0.0	0.0	0.0
$M_{Hdg} \rightarrow M_{Hdg}$	[] 1×1	1.0	0.06, 0.08	VI	1.0	0.0	0.0	0.0
$M_{Hdg} \rightarrow M_{Hdgi}$	[] 20×30	0.5	0.1, 0.2	VI	1.0	0.0	0.0	0.0
$M_{Hdgi} \rightarrow M_{Hdg}$	[] 1×1	1.0	-0.36, -0.5	VI	1.0	0.0	0.0	0.0
$HD \rightarrow ATN$	[] 30×2	0.2	0.01, 0.02	VI	1.0	0.0	0.0	0.0
$ATN \rightarrow ATNi$	[] 10×15	0.25	0.01, 0.02	VI	1.0	0.0	0.0	0.0
$ATNi \rightarrow ATN$	[] 1×1	1.0	-0.36, -0.5	VI	1.0	0.0	0.0	0.0
$4 \times WL \rightarrow SI\ (L/R)(T/M)$	[] 0×0	1.0	0.6, 0.75	VI	1.0	0.0	0.0	0.0
$4 \times SI(L/R)(T/M) \rightarrow$ itself	Θ 2×8	1.0	-0.45, -0.6	VI	1.0	0.0	0.0	0.0
$4 \times SI(L/R)(T/M) \rightarrow SII$	S2special	0.0350	0.25, 0.25	VI	1.0	0.0	0.0	0.0
$IT \rightarrow ECIN$	nontopo	0.0007	0.4, 0.5	VI	1.0	0.0	0.0	0.0
$PR \rightarrow ECIN$	nontopo	0.0007	0.4, 0.5	VI	1.0	0.0	0.0	0.0
$ATN \rightarrow ECIN$	nontopo	0.0007	0.2, 0.25	VI	1.0	0.0	0.0	0.0
$SMAP \rightarrow ECIN$	nontopo	0.0007	0.2, 0.25	VI	1.0	0.0	0.0	0.0
$SII \rightarrow ECIN$	nontopo	0.0007	0.2, 0.25	VI	1.0	0.0	0.0	0.0
$ECOUT \rightarrow IT$	nontopo	0.01	0.4, 0.45	VD	1.0	0.0	0.0	0.0
$ECOUT \rightarrow PR$	nontopo	0.01	0.4, 0.45	VD	1.0	0.0	0.0	0.0
$ECOUT \rightarrow ATN$	nontopo	0.01	0.4, 0.45	VD	1.0	0.0	0.0	0.0
$ECOUT \rightarrow SMAP$	nontopo	0.01	0.4, 0.45	VD	1.0	0.0	0.0	0.0

Table A.2. (*Continued*)

Projection	Arbor	P	$c_{ij}(0)$	Type	ϕ	η	k_1	k_2
ECOUT → *SII*	nontopo	0.01	0.4, 0.45	VD	1.0	0.0	0.0	0.0
ECIN → *ECINifb*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
ECINifb → *ECIN*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
ECOUT → *ECOUTifb*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
ECOUTifb → *ECOUT*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
DG → *DG*	[] 1 × 1	1.0	0.1, 0.14	VI	1.0	0.0	0.0	0.0
DG → *DGifb*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
DGifb → *DG*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
CA3 → *CA3ifb*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
CA3ifb → *CA3*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
CA1 → *CA1ifb*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
CA1ifb → *CA1*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
DG → *CA3iff*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
CA3iff → *CA3*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
CA3 → *CA1iff*	Θ 2 × 3	0.1	0.45, 0.6	VI	1.0	0.0	0.0	0.0
CA1iff → *CA1*	[] 1 × 1	1.0	−0.9, −1.20	VI	1.0	0.0	0.0	0.0
BF → *ECIN*	nontopo	0.05	−0.01, −0.02	VI	1.0	0.0	0.0	0.0
BF → *ECOUT*	nontopo	0.05	−0.01, −0.02	VI	1.0	0.0	0.0	0.0
BF → *DG*	nontopo	0.05	−0.01, −0.02	VI	1.0	0.0	0.0	0.0
BF → *CA3*	nontopo	0.05	−0.01, −0.02	VI	1.0	0.0	0.0	0.0
BF → *CA1*	nontopo	0.05	−0.01, −0.02	VI	1.0	0.0	0.0	0.0
ECIN → *ECOUT*	nontopo	0.05	0.04, 0.08	VI	1.0	0.0	0.0	0.0
ECOUT → *ECIN*	nontopo	0.05	0.04, 0.08	VI	1.0	0.0	0.0	0.0
IT → *PR*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
IT → *ATN*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
PR → *IT*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
PR → *ATN*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
ATN → *PR*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
ATN → *IT*	nontopo	0.005	0.01, 0.02	VI	1.0	0.0	0.0	0.0
$T+$ → *S*	nontopo	1.0	0.25, 0.25	VI	1.0	0.0	0.0	0.0
S → *CA1*	nontopo	1.0	0.001, 0.002	VD	1.0	0.0	0.0	0.0
S → *ATN*	nontopo	1.0	0.001, 0.002	VD	1.0	0.0	0.0	0.0
S → *IT*	nontopo	1.0	0.001, 0.002	VD	1.0	0.0	0.0	0.0
S → *PR*	nontopo	1.0	0.001, 0.002	VD	1.0	0.0	0.0	0.0
S → M_{Hdg}	nontopo	1.0	0.05, 0.06	VD	1.0	0.0	0.0	0.0

Sensory input

Vision

Visual images from Darwin XI's CCD camera were filtered for color and edges. The filtered output directly affected neural activity in area *V1*, which is composed of

Table A.3. Plastic connections between neural areas in Darwin XI's simulated nervous system. A presynaptic neuronal unit connects to a postsynaptic neuronal unit with a given probability (P) and a given projection topology (Arbor). This arborization shape can be rectangular with a particular height and width "$h \times w$", doughnut-shaped with an inner and outer radius "Θ $r1 \times r2$", nontopological "nontopo" in which any pair of presynaptic and postsynaptic neurons have an equal probability of being connected, or "S2Special" indicating that the post-synaptic neuronal unit took input from 3 neuronal units, each of which was in a different subarea of S1. The initial connection strengths $c_{ij}(0)$ are set with uniform random probability within the range given by (minimum, maximum) in that column. A negative value of $c_{ij}(0)$ denotes inhibitory connections. Connections can be either voltage-independent (VI) or voltage-dependent (VD). ϕ denotes the persistence of the synapse. A nonzero value of η, the learning rate parameter, signals a plastic connection that changes according to a modified BCM rule with parameters k_1 and k_2. An explanation of neural area abbreviations can be found in both App. A and in Fig. 2.

Projection	Arbor	P	$c_{ij}(0)$	Type	ϕ	η	k_1	k_2
$ECIN \rightarrow DG$	[] 3×3	0.1	0.45,0.6	VI	0.25	0.05	0.9	0.45
$ECIN \rightarrow CA3$	[] 3×3	0.05	0.15,0.2	VI	0.25	0.05	0.9	0.45
$ECIN \rightarrow CA1$	[] 3×3	0.04	0.3,0.4	VI	0.25	0.05	0.9	0.45
$DG \rightarrow CA3$	[] 3×3	0.03	0.45,0.6	VI	0.25	0.05	0.9	0.45
$CA3 \rightarrow CA3$	nontopo	0.1	0.15,0.2	VI	0.25	0.05	0.9	0.45
$CA3 \rightarrow CA1$	[] 3×3	0.08	0.45,0.6	VI	0.25	0.05	0.9	0.45
$CA1 \rightarrow ECOUT$	[] 3×3	0.25	0.6,0.75	VI	0.25	0.05	0.9	0.45
$CA1 \rightarrow S$	nontopo	1.0	0.01,0.02	VI	1.0	0.005	0.9	0.45
$CA1 \rightarrow M_{Hdg}$	nontopo	1.0	0.01,0.02	VD	1.0	0.05	0.9	0.45

functionally segregated subareas for color and shape. The CCD camera produced an 80×60 pixel RGB image. Different sized Gabor filters (2×2, 4×14, 16×16, and 32×32) were used to detect vertical edges of varying widths. The output of the Gabor function mapped directly onto the neuronal units of the corresponding *V1* sub-area (*V1-width2*, *V1-width4*, *V1-width16*, and *V1-width32*). The RGB video was transformed into YUV color space. Color filters in UV space (red, green, yellow, and blue) were applied to the image. The outputs of the color filters were mapped directly onto the neuronal units of *V1-red*, *V1-green*, *V1-blue*, and *V1-yellow*. *V1* color neuronal units projected non-topologically to inferotemporal cortex *IT*, and *V1* edge units projected retinotopically to parietal cortex *Pr*.

The use of Gabor filters to model the receptive fields of simple cells of primary visual cortex is widespread in the literature. The concept is due to [29] and was first tested rigorously by [21].

Head direction

A head direction system was modeled after areas of the rodent nervous system (e.g., anterior thalamic nuclei) that respond selectively to the animal's heading [36, 46]. Neurons in these areas are often called head direction cells. A magnetic compass was used to estimate current heading. This information was input into the head direction neural area (*HD*). Each of the 360 *HD* neuronal units had a cosine tuning

curve, which responded maximally to a preferred heading with a tuning width of π radians:

$$(\cos(HD_i - curr\_heading))^5, \tag{A.1}$$

where HD_i is a head direction cell with a preferred direction of $(\frac{i}{360}2\pi)$ and i ranges from 0 to 359.

The head direction cells projected topographically to an area analogous to the anterior thalamic nucleus ($HD \rightarrow ATN$) and to a motor area ($HD \rightarrow M_{HDG}$) used for selecting a new heading (see below).

The head direction system used in this model is intended to be functionally, rather than mechanistically equivalent to head direction cells found in rodents. More biologically realistic models that include mechanisms for the integration of vestibular and local visual cue input exist, e.g., [45, 32].

Whiskers

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

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

$$s_{ki}^{in}(t+1) = \begin{cases} 0.2 & s_{ki}^{in}(t) < 0.2, \overline{diff}_k(t) > 3.0 \\ 0 & s_{ki}^{in}(t) \geq \sigma^{fire} \\ (1+\psi_i)s_{ki}^{in}(t) & \text{otherwise,} \end{cases} \tag{A.2}$$

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

The output s_{ki} is calculated using:

$$s_{ki}(t+1) = \begin{cases} \tanh(10(\omega_i s_{ki}(t))) & s_{ki}(t) < \sigma^{fire} \\ \tanh(10(\omega_i s_{ki}(t) + (1-\omega_i)s_{ki}^{in}(t))) & \text{otherwise,} \end{cases} \tag{A.3}$$

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

Darwin XI had three sensory whiskers on each side arranged in a vertical stack. Whiskers are therefore referred to by the symbols LT, LM, LB to describe respectively the left top, left middle, and left bottom whiskers. Thalamic whisker lag (WL) and primary somatosensory (S1) areas in the simulation use these position suffixes to denote the whisker from which they receive input (e.g., *WLLT, S1LT*).

This whisker system model was originally used by Darwin IX to demonstrate that combining whisker deflection information at a variety of lags allowed successful texture discrimination [44]. Other models of rat somatosensory cortex have focussed on different transformations of several aspects of the whisker sensory signals [27, 1, 19].

Laser rangefinder

Darwin XI had a SICK LMS-200 laser rangefinder (SICK A.G., Waldkirch, Germany) that produced distance-to-wall data, which was used by the freely-available CARMEN software package (*http://carmen.sourceforge.net/*) to produce an estimate of device location. CARMEN took noisy laser and self-movement data and used a Bayesian particle filter to optimally combine these sources of information to localize the device against a prebuilt occupancy grid map of the arena. This location estimate was then turned into population-coded activity in a topologically-organized pseudo-cortical area (*SMAP*) that projected to EC_{in}. Each neuronal unit in *SMAP* responded preferentially to a particular place in the arena environment, and had a Gaussian tuning curve with 0.8 m standard deviation.

This is the only sensory input which is not directly analogous to an area of the vertebrate brain, but such information can be seen as similar to what an accurate path integration system, long hypothesized to be present in animals [6] would generate. It can be noted that although the encoding of spatial information is different than that used by various proposed path integration systems in the vertebrate brain, e.g., associative parietal [42] or medial entorhinal [31] cortices, the type of information present is equivalent. However, the purpose of the *SMAP* area is not to model an actual biological input stream, but merely to add another kind of input to investigate the multimodal formation of neural correlates of episodic memory. Nor is the rangefinder input necessary for the formation of place activity in the hippocampus. We have previously reported [25] on a similar neural model which lacks the *SMAP* input and yet still generates place activity in the *CA*1 hippocampal subfield with equivalent spatial information content.

The theory, technique and literature of robot localization using particle filters is reviewed in [48].

Hippocampus

The architecture of the simulated hippocampal formation was based on the known literature about the gross connectivity and microarchitecture of rodent neuroanatomy. This approach is distinct from other robotic navigation models of hippocampus function [3, 5, 10] in that it focusses on the how a large-scale model of this

anatomy can produce known single-unit phenomena with experience-dependent plasticity. In this model, the anatomical connectivity is fixed, as are synaptic strengths in the sensory input streams, but there are plastic connections within the hippocampus and between hippocampus and the neural areas responsible for action selection.

The input streams into the hippocampus are from the input areas of the simulation ($ATN \rightarrow EC_{IN}$, $IT \rightarrow EC_{IN}$, $PR \rightarrow EC_{IN}$). Parameter values for the neuronal units and connections in these areas were tuned such that each cortical area (*ATN*, *PR*, and *IT*) had an equivalent synaptic influence on EC_{IN}. The relative numbers of neuronal units in each area, and the intrinsic and extrinsic of connectivity of the hippocampus were implemented based on known anatomical measurements [2, 8, 50]. The perforant path projects mainly from entorhinal cortex to the dentate gyrus but also to the CA3 and CA1 subfields ($EC_{IN} \rightarrow DG$ $EC_{IN} \rightarrow CA3$, $EC_{IN} \rightarrow CA3$). The mossy fibers ($DG \rightarrow CA3$), Schaffer collaterals ($CA3 \rightarrow CA1$), and divergent projections from the hippocampus back to cortex ($CA1 \rightarrow EC_{OUT} \rightarrow ATN, IT, PR$) were also reflected in the neural simulation. Moreover, the prevalent recurrent connectivity found in the hippocampal formation was included in the model ($EC_{IN} \rightarrow EC_{OUT}$, $DG \rightarrow DG$, and $CA3 \rightarrow CA3$).

There are distinct patterns of intrinsic and extrinsic, feedback and feedforward inhibitory connections in the hippocampal circuitry [8, 16]. Feedback inhibitory connections ($EC \rightarrow EC_{FB} \rightarrow EC$, $DG \rightarrow DG_{FB} \rightarrow DG$, $CA3 \rightarrow CA3_{FB} \rightarrow CA3$, and $CA1 \rightarrow CA1_{FB} \rightarrow CA1$) and feedforward inhibitory connections ($EC \rightarrow DG_{FF} \rightarrow DG$, $DG \rightarrow CA3_{FF} \rightarrow CA3$, and $CA3 \rightarrow CA1_{FF} \rightarrow CA1$) were included in the model. These connections were important for separating inputs and maintaining network stability.

A simplified model of the basal forebrain provided an extrinsic theta rhythm for the neural simulation. The function of the simulated basal forebrain area was to gate input into the hippocampus and keep activity levels stable. The *BF* area had a rhythmic activity over 13 simulation cycles:

$$BF(t) = theta(t \bmod 13), \tag{A.4}$$

where $theta = \{0.01, 0.165, 0.33, 0.495, 0.66, 0.825, 1.00, 0.825, 0.66, 0.495, 0.33, 0.165, 0.01\}$. *BF* projected to all hippocampal areas with inhibitory connections ($BF \rightarrow EC_{IN}$, EC_{OUT}, *DG*, *CA3*, *CA1*). To allow this discretized sine wave to be smooth and also allow sufficient time for new input to propagate through the network, was set to 13. Owing to computational limitations, the duration of the cycle was not comparable to that of real nervous systems. The level of inhibition, which was adaptive, kept the activity in hippocampal regions within specific ranges:

$$\begin{aligned} \Delta sf_r(t) &= (s_r(t) - tgt_r), \\ BF_r(t) &= BF(t) + sf_r(t), \end{aligned} \tag{A.5}$$

where r denotes the region (i.e., EC_{IN}, EC_{OUT}, *DG*, *CA3*, *CA1*), $sf_r(t)$ is the scale factor at time t, $s_r(t)$ is the percentage of active neuronal units in region r at time t, tgt_r is the desired percentage of active units in area r ($EC_{IN} = 10\%$, $EC_{OUT} = 10\%$,

$DG = 20\%$, $CA3 = 5\%$, and $CA1 = 10\%$), and $BF_r(t)$ is the presynaptic neuronal unit activity for a connection from BF to hippocampus region r.

Action selection

Activity in the simulated value system (Area S, Fig. 2) signals the occurrence of salient sensory events and this activity contributes to the modulation of value-dependent connection strengths in synaptic pathways ($CA1 \rightarrow S$ and $CA1 \rightarrow M_{HDG}$). The projection from our simulated $CA1$ to the value and goal decision areas is consistent with the connectivity between $CA1$ and nucleus accumbens and frontal areas [34, 47]. Initially, S is activated by the hidden platform IR detector ($T^+ \rightarrow S$), causing potentiation of value-dependent connections. After experience, the value system could be activated by $CA1$. The magnitude of potentiation or depression is based on a neural implementation of a temporal difference (TD) learning rule [7, 35]. Other hippocampus models have also applied temporal difference learning for action selection in a spatial navigation task [4, 15]. The TD rule applied in this model is:

$$TD(t) = \begin{cases} T^+(t) - \overline{S(t-\tau)} & T^+ > 0 \\ \overline{S(t)} - \overline{S(t-\tau)} & \text{otherwise,} \end{cases} \tag{A.6}$$

where $\overline{S(t)}$ is the average activity of the value system at time t, τ is one theta cycle (13 simulation cycles), T^+ is positive reward and equal to 1 if the BBD is over the hidden platform. The basic idea of the temporal difference rule is that learning is based on the difference between temporally successive predictions of rewards. The goal of learning is to make the learner's current prediction of expected reward match more closely the actual expected reward at the next time interval (τ). If the expected reward value increases over τ, TD is positive and affected synaptic connections are potentiated, and if the change in value decreases, TD is negative and affected synaptic connections are depressed. Further details on how the temporal difference is applied to individual synaptic connections are given in the *Neuronal Dynamics* section below.

Darwin XI selected a new heading when reaching the choice point of the plus-maze. The device stopped moving forward, and the camera was panned 90° left of centerline, and waited for three seconds, then panned the camera 90° right of centerline and waited for a further three seconds. The average activity of M_{HDG} was calculated during the wait periods. A softmax algorithm was used to create a probability distribution for choosing a new heading:

$$p(newhdg) = \frac{\exp(40\overline{M_{HDG}(newhdg)})}{\sum_{h \in \{hdg-90,\, hdg+90\}} \exp\left(40\overline{M_{HDG}(h)}\right)}, \tag{A.7}$$

where $newhdg$ is a possible new heading for Darwin XI, $\overline{M_{HDG}(newhdg)}$ is the average activity of M_{HDG} at a possible new heading, hdg is the current heading, and h has two positions (current heading less 90°, current plus 90°).

Neuronal dynamics and synaptic plasticity

A neuronal unit in Darwin XI is simulated by a mean firing rate model, in which the mean firing rate variable of each unit corresponds to the average activity of a group of roughly 100 real neurons during a time period of approximately 200 ms. Synaptic connections between neural units, both within and between neuronal areas, are set to be either voltage-independent or voltage-dependent, and either plastic or non-plastic. Voltage-independent connections provide synaptic input regardless of postsynaptic state. Voltage-dependent connections represent the contribution of receptor types (e.g., NMDA receptors) that require postsynaptic depolarization to be activated [18, 53].

The mean firing rate (s) of each neuronal unit ranges continuously from 0 (quiescent) to 1 (maximal firing). The state of a neuronal unit is updated as a function of its current state and contributions from voltage-independent and voltage-dependent inputs (see Fig. 2). The voltage-independent input to unit i from unit j is:

$$A_{ij}^{VI}(t) = c_{ij}s_j(t), \tag{A.8}$$

where $s_j(t)$ is the activity of unit j, and c_{ij} is the connection strength from unit j to unit i. The voltage-independent postsynaptic influence, $POST_i^{VI}$, on unit i is calculated by summing over all the inputs onto unit i:

$$POST_i^{VI}(t) = \varphi(POST_i^{VI}(t-1)) + (1-\varphi)\left(\sum_{l=1}^{M}\sum_{j=1}^{N_l} A_{ij}^{VI}(t)\right), \tag{A.9}$$

where M is the number of different anatomically defined connection types (see Table A.2), N_l is the number of connections of type M projecting to unit i, and φ is the persistence of synaptic input.

The voltage-dependent input to unit i from unit j is:

$$A_{ij}^{VD}(t) = \Phi(POST_i^{VI}(t))c_{ij}s_j(t), \quad \text{where } \Phi(x)\begin{cases} 0 & x < \sigma_i^{vdep} \\ x & \text{otherwise,} \end{cases} \tag{A.10}$$

and σ_i^{vdep} is a threshold for the postsynaptic activity below which voltage-dependent connections have no effect (see Table A.1).

The voltage-dependent postsynaptic influence on unit i, $POST_i^{VD}$, is given by:

$$POST_i^{VD}(t) = \varphi(POST_i^{VD}(t-1)) + (1-\varphi)\left(\sum_{l=1}^{M}\sum_{j=1}^{N_l} A_{ij}^{VD}(t)\right). \tag{A.11}$$

The total post-synaptic influence on neuronal unit i is given by:

$$POST_i = \sum_{j=1}^{N_{VI}} POST_j^{VI}(t) + \sum_{k=1}^{N_{VD}} POST_k^{VD}. \tag{A.12}$$

The new activity is determined by the following activation function:

$$s_i(t+1) = \phi(\tanh\left(g_i\, POST_i + \omega\, s_i(t)\right)), \quad \text{where } \phi(x) \begin{cases} 0 & x < \sigma_i^{fire} \\ x & \text{otherwise,} \end{cases} \tag{A.13}$$

where g_i is a scaling constant and ω is a persistence constant for a given neuronal unit.

Synaptic strengths are subject to modification according to a synaptic rule that depends on the pre- and postsynaptic neuronal unit activities. Plastic synaptic connections are either value-independent (EC_{IN}→DG, $CA3$, $CA1$; DG→$CA3$; $CA3$→$CA1$; $CA1$→EC_{OUT}) or value-dependent ($CA1$→S, $CA1$→M_{HDG}). Both of these rules are based on a modified BCM learning rule [9], which has been shown to be equivalent to spike-timing dependent plasticity under certain conditions [20]. Synapses between neuronal units with strongly correlated firing phases are potentiated and synapses between neuronal units with weakly correlated phases are depressed; the magnitude of change is determined as well by pre- and postsynaptic activities. The specific parameter settings for fine-scale synaptic connections are given in the equations below and Table A.2.

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

$$\Delta c_{ij}(t+1) = \eta s_i(t)\, s_j(t)\, BCM(s_i). \tag{A.14}$$

where $s_i(t)$ and $s_j(t)$ are activities of post- and presynaptic units, respectively, and η is a fixed learning rate. The function BCM is implemented as a piecewise linear function, taking postsynaptic activity as input, which is defined by a sliding threshold, θ, two inclinations (k_1, k_2) and a saturation parameter ρ ($\rho = 6$ throughout):

$$BCM(s) = \begin{cases} -k_1 s & s \le \frac{\theta}{2} \\ k_1(s-\theta) & \frac{\theta}{2} < s \le \theta \\ \frac{k_2}{\rho}\tanh\rho(s-\theta) & \text{otherwise,} \end{cases} \tag{A.15}$$

The threshold is adjusted based on the postsynaptic activity:

$$\Delta\theta = 0.25(s^2 - \theta). \tag{A.16}$$

Value-independent plasticity was subject to weight normalization to prevent unbounded potentiation:

$$c_{ij} = \frac{c_{ij}}{\sqrt{\sum_{k=l}^{K} c_{kj}^2}}, \tag{A.17}$$

where c_{ij} is a particular connection, and K is the total number of connections onto unit j.

The rule for value-dependent plasticity differs from the value-independent rule in that synaptic change is governed by the presynaptic activity, postsynaptic activity,

and temporal difference derived from the value system. The synaptic change for value-dependent synaptic plasticity is given by:

$$\Delta c_{ij}(t+1) = \eta\, s_i(t)\, TD(t) - 0.002(c_{ij}(t) - c_{ij}(0)), \tag{A.18}$$

where $TD(t)$ is the temporal difference value at time t (see Eq. (A.6)). The second term in Eq. (18) generates synaptic decay towards the initial weight. This was necessary to enable the device to reverse its behavioral choice in the plus-maze task.

References

[1] Ahissar E, Arieli A, Figuring space by time, *Neuron* **32**(2):185–201, 2001.

[2] Amaral D, Ishizuka N, Claiborne B, Neurons, numbers, and the hippocampal network, *Prog Brain Res* **83**:1–11, 1990.

[3] Arleo A, Gerstner W, Spatial cognition and neuromimetic navigation: A model of hippocampal place cell activity, *Biol Cybern* **83**(3):287–299, 2000.

[4] Arleo A, Smeraldi F, Gerstner W, Cognitive navigation based on nonuniform Gabor space sampling, unsupervised growing networks, and reinforcement learning, *IEEE Trans Neural Netw* **15**(3):639–652, 2004.

[5] Banquet JP, Gaussier P, Quoy M, Revel A, Burnod Y, A hierarchy of associations in hippocampo-cortical systems: Cognitive maps and navigation strategies, *Neural Comput* **17**(6):1339–1384, 2005.

[6] Barlow JS, Inertial navigation as a basis for animal navigation, *J Theor Biol* **6**(1):76–117, 1964.

[7] Barto AG, Sutton RS, Time-derivative models of pavlovian reinforcement, in *Learning and Computational Neuroscience: Foundations of Adaptive Networks*, MIT Press, Cambridge, Massachusetts, 1990.

[8] Bernard C, Wheal H, Model of local connectivity patterns in CA3 and CA1 areas of the hippocampus, *Hippocampus* **4**:497–529, 1994.

[9] Bienenstock EL, Cooper LN, Munro PW, Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex, *J Neurosci* **2**:32–48, 1982.

[10] Burgess N, Donnett JG, Jeffery KJ, O'Keefe J, Robotic and neuronal simulation of the hippocampus and rat navigation, *Philos Trans R Soc Lond B Biol Sci* **352**(1360):1535–1543, 1997.

[11] Cooper BG, Mizumori SJ, Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus, *J Neurosci* **21**(11):3986–4001, 2001.

[12] Edelman GM, Gally JA, Degeneracy and complexity in biological systems, *Proc Natl Acad Sci USA* **98**(24):13763–13768, 2001.

[13] Ferbinteanu J, Shapiro ML, Prospective and retrospective memory coding in the hippocampus, *Neuron* **40**(6):1227–1239, 2003.

[14] Fleischer JG, Gally JA, Edelman GM, Krichmar JL, Retrospective and prospective responses arising in a modeled hippocampus during maze navigation by a brain-based device, *Proc Natl Acad Sci USA* **104**(9):3556–3561, 2007.

[15] Foster DJ, Morris RG, Dayan P, A model of hippocampally dependent navigation, using the temporal difference learning rule, *Hippocampus* **10**(1):1–16, 2000.

[16] Freund TF, Buzsaki G, Interneurons of the hippocampus, *Hippocampus* **6**(4):347–470, 1996.

[17] Fyhn M, Hafting T, Treves A, Moser MB, Moser EI, Hippocampal remapping and grid realignment in entorhinal cortex, *Nature* **446**(7132):190–194, 2007.

[18] Grossberg S, The link between brain learning, attention, and consciousness, *Conscious Cogn* **8**(1):1–44, 1999.

[19] Hipp J, Arabzadeh E, Zorzin E, Conradt J, Kayser C, Diamond ME, Texture signals in whisker vibrations, *J Neurophysiol* **95**(3):1792–1799, 2006.

[20] Izhikevich EM, Desai N, Relating STDP to BCM, *Neural Comput* **15**:1511–1523, 1982.

[21] Jones JP, Palmer LA, An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex, *J Neurophysiol* **58**(6):1233–1258, 1987.

[22] Knierim JJ, Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells, *J Neurosci* **22**(14):6254–6264, 2002.

[23] Krichmar J, Edelman G, Machine psychology: Autonomous behavior, perceptual categorization and conditioning in a brain-based device, *Cereb Cortex* **12**(8):818–830, 2002.

[24] Krichmar J, Edelman G, Brain-based devices for the study of nervous systems and the development of intelligent machines, *Artif Life* **11**(1–2):63–78, 2005.

[25] Krichmar JL, Nitz DA, Gally JA, Edelman GM, Characterizing functional hippocampal pathways in a brain-based device as it solves a spatial memory task, *Proc Natl Acad Sci USA* **102**(6):2111–2116, 2005.

[26] Krichmar JL, Seth AK, Nitz DA, Fleischer JG, Edelman GM, Spatial navigation and causal analysis in a brain-based device modeling cortical-hippocampal interactions, *Neuroinformatics* **3**(3):197–221, 2005.

[27] Kyriazi HT, Simons DJ, Thalamocortical response transformations in simulated whisker barrels, *J Neurosci* **13**(4):1601–1615, 1993.

[28] Lavenex P, Amaral D, Hippocampal-neocortical interaction: A hierarchy of associativity, *Hippocampus* **10**(4):420–430, 2000.

[29] Marcelja S, Mathematical description of the responses of simple cortical cells, *J Opt Soc Am* **70**(11):1297–1300, 1980.

[30] McKinstry JL, Edelman GM, Krichmar JL, A cerebellar model for predictive motor control tested in a brain-based device, *Proc Natl Acad Sci USA* **103**(9):3387–3392, 2006.

[31] McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB, Path integration and the neural basis of the "cognitive map", *Nat Rev Neurosci* **7**(8):663–678, 2006.

[32] McNaughton BL, Chen LL, Markus EJ, Dead reckoning, landmark learning and the sense of direction: A neurophysiological and computational hypothesis, *J Cogn Neuro* **3**(2):190–202, 1991.

[33] Mizumori SJ, McNaughton BL, Barnes CA, Fox KB, Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of cac output: Evidence for pattern completion in hippocampus, *J Neurosci* **9**(11):3915–3928, 1989.

[34] Mogenson GJ, Nielsen M, Neuropharmacological evidence to suggest that the nucleus accumbens and subpallidal region contribute to exploratory locomotion, *Behav Neural Biol* **42**(1):52–60, 1984.

[35] Montague PR, Dayan P, Sejnowski TJ, A framework for mesencephalic dopamine systems based on predictive hebbian learning, *J Neurosci* **16**(5):1936–1947, 1996.

[36] Muller RU, Ranck JBJ, Taube JS, Head direction cells: Properties and functional significance, *Curr Opin Neurobiol* **6**(2):196–206, 1996.

[37] O'Keefe J, Conway DH, Hippocampal place units in the freely moving rat: Why they fire where they fire, *Exp Brain Res* **31**:573–590, 1978.

[38] O'Keefe J, Dostrovsky J, The hippocampus as a spatial map preliminary evidence from unit activity in the freely-moving rat, *Brain Res* **34**(1):171–175, 1971.

[39] O'Reilly RC, McClelland JL, Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off, *Hippocampus* **4**(6):661–682, 1994.

[40] Paz-Villagran V, Save E, Poucet B, Independent coding of connected environments by place cells, *Eur J Neurosci* **20**(5):1379–1390, 2004.

[41] Quirk G, Muller R, Kubie J, The firing of hippocampal place cells in the dark depends on the rat's recent experience, *J Neurosci* **10**(6):2008–2017, 1990.

[42] Save E, Guazzelli A, Poucet B, Dissociation of the effects of bilateral lesions of the dorsal hippocampus and parietal cortex on path integration in the rat, *Behav Neurosci* **115**(6):1212–1223, 2001.

[43] Save E, Nerad L, Poucet B, Contribution of multiple sensory information to place field stability in hippocampal place cells, *Hippocampus* **10**(1):64–76, 2000.

[44] Seth AK, McKinstry JL, Edelman GM, Krichmar JL, Spatiotemporal processing of whisker input supports texture discrimination by a brain-based device, in *From Animals to Animats 8: Proc of the 8th Intl Conf on the Simulation of Adaptive Behavior*, pp. 130–139, 2004.

[45] Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL, A model of the neural basis of the rat's sense of direction, *Adv Neural Inf Process Syst* **7**:173–180, 1995.

[46] Taube JS, Head direction cells and the neurophysiological basis for a sense of direction, *Prog Neurobiol* **55**(3):225–256, 1998.

[47] Thierry AM, Gioanni Y, Degenetais E, Glowinski J, Hippocampo-prefrontal cortex pathway: Anatomical and electrophysiological characteristics, *Hippocampus* **10**(4):411–419, 2000.

[48] Thrun S, Burgard W, Fox D, *Probabilistic Robotics*, MIT Press, 2005.

[49] Touretzky DS, Muller RU, Place field dissociation and multiple maps in hippocampus, *Neurocomputing* **69**(10–12):1260–1263, 2006.

[50] Treves A, Rolls E, Computational analysis of the role of the hippocampus in memory, *Hippocampus* **4**:374–391, 1994.

[51] Witter M, Naber P, van Haeften T, Machielsen W, Rombouts, Barkhof F, Scheltens P, Loper da Silva F, Corticohippocampal communication by way of parallel parahippocampal-subicular pathways, *Hippocampus* **10**(4):398–410, 2000.

[52] Witter M, Wouterlood F, Naber P, Van Haeften T, Anatomical organization of the parahippocampal-hippocampal network, *Ann NY Acad Sci* **911**(1):1–24, 2000.

[53] Wray J, Edelman GM, A model of color vision based on cortical reentry, *Cereb Cortex* **6**(5):701–716, 1996.