The following is a PDF copy of a chapter from this cyberbook
Not all elements of the chapter are available in PDF format
Please view the cyberbook in its online format to view all elements

Animal Spatial Cognition:

Comparative, Neural & Computational Approaches

Edited and Published by

Michael F. Brown
Department of Psychology
Villanova University

and

Robert G. Cook
Department of Psychology
Tufts University

In cooperation with Comparative Cognition Press
of the Comparative Cognition Society
October, 2006

This Cyberbook is Available to All at
www.comparativecognition.org
or
www.pigeon.psy.tufts.edu/asc

Exploration and Navigation Using Hierarchical Cognitive Maps

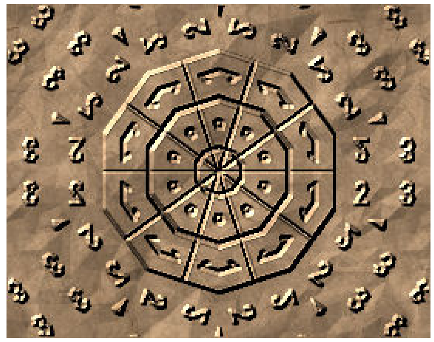

Nestor Schmajuk & Horatiu Voicu
Duke University & University of Memphis

Abstract

We describe a spatial navigation model that uses a hierarchical representation of space. The model is able to explore exhaustively large environments and to plan paths between distant places, even though it uses a working memory with a limited capacity. By using a hierarchical structure, an agent with a limited working memory capacity has (a) a comparable performance to that of an agent with no working memory limitations, (b) a reduced number of connection values and (c) a shorter decision time. The model describes many properties of hierarchical spatial behavior in humans and animals. Potentially, the approach can also be applied to the design of robots able to navigate in large environments.

Chapter Outline & Navigation

I. Introduction

Even if during a purported trip from a house in one city in one continent to an office in a different city in another continent we need to traverse a large number of roads and streets, it is impossible to keep all the locations simultaneously in working memory (WM). Just looking at a map of the East Coast of the United States with street resolution would be a daunting task. Therefore, the trip should be planned first at a high level that tell us which continents are to be examined, then at an intermediate level showing what cities are to be called upon (road maps), and finally at a low level that specifies which streets to take (street maps). In all cases, it is possible that only a part of each type of map can be processed at one time.

From a psychological perspective, the problem is to explain how navigation is accomplished by humans or animals in a large environment, i.e., one that contains a large number of places. From a robotics perspective, the question is how to build a robot that navigates in such large environments. It has been suggested (McNamara, Hardy, & Hirtle, 1989) that humans solve the problem by building a hierarchical cognitive map that contains multiple levels representing the same large environment at different resolutions. In contrast, autonomous robots have difficulty navigating in large environments because path planning algorithms for large environments are slow for real time applications (Nehmzow, 1995).

In this paper, we introduce a neural network model of spatial navigation that incorporates a hierarchical cognitive map. First, we analyze the benefits of using hierarchical spatial representations. Then, we describe the model and present computer simulations that illustrate how it builds and uses the hierarchical map. Finally, we discuss how the model describes human and animal experimental data and compare it to other models.

The need for hierarchical maps.

Voicu and Schmajuk (2002) described a model of spatial navigation that uses a topological representation of the environment structured as a grid (see Figure 1). The cells of the grid represent unique places that the agent can identify and are approximately the size of the agent. Although cells in the grid are assumed to be square, any shape that preserves the continuity of space can be used. The location of the agent in the grid is determined by a localization system which, by making use of the distance to cues as coded by their visual angles (Schmajuk, 1990; Zipser, 1985), determines the agent's location in space. The model assumes that all cells on the grid representing adjacent places are connected even before the agent enters and explores the environment, i.e., it assumes the continuity of space.

The model shown in Figure 2 includes a goal-seeking system with goals set by a motivational system, a cognitive system that has the role of a long-term memory and a WM that has the role of a short-term memory (Atkinson & Shiffrin, 1968). While the cognitive map represents the connectivity between places, the connectivity between places and goals, and the connectivity between goals and places, the WM is used for path planning once the relevant information from the cognitive maps has been stored in it. The information stored in WM is used by the goal-seeking system to produce motor action upon the environment. The model also includes a localization system that tells where in the environment the agent is located and a control system that has the role of supervising the flow of information between the components of the model. The control system can be viewed as an executive system (Baddeley, 1974, 1986) that allocates processing resources to different slave systems.

Exploration.

The system guides the exploration of an unknown environment by assuming that, as suggested by Berlyne (1950), all unexamined places on the grid are goals for the motivational system driven by curiosity. In addition, as mentioned above, we assume that all these points are connected. Therefore, the system will direct the agent to visit all places connected to the exploration goal, decreasing the connection between that goal and each place whenever the place is examined and the agent is no longer curious about it (see Figure 3). Thus the model can perform a complete exploration of the environment.

When all unexamined places initially connected to the exploration goal are examined, exploration ceases. Exploration also ceases when, after certain time using

A.

B.

Figure 1. The Canvas. (A) Squares in broken lines represent the places to be explored. Solid lines represent connections between places. The empty canvas is a lattice representing the potential continuity of the space to be explored. Adjacent places are assumed to be linked and each place is designated as unexamined (represented by an open circle). (B) Numbers of the places in the nonhierarchical case.

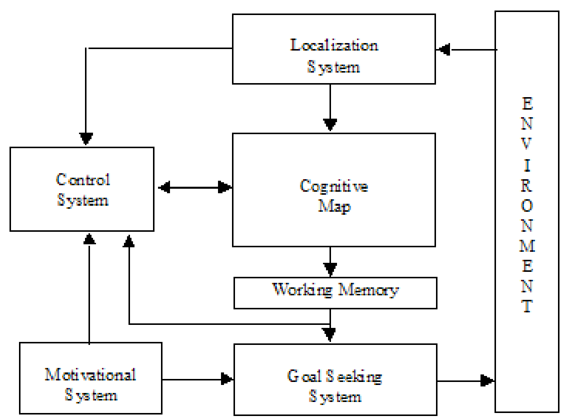

Figure 2. Block diagram of the model. It shows the interaction between the Control System, Localization System, Cognitive map (see Figure 3 for a detailed version), Working Memory, Goal Seeking System and Environment.

the cognitive map (see below), the system cannot decide how to approach an unexamined place. Voicu and Schmajuk (2001a) showed that this curiosity-directed exploration is more efficient, i.e., takes less movements, than random exploration. The system represents the actual structure of the environment by updating the connections between initially connected places in the cognitive map whenever an obstacle is found.

The cognitive map built by the network is a topological map, i.e., it represents only the connectivity, not distance or direction, between places. Associations between place i and place j, $V_{i,j}$, are the elementary internal learned representations of the links in the external world. These associations are stored in modifiable synapses, indicated by open triangles in Figure 3. Whereas a positive $V_{i,j}$ association, means that place j can be accessed from place i, a positive $V_{j,i}$ association, means that place i can be accessed from place j. In both cases, $V_{i,j} = V_{j,i} = 0$ means that each place cannot be accessed from the other. When the agent is in place i, activation a_j is given by $a_j = p_i V_{i,j}$, and this activity indicates whether place j is accessible from place i. When place j, adjacent to place i currently occupied by the agent, is accessed by the agent, then $V_{i,j}$ and $V_{j,i}$ are incrementally increased to a asymptotic value larger than 1. Therefore, in the cognitive map trodden (strengthen) links between places are represented by stronger connections ($V_{i,j} = V_{j,i} > 1$) than unexplored links ($V_{i,j} = V_{j,i} = 1$) or links between places that have been verified to be accessible ($V_{i,j} = V_{j,i} = 1$). A detailed description of the model is given in Voicu and Schmajuk (2002).

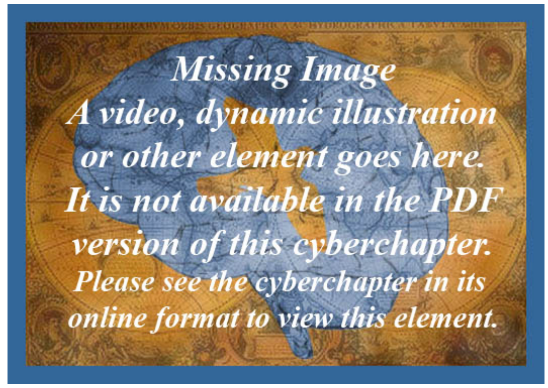

Figure 3. Cognitive System. The cognitive map is implemented by a neural network that stores the links between a place and its neighboring places. The prediction of neighboring place j, p_j, is fed back into the neuron representing place j as a current place. V_{ij}: Association between place i and place j. W_{hi} : Association between goal h and place i. W_{ih} : Association between place i and goal h. g_h: Activation of goal h. Examples of goals are: food, visited places, unvisited places. Arrows: Fixed excitatory connections. Triangles: Variable excitatory connections. *Click on image for animation.*

Navigation.

When an examined place is a goal for a motivation other than curiosity, e.g., when the agent has already found an appetitive stimulus at that location, the agent can use the cognitive map to navigate from its present position to the goal. If the environment has not been explored and mapped before navigation, the system will use its a priori representation and assume that all adjacent places are connected. If the agent does not find these connections during navigation, the map is modified accordingly. If the environment is known, the system will use the map representing the actual structure of the environment.

Either during exploration or navigation, the goal-seeking system will guide the agent to approach a neighboring place in which a goal is found. When the goal cannot be perceived from the current location of the agent, the system will guide the agent to move towards the goal using the cognitive map. There are two methods of using the cognitive map.

One method, referred to as next-place activation method (Voicu & Schmajuk, 2001a), consists of briefly entering all neighboring places in succession, thereby activating their representation in the cognitive map, and spreading this activation forward until the representation of the goal becomes active. In order to make the activation of the goal proportional to the distance (measured by the number of intermediate places from the current position), we assume that this activation is attenuated as it spreads from the representation of one place to another. Therefore, neighboring places that are closer to the goal will activate more strongly the representation of the goal.

A second method, referred to as goal-activation method (Staddon, 2001; Voicu & Schmajuk, 2002), consists of activating the goal in the cognitive map, which will activate the place where it is located, and spreading this activation until the representation of the current location of the agent becomes active. Again, because we assume that this activation is attenuated as it spreads from one place to another, places that are closer to the goal are more active than places that are more removed from it.

Next-place and goal activation procedures are similar in some aspects and different in others. As mentioned, both methods measure the distance between the current position of the agent and the goal by attenuating an initial activation as it spreads over the cognitive map. Therefore, if these two points are separated by a very large number of intermediate places, the attenuated activity will become too weak to be used in the decision process.

The next-place and goal activation methods differ in speed. With the next-place activation method, the system needs to repeat the activation of the cognitive map at each location on its way to the goal, which makes the total decision time proportional to the square of the length of the path (the time required to compute all the moves to the goal is proportional to $l + (l - 1) + . . . 2 + 1 = l\,(l + 1)/2$, where l is the length of the path). Instead, with the goal-activation method, the total decision time is proportional to the length of the path. The information needed to generate the moves to the goal is produced by spreading the activation l times in the network, where l is the length of the path. The next-place and goal activation methods also differ in their memory requirements. The WM required by the next-place activation procedure can be either large or small. The size is large if the WM stores the associations of every place in the environment with the prediction of the goal generated when the agent enters a given place. Alternatively, this size can be relatively small if the WM stores only the associations of (a) the alternative directions in which the agent can move with (b) the prediction of the goal produced when the agent moves in a given direction.

The goal-activation method can plan a complete path between a goal and the current place by storing the activation of all places in the environment in WM and then choosing which place to move next by performing gradient ascent on this activity surface. Once the activation of all places is stored in WM, this information can be used at every location on the way to the goal. Therefore, the goal activation method, though faster and able to plan a complete path from the current position of the agent to the goal, always requires a WM of size equal to the number of places in the environment.

II. The Hierarchical Approach: Path Planning in Large Environments

There are two important advantages that a hierarchical structure provides: (a) a reduced number of connection values and (b) a shorter decision time. Both advantages are the result of storing the same information at multiple scales. The number of connections is reduced because there are no connections between places that belong to different regions (except adjacent places in adjacent regions), and decision time is improved because coarser representations of the environment involve fewer nodes and therefore faster decision times. More details are presented in appendices A and B. Appendix A shows that the use of a hierarchical cognitive map reduces the number of connection values in a nonhierarchical map. Appendix B shows that the decision time is shorter in a hierarchical than in a nonhierarchical cognitive map.

In this section we address the question of how path planning can be accomplished in a large environment when two restricting factors are considered: (a) a limited WM capacity and (b) an adequate attenuation when a spreading activation method is used.

With a WM of limited capacity.

As shown, the goal activation method, which is relatively fast and able to plan a complete path to the goal, always requires a WM of size equal to the number of places in the environment. It has been suggested that WM capacity is relatively small, ranging between five and nine items in humans (Miller, 1956). Therefore, an important question is how path planning in a large environment can be accomplished using a limited WM capacity. McNamara, Hardy and Hirtle (1989) suggested that organisms that move in large-scale environments should have their spatial memory structured in units compatible in size to that of the WM. How does this idea apply to the system described above?

As mentioned, planning a complete path from the current position of the agent to the goal can be achieved using the goal-activation method, which requires a WM of size W equal to the number of places N in the environment ($W = N$). If the capacity of the WM is smaller than the number of places N ($W < N$), path planning cannot be completed. This happens because the goal-seeking mechanism needs the activation of both the current place and the goal to be in WM in order to produce motor actions. Consequently, one way to achieve path planning between two places is by dividing the environment in a number of regions n that is equal to or smaller than W ($W \geq n$). Because $n < N$, the size s of each of the n regions will be larger than the size S of each of the N places (which are approximately the size of the agent), that is, $s > S$. Finally, the number of places of size S in each region of size s is given by $m = s / S$.

Even though it is possible to plan a path in an environment divided in $n = W$ regions, these regions are larger than the size of the agent and therefore inadequate to exactly define its spatial location. Consequently, once the path through the n regions is defined, we need to define the paths through the places inside each of these regions. As described above, planning a complete path between two points inside a region requires a WM of size W equal to or greater than the number of places m inside a region ($W \geq m$). If $W < m$, then another division of each region will be necessary. If the environment is represented by a hierarchy with L levels then the total number of places covered is given by $N = W^L$ Therefore, the number of levels L in which an environment of N places should be divided in order to fit a WM of capacity W, is given by $L = \log N / \log W = \log 36 / \log 6 = 2$ (for an example with W=6, N=36 see Appendix A and Figure 4).

A.

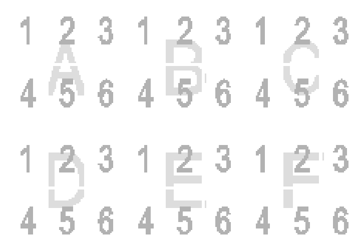

B.

Figure 4. (A) Canvas for a two level hierarchy. The environment is divided in 6 regions (A, B, C, D, E, and F) and each region is divided in six places. The advantage of this representation is that a WM of size 6 is enough for planning paths in the environment. (B) Squares in broken lines represent the regions to be explored. Solid lines represent connections between regions.

With a spreading activation method.

Even if there were no potential limitations to WM capacity, the use of both the goal activation and the next-place methods is still problematic. Both methods are required to spread activation over a limited number of places in order to avoid excessive attenuation of the spreading activity, thereby restricting how far the goal can be from the location of the agent. It can be argued that this problem can be avoided if the attenuation rate, i.e., the attenuation at each reinjection in the cognitive map, is small enough. However, small attenuation rates can result in poor discrimination between two alternative next places when the goal is far from the location of the agent. This poor discrimination might be caused by a baseline noise that is larger than the difference between activities at adjacent places far removed from the present position of the agent.

Therefore, a compromise should be reached by combining a minimum attenuation rate with a maximum number of reinjections that still generates a significant activity at the current location of the agent when the goal is activated. Let this maximum number of reinjections be A. Consequently, in order to obtain a useful activity from the most distant confines of the environment, we should divide the environment in such a way that in the worst case the agent is A regions, and A reinjections, away from the goal.

Whereas A is limited directly by the physical implementation of the spreading activation process, the limited WM capacity could be either a direct result of the limitation of A or a consequence of how the agent is designed.

III. The Hierarchical Cognitive Map

Here we introduce a hierarchical version of the cognitive map presented by Voicu and Schmajuk (2002) and described in the section on The Need for Hierarchical Maps. In the hierarchical version, spatial information is structured on a two-level hierarchy (see Figures 5 and 6). Each level is represented by an associative network (Kohonen, 1977; see Figure 3). The first level, or lower level map (LLM), has a high resolution (locations are relatively small), which permits the representation of all environmental details (such as obstacles). For each region, the LLM stores associations between places and places, places and goals, and goals and places. Associations between places represent possible transitions between spatial locations, associations between goals and places define where a given goal is located, and associations between places and goals define what goals can be found in a given location. This LLM is identical to the nonhierarchical cognitive map described before except that it has fewer connections due to the lack of interregional connections (see Appendix A).

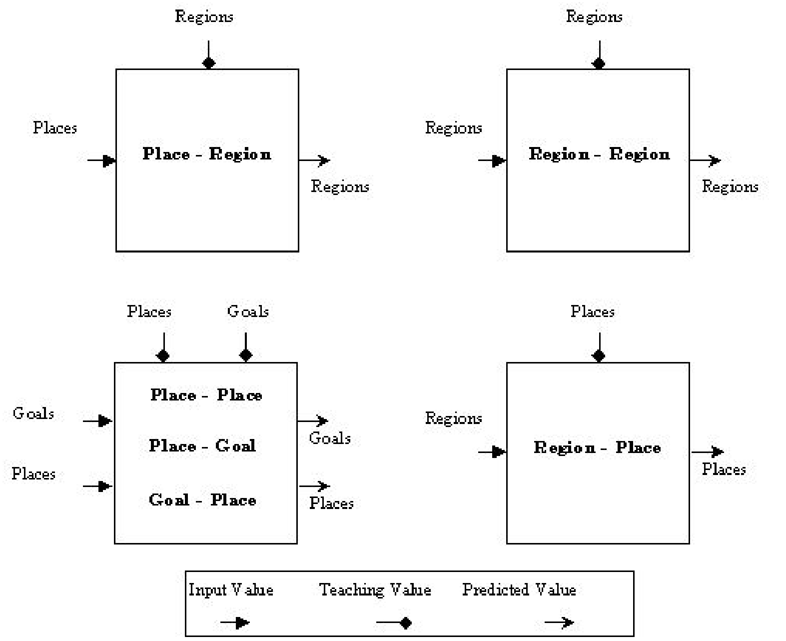

Figure 5. Memory modules used in the Hierarchical Cognitive Map. The place-region, region-region, place-place, place-goal, goal-place, and region-place connections are stored in four associative networks. These networks and the flow of information among them are manipulated by the control system. This figure shows only the inputs and outputs of these memory modules. Connections between modules are shown in Figure 6. Appendix C shows the equations used for updating the associations and outputs in each of the blocks presented above.

The second level, or upper level map (ULM), has a low resolution (relatively large locations) that permits the representation of the whole environment. The ULM stores associations between regions and regions. Again, associations between regions represent possible transitions between them.

Two additional associative networks (see Figure 5) are needed to represent the associations between places and regions (defining the region where a place is located) and between regions and places (defining the places contained in a given region).

The control system acts on the four associative networks (place-region, place-place, region-region, and region-place) by directing the flow of information between them and between the networks and the goal-seeking system (see Figures 2 and 6). This control system determines how the networks are used during exploration and navigation.

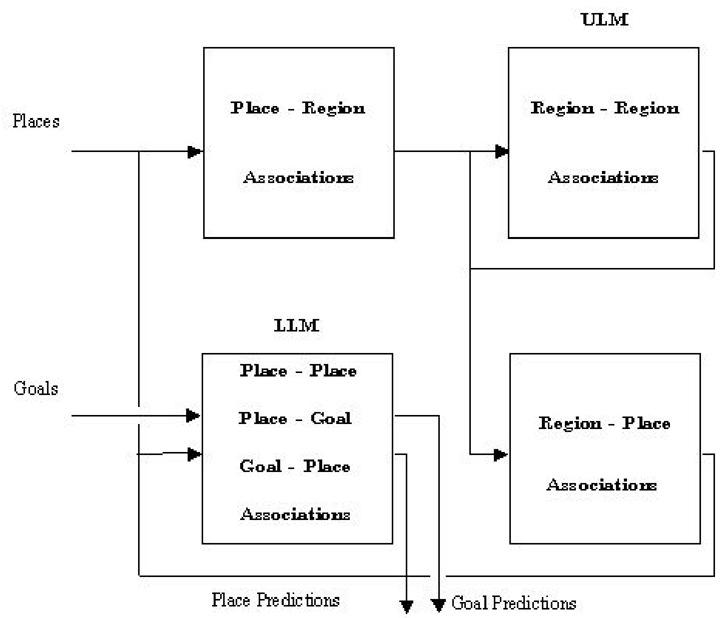

Figure 6. Hierarchical cognitive map with two levels, lower level map (LLM) and upper level map (ULM). This figure shows the connections among the four memory modules that allow the control system to read the hierarchical structure and plan paths in the environment. Connections related to the update of the memories are not shown.

Exploration and the updating of the hierarchical map.

As in the case of nonhierarchical maps described in the section on The Need for Hierarchical Maps, in order to build a hierarchical map of the environment the agent has to examine each place in the environment and enter the result of the inspection in the cognitive map.

As shown in Figure 4, the agent is provided with a priori information. (1) The environment to be explored contains a partition of disjoint (non overlapping) regions that covers it completely. (2) These regions are initially interconnected and preserve the continuity of space. (3) In turn, each region contains a partition of disjoint places that covers it completely. (4) These places are initially interconnected and preserve the continuity of space. For simplicity, it is assumed that regions and places are rectangular; however, as in the nonhierarchical case, any type of shape is suitable as long as it preserves the continuity of space. In addition, it is assumed that any two adjacent place fields are connected and any two adjacent regions are connected.

In order to completely explore the environment, the motivational system defines as goals for exploration all the unexamined regions in which the environment is divided. Alternatively, one or a few regions can be defined as goals. In turn, all the places in the region(s) to be explored are goals for exploration. Because the agent can be surrounded by unexplored regions or places, the decision of which region or place to enter first is made according to a set of priorities. In our case, these are North, West, East, North-West, North-East, South, South-West, and South-East.

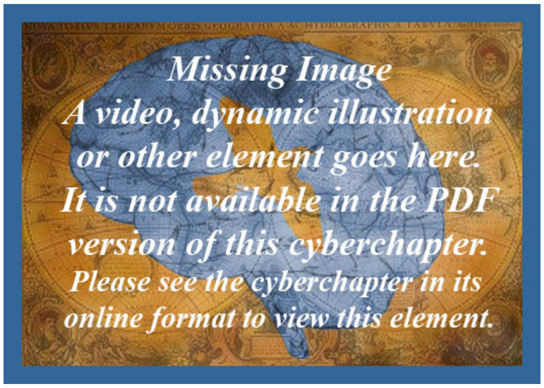

Figure 7. Exploration. The model can exhaustively explore a large environment by applying the exhaustive exploration strategy at all levels of the hierarchy. All places are equally attractive at the outset.

Exploration of a region.

If obstacles are found in the explored region, then the LLM that represents the connectivity between places is updated according to a Hebbian rule (see Appendix C). The control system knows that a region has been completely explored, i.e., all places have been examined, when all place-exploratory goal associations are equal to 0. After exploration, a region is defined as composed by all places in a region that can be examined without leaving the original region.

Several examples of how different types of regions are explored are described in Appendix D. Appendix E describes how the a priori shape of a region and its a priori connections can be changed according to how obstacles are positioned in the environment (region redefinition). If an obstacle divides a region in two, then during the exploration of that particular region the model creates two smaller new regions and connects them to the hierarchical structure. The regions created are disjointed and separated by the obstacle. This division process can continue an arbitrary number of times so that the map adapts to how obstacles are positioned in the environment.

Exploration of the environment.

Figure 8. Path planning in a known environment. Starting and ending positions are A1 and F4 respectively (black circles). Obstacles are shown in gray. The arrows show the trajectory of the agent within each region. Since the layout of the environment is known the agent is building an optimal path. *Click on image for animation of process shown in Figures 8, 9, & 10.*

The ULM is updated only if the newly found obstacles prevent the agent from reaching the next region to be explored. The same Hebbian rule, described above, used for updating the connections between places within a region in the LLM is used for updating the connections between regions in the ULM (see Appendix C). Connections between regions are assumed to be weaker than connections between places within a region.

Figure 7 illustrates the case of an environment to be explored. Exploration of the regions is done following the priorities described before, so if the agent starts at place A1, it will move to places A2, A3, A6, A5, and A4. From region A the agent will move to regions B, C, F, E, and D, where it will attempt to visit the places in the order followed when exploring region A.

Finding a path between two places in an explored environment.

Once exploration is complete, and both ULM and LLM accurately reflect the structure of the environment, the agent is ready to navigate between any two points in the environment. If the place where the agent is currently located and the goal are in the same region (a fact detected by spreading the activation from the goal a fixed number of times and succeeding in activating the current location of the agent), then only the LLM of that region is used for path planning. Instead, if as illustrated in Figure 8, they are in different regions (a fact detected by spreading the activation from the goal a number of times proportional to the size of WM and failing to activate the current location of the agent), then planning starts in the ULM.

Figure 9. Upper Level Map. (A) Planning done at the ULM for traveling from region A to region F with a priori connections. (B) New planning once the connection between B and F has been removed.
Click on image for animation of process shown in Figures 8, 9, & 10.

Figure 8 illustrates the case in which the agent is initially located in place A1 and the goal in place F4. As described below, the execution of each step in ULM is accompanied by path planning at LLM. Whenever a subgoal at the local level is reached the system proceeds with the next step at the ULM.

Step 1. Lower Level Map. Because the starting place and the goal are not in the same region the control system plans a path at the ULM.

Step 2. Upper Level Map. The agent combines the present region A with the goal region F in the ULM and plans the upper level path (ULP) between them, ULP = A, B, C, F, using the activity surface in WM (see Figure 9B).

Step 3. Lower Level Map. Once the path is planned at the Upper Level, the detailed navigation at the Lower Level is planned and achieved. As before, starting at place A1, the goal for the agent is the closest place that belongs to B, that is B1. This is accomplished by activating all places in region B, and spreading this activation over the network, including places A3 and A6. Because the resulting activation surface shows similar activity levels for places A3 and A6, the agent uses the above priorities to determine lower-level path (LLP), LLP = A1, A2, A3, B1. This path is stored in WM and is overwritten the next time the system plans a lower level path.

Step 4. Lower Level Map. The agent is initially at place B1. In order to move from

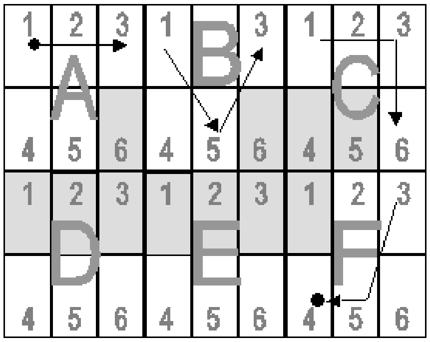

Figure 10. Path planning in an unknown environment. The arrows show the trajectory of the agent within each region. Since the layout of the environment is not known the agent has to detour at the border between region B and F. *Click on image for animation of process shown in Figures 8, 9, & 10.*

region B to region C, the agent plans the LLP = B1, B2, B3, C1.

Step 5. Lower Level Map. The agent is initially at place C1, and the goal, defined by the next region in the upper level path B, C, F, is the closest place that belongs to F, that is F3. Due to its navigational priorities defined above, the agent will plan the LLP = C1, C2, C3, C6, and F3, and move through it.

Step 6. Lower Level Map. The agent is initially at place F3 and the goal is F4. Because activation of the goal activates the representation of the current place, the LLM is used. The LLP planned and followed is LLP = F3, F5, F4. Because the last region, F, has been reached, the task has been completed.

Finding a path between two places in an unexplored environment.

The agent can navigate between any two points in the environment even when exploration has not been completed and both ULM and LLM just reflect the a priori assumptions about the structure of the environment. Figure 9A illustrates the ULM with six regions. As in the previous case, the agent has to move from place A1 in region A to place F4 in region F.

Figure 10 illustrates the representation of the complete environment, gray squares indicating obstacles, showing place 1 in region A (place A1) and place 4 in region F (place F4). Unlike the previous case, the model plans the following path at the ULM: A, B and F (see Figure 9A). Once the agent is in place B5 it updates the LLM and the ULM, and path planning starts again. The agent moves from B5 to C1, then from C1 to B3 and finally reaches F4.

In the case of navigation between two points in the environment, the continuity of each region is not critical. If a region is actually composed of two or more subregions, the agent will find that the predicted path is not executable, disconnect the regions in the ULM, and plan an alternative path.

Computer simulations.

In our simulations, we assumed that the capacity of the WM is W= 1120 and the environment is divided into 30 regions. Therefore, the agent navigates in an environment that contains N = 33,600 places. All places are of approximately the size of the agent. The simulated environment is based on the topography of the Duke University West Campus. It includes 12 buildings and their surrounding areas. All buildings have their entrances open and no details about the inside structure is included. Only the outside walls of the buildings appear as obstacles in the simulated environment.

First, we let the agent perform a complete exploration of the environment and then plan paths between different locations in the environment. Figure 11 shows the result of a complete exploration of the environment. The grid shown in thin black lines illustrates the initial division of the environment. Regions created during exploration contain dotted lines and are delineated by gridlines and obstacles.

One important aspect of exhaustive exploration is the inability to examine all places in one sweep. This is most likely to happen in environments that contain many cul de sacs. In the case of an agent that does not use a hierarchical representation of space this problem is acute since the agent has no means of detecting if there are unexamined cul de sacs close to its current position. This happens because the agent always moves towards the closest unexamined place with respect to the whole environment. In contrast, an agent that uses a hierarchical representation of space always moves to the closest unexamined place with respect to the current region. Therefore, unvisited cul de sacs that are in the current region will be visited even if they are further away than unexamined places located in an adjacent region. In the computer simulations shown in Figure 11, with few exceptions, all cul de sacs (i.e. all buildings) are completely visited in one sweep.

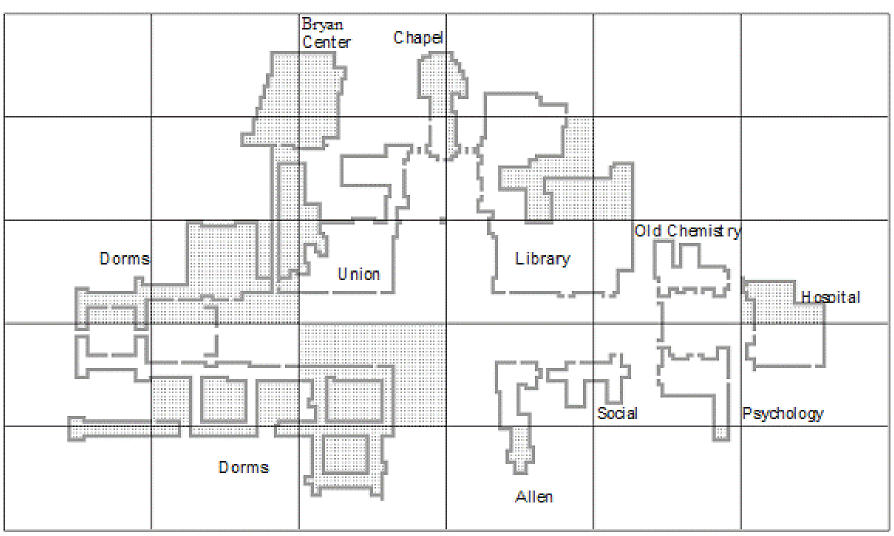

Figure 11. Exhaustive exploration. Gray lines represent either obstacles or building contours. The grid represents the initial division of space. The map shows a schematic representation of Duke West Campus. Regions that contain dotted lines are created during exploration. *Click on image for animation of process shown in Figures 11, 12, & 13.*

After performing a complete exploration of the environment we let the model plan paths between different locations. Figure 12 shows the path between two places that belong to regions created during exploration. The starting place is located in the Psychology building and the goal is located in Bryan Center. Figure 13 shows two paths between places that belong to regions defined a priori. The first path Start1-Goal1 crosses the campus from North-West to South-East and the second path crosses the campus from South-West to North-East. It can be noticed that the agent takes advantage of shortcuts contained inside regions.

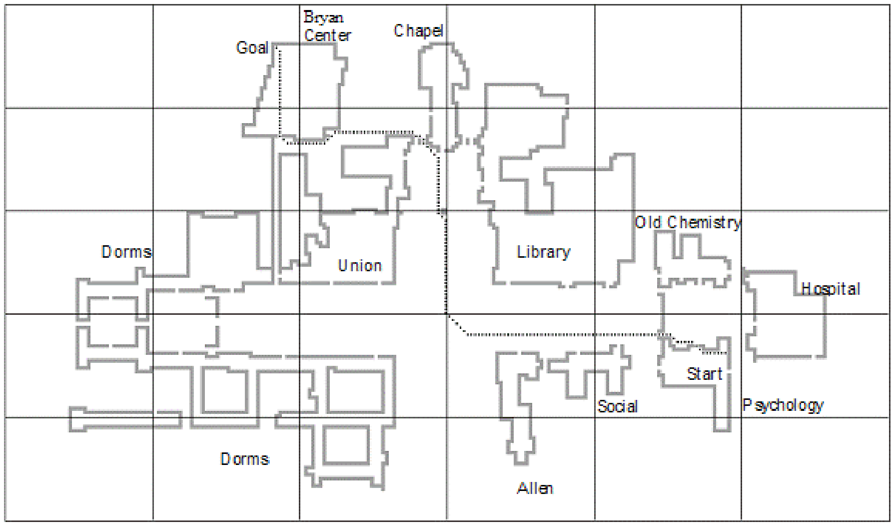

Figure 12. Path planning. Gray lines represent obstacles. The grid represents the initial division of space. The map shows a schematic representation of Duke West Campus. Dotted line represents the path planned by the agent between Psychology building and Bryan Center. *Click on image for animation of process shown in Figures 11, 12, & 13.*

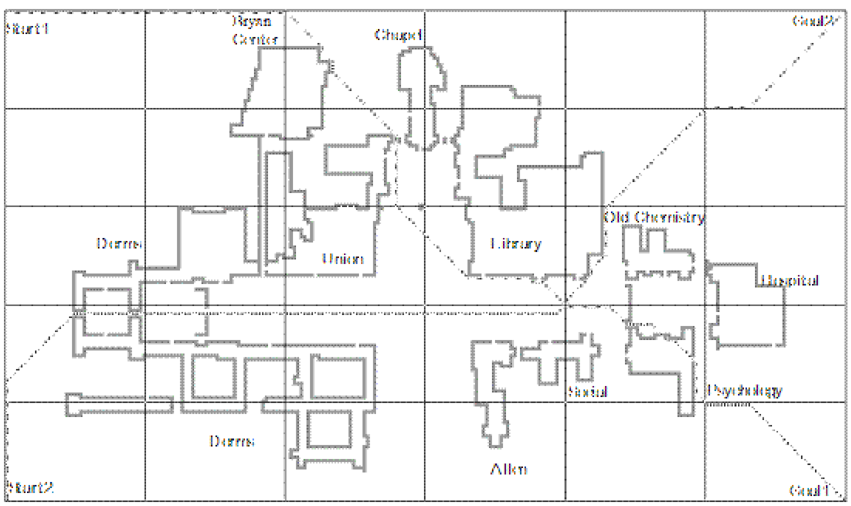

Figure 13. Path planning. Gray lines represent obstacles. The grid represents the initial division of space. Dotted lines represent paths between pairs of places Start1-Goal1 and Start2-Goal2. *Click on image for animation of process shown in Figures 11, 12, & 13.*

IV. Discussion

This chapter shows that agents with a limited WM capacity and that use a hierarchical representation of space, can plan navigation and exploration of large environments. Agents plan navigation by spreading activation between the representations of two places in the cognitive map and storing the resulting activities in a WM. Then, they perform gradient ascent on the activity surface stored in WM. Agents explore the environment by setting as goals places that have not been examined before. During exploration agents update their internal representation (all levels of the hierarchy) to reflect the spatial layout of the environment. Computer simulations show how the model is applied to an environment with 33,600 places.

Comparison with other spatial representations.

Our model represents a given environment by determining the occupancy at points equally spaced on a grid. The space between points is equal to the size of the agent or that of the smallest object in the environment. This approach has been used both in building robots (e.g. Schultz et al., 1999) and developing models of animal navigation (e.g. Reid & Staddon, 1998; Trullier & Meyer, 1998). If, instead of being equally spaced, the points on the grid are unequally separated, then the resulting spatial structure can take advantage of the size of the occupied or unoccupied areas. Unequally spaced grid structures save memory when large, either occupied or unoccupied regions are present (Samet, 1980; Hirtle, 1995). These type of structures have also been used in robots (e.g., Arleo, Millan, & Floreano, 1999).

Another type of spatial representation results from partitioning the environment into disjoint parcels. The shape and the location of each partition are determined by the angles and distances from the agent to salient features of the environment. Experimental data shows that place cells in the rat hippocampus provide a partition of the environment that has fields approximately two and a half times larger than the size of the animal (O'Keefe & Dostrovsky, 1971). The firing pattern of the place cells seems to be modulated in part by the positions of cues and landmarks external to the maze in which the subjects perform. When compared to the grid methods defined above, the partition method lacks rules to generate a unique division of space. This method has been used both in robots (Bachelder & Waxman, 1994) and models of animal navigation (Sharp, 1991). Vectorial representations describe the environment by making use of geometric objects, such as points, circles or polygons (Latombe, 1991). Each of these objects contains a vector that defines its position and a set of mathematical equations that define its shape. The most important advantage of this type of representations is the ability to use a variety of geometrical operators, such as magnification, rotation, and translation, to manipulate the spatial representation. Its most important drawback is the difficulty of constructing the objects directly from sensory data (Dudek & Jenkin, 2000). These methods have been used in robots (Latombe, 1991) and models of animal navigation (Cartwritght & Collet, 1983).

Grid representations are better suited than vectorial representations for exhaustive explorations of an environment. If an agent uses the second approach, then only salient objects would be coded and the space between objects would not have an internal representation. Therefore, the agent will probably fail to explore some part of the space between objects. On the other hand, if the agent uses a grid, every place the size of the footprint of the agent or of a small object would have an internal representation.

In contrast to the metric representations (grid, partition, or vectorial) mentioned above, topological representations describe only the connectivity, not direction and distance, between places. Since path planning on a topological structure is efficient, many models use this type of representation (Chown, Kaplan, & Kortenkamp, 1995; Mataric, 1991; Schmajuk & Thieme, 1992).

Because our model describes space by the occupancy of points on a grid and, in addition, defines the connectivity between these points, it provides information about the topological aspects of the environment. Furthermore, as shown by Voicu and Schmajuk (2002), even if metric data are not explicit in the topological representation of our model, the representation contains implicit metric information because nodes that are further away in its structure are also further away in the real environment. Other models also combine metric and topological information (e.g., Kuipers, 2000).

How the model explains experimental data.

Supporting the idea that humans use hierarchical spatial representations, Stevens and Coupe (1978) reported that people misjudge the relationship between geographical locations or artificial stimuli based on super-ordinate spatial relations. For example, because most of the US is south of Canada, Seattle is usually judged to be south (and west) of Montreal even though it is actually north (and west). According to our model, in the absence of information about the relationship between two places that belong to two different regions (e.g., Seattle and Montreal), the agent relies on the relation between the regions where those places are located (e.g., US and Canada). The latter relation, however, might contain inadequate information since the relationship between regions might be different from the relationship between places. For example, in Figure 7, place A1 is north of place B4. However, in the absence of information about their relative location and based on the location of regions A and B, A1 and B4 are judged to be at the same latitude.

By adapting the ordered-tree algorithm (Reitman & Rueter, 1980) to detect underlying organization of spatial memory, Hirtle and Jonides (1985) showed that the spatial representation in humans is organized hierarchically (see Taylor for more information on human spatial memory). Their subjects were tested on free recall, map drawing and relative distance judgment regarding a region in which subjects have been living for at least 2 years. The study shows that, when the points belong to different clusters, physical distances tend to be estimated as longer than that estimated when the points belong to the same cluster (Figure 5 in Hirtle & Jonides, 1985).

In its present form our model can eventually explain Hirtle and Jonides's results by assuming that associations between places within a region are stronger than those between places located in different regions. For example, if there are three places (two located in the same region and one located in a different region) that have the same physical distance between them, the activity that spreads between the places that belong to the same region is attenuated less than the activity that goes between two places that belong to different regions. Since, as mentioned in the section on The Need for Hierarchical Maps, our model evaluates the distance between two places by the level of activation of one place after activity is spread in the network from the other place, places within one region will be judged to be closer than places across regions (for a different approach see Voicu, 2003).

Still another study, conducted by McNamara et al. (1989) using human subjects, showed that spatial memory is hierarchically structured even when objects are uniformly distributed on a map (see Taylor for more information on human spatial memory). The study suggests that hierarchies are present even when there are no physical and perceptual boundaries in the layout of the environment. This hierarchical organization might be implicit in the structure of the human visual system, where external images are represented in a number of different resolutions (Watson & Robson, 1981). This evidence might supports our assumption that the environment is partitioned a priori in a hierarchical structure that has different resolutions at different levels.

Further evidence for a hierarchical representation of spatial information has been provided by Holding (1994). Subjects learned an imaginary map that contains houses, streets and towns. An analysis using the ordered-tree algorithm shows that the hierarchical structure groups houses together and streets together. Similarly, in our model streets could be represented at the same level in the hierarchical map. Moreover, the experimental results suggest that these hierarchical spatial representations support semantic priming. Our model can explain priming in terms of one item activating the representation of a second one in the cognitive map and, therefore, decreasing response time to the second item.

The animal literature also provides evidence that suggests that spatial memory in animals uses multiple representations with different resolutions. For example, Cartwritght and Collet (1983) concluded that bees navigate using both contours of objects (located within tens of meters) and tree lines (located within hundred of meters). It is unclear, however, if these representations are combined in a hierarchical way or if, instead, bees use the objects when close to them and the tree lines when far from the objects.

Other experimental studies (Dallal & Meck, 1990; Fountain & Rowan, 1995; Macuda & Roberts, 1995; Meck & Williams, 1997; Roberts, 1979;

Terrace, 1991) suggest that animals chunk information and form hierarchical representations to facilitate learning. For example, experimental data (Roberts, 1979) suggest that when rats are tested in a hierarchical radial arm maze (a regular radial arm maze that has three additional small aisles at the end of each arm), they use separate memories for primary and secondary alley choices. In our model, the secondary alleys and the central platform can be seen as regions, each one represented by its own LLM. In turn, regions are connected by the primary alleys, and this spatial structure is represented by a ULM. (See Mizumori for possible novel representations.)

Comparison with other hierarchical cognitive maps.

Like the model presented in this paper, several approaches to spatial navigation take advantage of hierarchical spatial representations for planning action in a given environment. Niizuma et al. (1992) proposed a hierarchical system in which each control level uses its own map of the environment. The hierarchical structure contains a connection map, a geometrical map, a local map and a sensing map. Only one level of the hierarchy is active at one time and, when one level fails to accomplish the navigational task, the control is passed to a higher level. Whereas in the Niizuma et al. (1992) model each control level uses its own map of the environment, in our model the control system has access to all levels in the hierarchical structure.

Another approach to hierarchical mapping has been proposed by Donnart and Meyer (1996). Their model uses a production rule system to build a hierarchical map of the environment that includes not only spatial information, but also information related to spatial tasks. The model learns a hierarchical organization of landmarks that helps in landmark recognition. The main benefit of using this type of hierarchy is that the localization system is more robust by having additional contextual information. Issues related to the size of the explored environment are not addressed by the authors.

Car (1998) proposed an approach to hierarchical mapping that describes a method for finding a minimum path between places in a large environment using a hierarchical structure. In this hierarchical structure, the environment is divided in regions that contain exit and entry points. These points are used as goals when moving from one region to another. The method first determines the level of the hierarchy that includes both the goal and the current position. For this upper level, the method finds the shortest path using Dijkstra's (1959) algorithm. Then, using the same algorithm at the lower level, the method finds paths between the current position and the exit point in the first region, the entry point and the exit point in the second region, and so forth. Because the method uses fixed entry and exit points, the path generated combining the upper and lower levels is seldom identical to the shortest path generated when Dijkstra's algorithm is directly applied to the lower level. However, these results are improved when several adjacent regions are collapsed into a single one.

Car's approach is similar to ours in the way spatial information is represented and used; however, there are two important differences. First, our approach addresses the issue of how the hierarchical map is updated. Second, because in our method any point on the border between regions can be an exit or an entry point, the path generated combining the upper and lower levels better approximates the path generated when path planning is directly applied to the lower level.

Still another approach that uses hierarchical structures is the Spatial Semantic Hierarchy (SSH) proposed by Kuipers (2000). The SSH comprises five levels, including sensory, control, causal, topological, and metric. The sensory level provides the sensory information in a readable form for all the other levels. Basically, variables like distance, intensity of light and sound are converted in an appropriate format. The control level describes the environment by using control laws that are associated with conditions for their utilization. A particular law (e.g. follow wall, approach landmark) is used within a segment of the environment. The causal level summarizes the environment as a discrete model in terms of sensory views and actions and causal relations among them. The topological level contains the connections among places, paths and regions in the environment. Finally, the metric level contains a geometrical map of the environment using a single frame of reference. This type of information is seldom useful but it can prove essential in certain scenarios.

Kuipers (2000) argues for the advantage of a hierarchical representation of the topological level versus a linear representation. One concern is, as mentioned before, that a hierarchical representation usually produces a path that is far from the optimal solution. However, even if hierarchical representations are not able to provide the optimal path, they can save time by providing a computationally inexpensive suboptimal path, which sometimes is more important than traveling on the shortest path.

A hierarchical structure has also been used by Graham, Joshi, and Pizlo (2000) to build a model that solves the traveling salesman problem (TSP). The TSP is the problem of finding the shortest tour that connects a number of cities that a salesman is interested in visiting. Graham et al. (2000) compared the performance of human subjects, their model and well-known algorithms in the artificial intelligence (AI) literature on TSP. Their results showed that, unlike the AI algorithms, human subjects and the hierarchical model show a linear dependency between the size of the problem, i.e. number of cities, and the error of the solution, that is the difference between the solution found by the human subject or the computational model and the actual optimal solution (the shortest tour).

Graham et al. (2000) represent the TSP as a high-resolution grid where cells that encode cities have values greater than zero whereas all other cells have a value of zero. By using a process of filtering and compression a succession of smaller grids are produced. Each grid looks like the initial grid only it has a different resolution. All these grids are used to build step by step a solution to the TSP. First, a partial solution of the problem is produced at the grid with the lowest resolution. Then, this solution is expanded by including partial solutions taken from the grids with higher resolutions.

Notice that although the representation used by Graham et al. (2000) is very similar to ours, the information stored in it is used for different purposes. Whereas their representation builds to solve a particular problem, namely the TSP, our representation is used to explore and navigate large environments.

V. Summary and Conclusion

Voicu and Schmajuk (2001a) presented a neural network model of spatial navigation that includes an action system and a cognitive map. By spreading activation between the representations of two places in the cognitive map and storing the resulting activities in a WM, the action system can plan the shortest path between those places. However, the capacity of the WM might be too small to store the activity of all intermediate places or the attenuation might be too large to be used in the decision process when navigating in a large environment. These restriction call for the use of a hierarchical cognitive map.

In the hierarchical map, the environment is represented at multiple levels. At the highest level, the environment is divided in a number of parts equal to the size of WM. At the lowest level, each part of the previous level is divided in parts equal to the size of the agent. Between the highest and the lowest level, each part of the previous level is divided in a number of parts equal to the size of the WM. Path planning starts at the level that contains the points between which navigation is desired, and ends at the lowest level, at which motion is produced.

Computer simulations show that a model, with a WM capacity smaller than the total number of places in the environment, explores, creates new regions and navigates between places in the environment. The model describes some, but not all, of the properties of hierarchical spatial behavior in humans and animals. Potentially, the approach can also be applied to the design of robots able to navigate in large environments.

VI. References

Arleo, A., Millan, J., & Floreano, D. (1999). Efficient learning of variable-resolution cognitive maps for autonomous indoor navigation. *IEEE Transactions on Robotics and Automation, 15*, 990-1001.

Atkinson, R.C., & Shiffrin, R.M. (1968). Human memory: A proposed system and its control processes. In K.W. Spence and J.T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2, pp. 89-195). New York: Academic Press.

Bachelder, I.A., & Waxman, A.M. (1994). Mobile robot visual mapping and localization: A view based neurocomputational architecture that emulates hippocampal place learning, *Neural Networks, 7*, 1083-1099.

Baddeley, A. (1986). *Working memory*. Oxford: Clarendon Press.

Berlyne, D.E. (1950). Novelty and curiosity as determinants of exploratory behavior. *British Journal of Psychology, 41*, 68-80.

Car, A. (1998). Hierarchical spatial reasoning: A geocomputation method. *Proceedings of the 3rd International Conference on GeoComputation, 17-19 September, Bristol*.

Cartwright, B.A, & Collett, T.S. (1983). Landmark learning in bees: Experiments and models. *Journal of Comparative Physiology, 151, 5*21-543.

Chown, E., Kaplan, S., & Kortenkamp, D. (1995). Prototypes, location, and associative networks (PLAN): Toward a unified theory of cognitive mapping. *Cognitive Science, 19*, 1-51.

Dallal, N.L., & Meck, W.H. (1990). Hierarchical structures: Chunking by food type facilitates spatial memory. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 69-84.

Dijkstra, E.W. (1959). A note on two problems in connection with graphs. *Numeriske Mathematik, 1*, 269-271.

Donnart, J.Y., & Meyer, J.A. (1996). Hierarchical-map building and self-positioning with MonaLysa. *Adaptive Behavior, 5*(1) 29-74.

Dudek, G., & Jenkin, M. (2000). *Computational principles of mobile robotics*. Cambridge University Press.

Fountain, S.B., & Rowan, J.D. (1995). Coding for hierarchical versus linear pattern structure in rats and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 21*, 187-202.

Graham, S.M., Joshi, A., & Pizlo, Z. (2000). The travelling salesman problem: A hierarchical model. *Memory and Cognition. 28*(7), 1191-1204.

Hirtle, S.C. (1995). Representational structures for cognitive space: Trees, ordered trees and semi-lattices. *Lecture Notes in Computer Science, 988*, 327-340.

Hirtle, S.C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory and Cognition, 13*(3), 208-217.

Holding, C.S. (1994). Further evidence for the hierarchical representation of spatial information. *Journal of Environmental Psychology, 14*(2), 137-147.

Kohonen, T. (1977). *Associative memory*. Berlin, Germany: Springer Verlag.

Kuipers, B. (2000). The spatial semantic hierarchy. *Artificial Intelligence, 119*, 191-233.

Latombe, J.C. (1991). *Robot motion planning*, Kluwer Academic Publishers.

Macuda, T., & Roberts, W.A. (1995). Further evidence for hierarchical chunking in rat spatial memory. *Journal of Experimental Psychology: Animal Behavior Processes, 21*, 20-32.

Mataric, M.J. (1991). Navigating with a rat brain: A neurobiologically-inspired model of robot spatial representation. In *From Animals to animats* (Vol. 1, pp. 169-175). Cambridge, MA: MIT Press.

McNamara, T.P., Hardy, J.K., & Hirtle, S.C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 211-227.

Meck, W.H., & Williams C.L. (1997). Perinatal choline supplementation increases the threshold for chunking in spatial memory, *Neuroreport, 8* (14), 3053-3059.

Miller, G.A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63*, 81-97.

Miller, G.A., Galanter, E., & Pribram, K.H. (1960). *Plans and the structure of behavior*. New York: Holt.

Nehmzow, U. (1995). Animal and robot navigation. *Robotics and Autonomous Systems, 15*, 71-81.

Niizuma, M., Kawano, Y., Tomizawa, M., Sugiyama, M., & Degawa, S. (1992). Flexible control of autonomous vehicle utilizing hierarchical map and planning, *Proceedings of the IMACS/SICE International Symposium on Robotics, Mechatronics and Manufacturing Systems, Japan.*

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research, 34*, 171-175.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford University Press.

Reid, A.K., & Staddon, J.E.R. (1998). A dynamic route finder for the cognitive map. *Psychological Review, 105*, 585-601.

Reitman, J.S., & Rueter, HR. (1980). Organization revealed by recall orders and confirmed by pauses. *Cognitive Psychology, 12*, 554-581.

Roberts, W.A. (1979). Spatial memory in the rat on a hierarchical maze. *Learning and Motivation, 10*, 117-140.

Samet, H. (1980). Region representation: Quadtrees from boundary codes. *Communications of the ACM, 23*(2), 163-170.

Schmajuk, N.A. (1990). Role of the hippocampus in temporal and spatial navigation: An adaptive neural network. *Behavioural Brain Research, 39*(3), 205-229.

Schmajuk, N.A., & Thieme A.D. (1992). Purposive behavior and cognitive mapping: An adaptive neural network. *Biological Cybernetics, 67*, 165-174.

Schultz, A.C., Adams, W., & Yamauchi B. (1999). Integrating exploration localization and planning with a common representation. *Autonomous Robots, 6*, 293-308.

Sharp, P.E. (1991). Computer-simulation of hippocampal place cells. *Psychobiology, 19*(2), 103-115.

Staddon, J.E.R. (2001). *Adaptive dynamics: The theoretical analysis of behavior*. Cambridge, MA: MIT Press.

Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations, *Cognitive Psychology, 10*(4), 422-437.

Terrace, H.S. (1991). Chunking during serial-learning by a pigeon. *Journal of Experimental Psychology: Animal Behavior Processes, 17*(1), 81-93.

Trullier, O., & Meyer, J.A. (1998). Animat navigation using a cognitive graph. In Blumberg et al. (Eds.). *From animals to animats 5: Proceedings of the fifth international conference on simulation of adaptive behavior*. Cambridge, MA: MIT Press.

Voicu, H. (2003). Hierarchical cognitive maps. *Neural Networks, 16*, 569-576.

Voicu, H., & Schmajuk, N.A. (2001a). Three-dimensional cognitive mapping with a neural network. *Robotics and Autonomous Systems, 35*, 21-35.

Voicu, H., & Schmajuk, N.A. (2001b). Hierarchical cognitive maps. *Fifth international conference on cognitive and neural systems, ICCNS '01, Boston, USA*.

Voicu H., & Schmajuk, N.A. (2002). Exploration, navigation and cognitive mapping. *Adaptive Behavior, 8(3-4)*, 207-223.

Watson, A. B., & Robson, J.G. (1981). Discrimination at threshold: Labeled detectors in human vision. *Vision Research, 21*, 1115-1122.

Zipser, D. (1985). A computational model of hippocampal place fields. *Behavioral Neuroscience, 99*, 1006-1018.

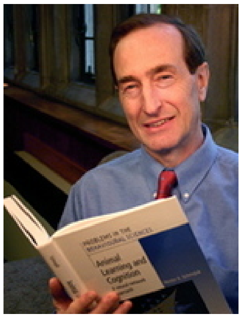

Acknowledgements

The authors would like to thank Jose Larrauri, Matt Serra and Ron Parr for comments and suggestions of previous versions of the manuscript.

Exploration and Navigation Using Hierarchical Cognitive Maps

Nestor Schmajuk[1] and Horatiu Voicu [2]
1 Department of Psychological and Brain Sciences, Duke University
2 Institute for Intelligent Systems, University of Memphis

Chapter Outline & Navigation

Appendix A: Cognitive Map Memory Size

The use of a hierarchical cognitive map to represent an environment with N places, reduces the number of connection values in a nonhierarchical map to a considerable smaller number (see Voicu & Schmajuk, 2001b). Figure 1B illustrates a canvas for an environment subdivided in N (36) places, and Figure 3 illustrates how the environment is represented in a nonhierarchical cognitive map. The cognitive map represents all the connections between a given place in the environment with all other possible places. Because the connection of a place with itself (always equal to 1) is not stored in the map, the number of connections equals $N^2 - N$. That is 1260 connections for an environment with 36 places.

Assuming a WM capacity W = 6, then the number of levels needed to represent this environment is given by L = log N/ log W = log 36 / log 6 = 2. Figure 4 shows how the original environment is represented in a hierarchical cognitive map in which the canvas is subdivided in 6 regions with n (6) places each. Now, instead of representing all the connections between a given place with all other possible places in the environment, we represent the connections (a) between all places in a region plus the connections (b) between places at the borders of that region with those places in neighboring regions that are adjacent to them.

If n (6) is the number of regions, the number of places in a region is N/n (6), and the number of connections between places inside the region is $(N/n)^2$ - N/n. That is, for an environment with N (36) places, subdivided in n (6) regions, the number of connections is 30 in each region. The number of connections at the borders between regions is calculated as follows. The 4 places located at each corner of a region have 5 connections to places located in neighboring regions (see Figure A1), thus the total number of corner connections is 20. Non-corner places located at the borders of a region have only 3 connections, thus the total number of non-corner connections for a region with 6 places is given by the number of non-corner places (6 – 4 = 2) multiplied by 3, i.e., 6. Therefore, the total number of connections at the borders of a region with 6 places surrounded by other regions of the same size is 26. Therefore, the total number of connections for the lower level is given by $[(N/n)^2 - N/n] \times n + 26 \times n = 180 + 156 = 336$.

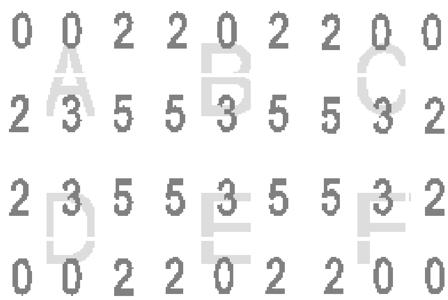

Figure A1. Numbers represent the number of connections from a place in one region to adjacent places in other regions.

Once the number of connections at the lower level have been calculated, we go to the second level and represent the connection between regions. Applying the same rule we applied before, this number is n^2 - n. For n (6) regions, the number of connections for the second level is 30.

Finally, we need to connect the representation of each region in the second level with their respective places in the first level (see Figure 4). For n (6) with N/n (6) that is 36 connections. Therefore, the total number of connections in the hierarchical map is 336 + 30 + 36 = 402, much smaller than the 1260 connections needed to represent the same environment in a nonhierarchical cognitive map. However, as shown in Figure A1, no region in this environment is surrounded by other regions, and therefore, the number of connections (b) at the borders of a region is smaller. This number is 82 for the environment represented in Figure A1. The total number of connections in this case is 328.

Exploration and Navigation Using Hierarchical Cognitive Maps

Nestor Schmajuk[1] and Horatiu Voicu [2]
1 Department of Psychological and Brain Sciences, Duke University
2 Institute for Intelligent Systems, University of Memphis

Chapter Outline & Navigation

Appendix B: Decision Time

In this section we compare the decision times in non-hierarchical and hierarchical representations. For a one-dimensional environment (see Figure B1) that contains 32 places we build a hierarchical representation shown in Figure B2. This hierarchy contains multiple representations of the same environment at different resolutions. On one hand, the time to spread the activity between place 1 and place 32 only by using the representation with the highest resolution is proportional to the number of connections in between place 1 and place 32, namely 31. On the other hand, the time to spread the activity between place 1 and place 32 when all levels of the hierarchy are used is proportional to twice the depth of the hierarchy, namely 10 (= 2 * 5). This is the number of steps required by the spreading activity from place 32 to reach place 1. We can extrapolate that in general, whereas the reaction time for spreading the activity between two places in the case of a nonhierarchical representation of space is linear with the distance between the two places, in the case of a hierarchical representation of space it varies logarithmically with the distance between the two places (see Figure B3). Therefore, by using hierarchical, low-resolution representations of the environment the agent can speed the spatial planning process.

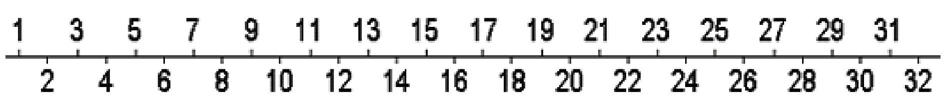

Figure B1. One-dimensional environment that contains 32 places.

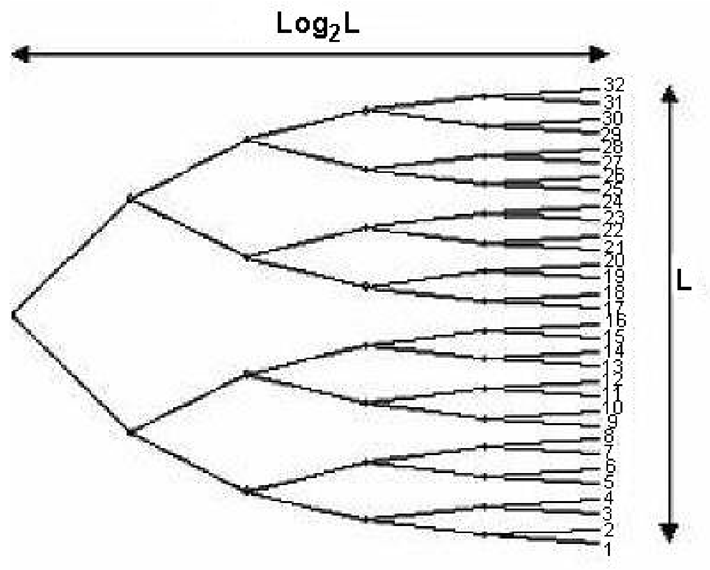

Figure B2. Multiple level representation of a one-dimensional environment that contains 32 places.

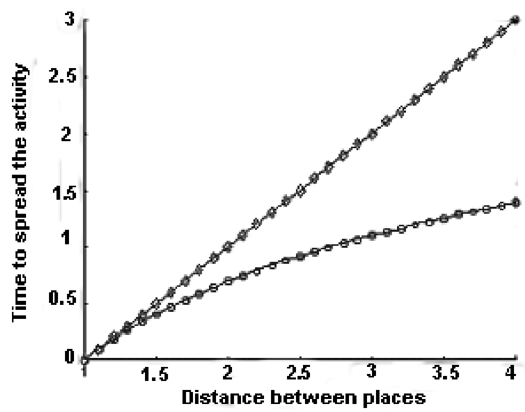

Figure B3. Linear and logarithmic reaction times as a function of the distance between places. Circles represent time to spread the activity in a hierarchical representation of space; diamonds represent time to spread the activity in a non-hierarchical representation of space.

Exploration and Navigation Using Hierarchical Cognitive Maps

Nestor Schmajuk[1] and Horatiu Voicu [2]
1 Department of Psychological and Brain Sciences, Duke University
2 Institute for Intelligent Systems, University of Memphis

Chapter Outline & Navigation

Appendix C: Description of the Associative Networks

In the associative network shown in Figure 3, the connections $V_{i,j}$ and $V_{j,i}$- between the neuron representing place (or region) i (where the agent is located) and each neuron representing neighboring places (or regions) j are updated as follows

$$V_{i,j} = V_{j,i} -= 1 \text{ if place or region j can be accessed from place or region i,}$$

$$V_{i,j} = V_{j,i} -= 0 \text{ otherwise} \qquad [1]$$

At each location, the Motivation System activates the desired goal (e.g., unvisited places or regions), and the goal activates Place (Region) j where it is found,

$$p_i = W_{i,h} \text{ Goal}_h \qquad [2]$$

This activation spreads through the network by using the following equation

$$p_j = \eta \Sigma_i V_{i,j} p_i \qquad [3]$$

When the location (Place or Region) of the agent is activated or the number of iterations is greater than R, the maximum number of iterations the spreading of activation stops. The activation of all neurons represents a gradient that is used for planning a path to the closest goal. For each place m, this is accomplished by choosing

$$\text{Next Place (Region)} = \text{Max}_n\, p_n \qquad [4]$$

where n denotes a neighbor of place (region) m.

Equation 4 ensures that the agent will move in the direction of the nearest goal if all goals are of the same magnitude. When several neighboring places i have an identical Goal value, priorities are used to decide the next place. Priorities are given in the following order: North, West, East, North-West, North-East, South, South-West, and South-East.

Parameter η, ($\eta = .5$) which controls the attenuation at each reinjection in the network, was chosen to obtain an adequate signal at the neurons representing the goal(s). R, the maximum number of reinjections, (R= 625) was selected to ensure that the representation of the location of the agent is activated from any place in the environment where the goal can be found.

Exploration and Navigation Using Hierarchical Cognitive Maps

Nestor Schmajuk[1] and Horatiu Voicu [2]
1 Department of Psychological and Brain Sciences, Duke University
2 Institute for Intelligent Systems, University of Memphis

Chapter Outline & Navigation

Appendix D: Basic Procedures

Under the supervision of the control system, different blocks are activated, thereby processing information in alternative ways in order (a) to decide whether all places in a given region have been examined and (b) to create new regions.

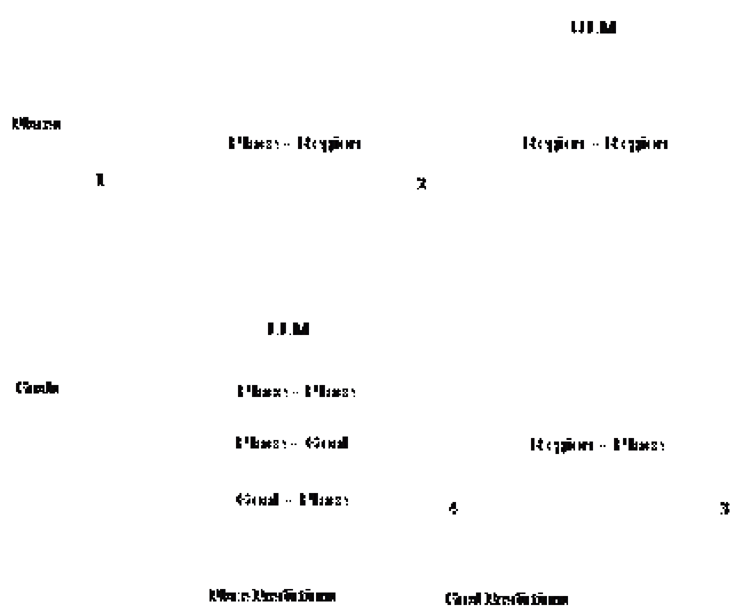

Figure D1. Examined places. Under the supervision of the control system, the current place 1 activates the following representations: current region 2 , all places in current region 3 , the goal of all places in the current region being examined 4 . Thick lines represent flow of information.

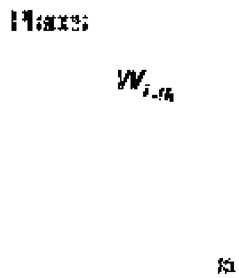

Figure D2. Examined places. The activation of g_1 depends on the active places and the values of the associations between places and the exploratory goal.

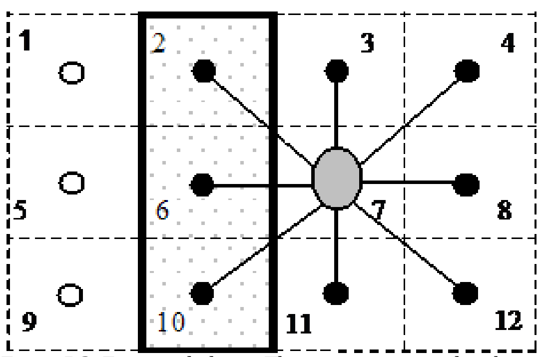

Figure D3. Examined places. The agent represented with a gray oval examines places 2, 3, 4, 6, 7, 8, 10, 11, and 12. Places 2, 6, and 10 are occupied by obstacles. Places 1, 5, and 9 remain unexamined. Black lines radiating from the agent represent sensors, black circles represent examined places, and white circles represent unexamined places. Places 2, 6 and 10 contain an obstacle.

Examined places. While moving through a region, the representation of the current place (see 1 in Figure D1) activates the representation of the region where that place is located through place-region associations (see 2 in Figure D1). In turn, region-place associations activate the representations of all places in that region (see 3 in Figure D1) and, as explained below, through place-exploratory goal associations the agent is able to determine whether or not all places have been examined (see 4 in Figure D1).

Figure D2 shows the unit that stores the connections, $W_{i,g1}$, between places and the goal being examined (see Figure 3). When all the places that belong to the current region are activated, if all $W_{i,g1}$ are zero, then $g_1 = \sum_i p_i W_{i,g1} = 0$. This means that the whole region has been examined. If, on the other hand, there is at least one place for which $W_{i,g1} > 0$, then there still are unexamined places.

Figure D3 shows that places adjacent to examined places and occupied by obstacles are counted as examined (places 2, 6, and 10). Places that are not adjacent to examined places due to the presence of obstacles and therefore out of the reach for the agent's sensors, are counted as unexamined places (places 1, 5, 9).

Creation of a new region. If not all places have been examined ($g_1 > 0$), then some unexplored places in the region are unreachable when the agent is at any place of the already explored places. Therefore, a new subregion should be defined.

There are three possible cases concerning the status of an explored region.

Case 1: All places can be examined, there are no obstacles. Figure D4 shows 6 accessible places in a region that has not been explored. Because the agent can enter all the places in the regions the weights connecting all place representations remain equal to 1. When the agent applies the examined places procedure mentioned above, g_1 equals 0.

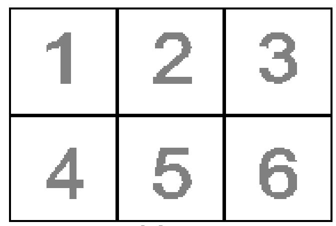

Figure D4. Accessibility: Continuous region. All places are examined and the region remains the same.

Case 2: All places can be examined, some contain obstacles. Figure D5 shows 5 accessible places in an unexplored region. Because place 2 cannot be entered, the connections between the representations of all places and place 2 are set to 0. All other connections remain equal to 1. As before, when the agent applies the examined places procedure, g_1 equals 0. Even if place 2 is occupied by an obstacle, it is considered as examined.

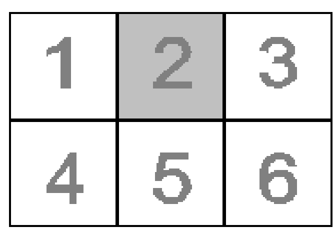

Figure D5. Accessibility: Continuous region with one obstacle. All places are examined. The region remains the same.

Case 3: Not all places can be examined. Figure D6 shows 4 accessible places in a region that has not been explored. When the agent is in place 1 or place 4, because places 2 and 5 are not accessible their connectivity with places 1 and 4 are made equal to 0, whereas the connection between place 1 and place 4 stays equal to 1. As mentioned, places 2 and 5 are occupied by obstacles and are considered as examined. Because they are out of reach for the sensors of the agent, place 3 and 6 are considered unexamined. Therefore, when the agent applies the examined places procedure, g_1 is greater than 0. This result indicates that there are unexamined places that are disconnected from places 1 and 4, and therefore, two new subregions should be defined.

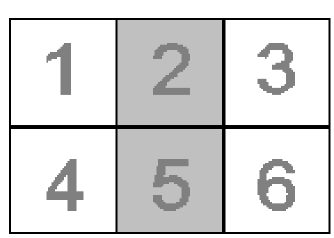

Figure 6. Accessibility: Continuous region. All places are examined and the region remains the same.

In order to update the place-region and region-place connections, the system activates the representation of the examined places (see Figure D7) and the representation of region E, and as a consequence, places E1, E2, E4 and E5 remain part of that region. At the same time because the representations of the unexamined places are not active when the representation of region E is active, unexamined places E3 and E6 are disconnected from region E. When the previous step is completed, because region E has not been completely explored, a new region E' is connected to the unexamined places E3 and E6.

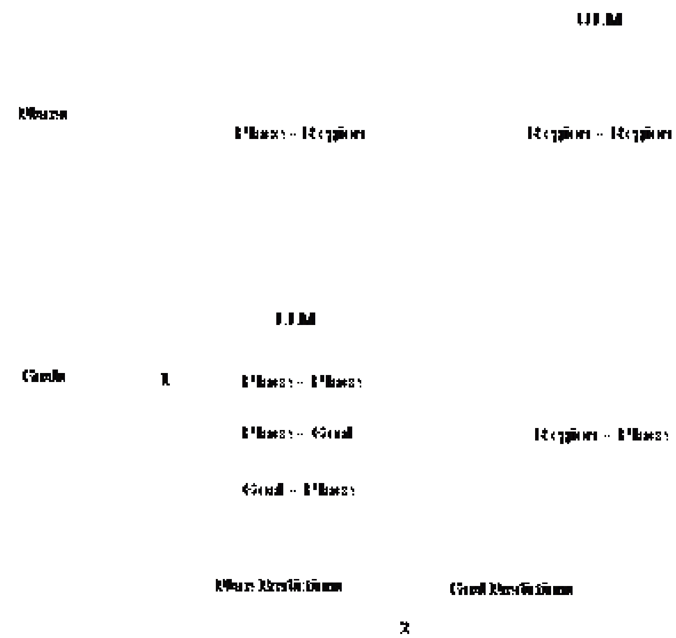

Figure D7. Connected places. By using the LLM the control system activates the following representations: the goal of examined places in the current region 1 and then all examined places in the current region 2. Thick lines represent flow of information.

Exploration and Navigation Using Hierarchical Cognitive Maps

Nestor Schmajuk[1] and Horatiu Voicu [2]
1 Department of Psychological and Brain Sciences, Duke University
2 Institute for Intelligent Systems, University of Memphis

Chapter Outline & Navigation

Appendix E: Updating the ULM

Under the supervision of the control system, different blocks are activated to update the ULM. Let's assume that the region represented in Figure C6 is region E, located in the environment as indicated in Figure E1.

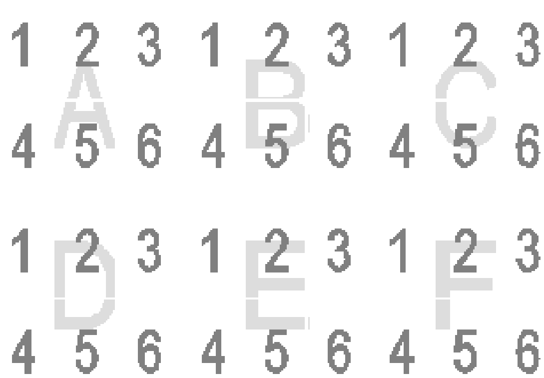

Figure E1. Accessibility: Discontinuous region. Places E1 and E4 are examined. Places E2 and E5 (shown in gray) contain obstacles. After exploration the initial region is divided in two. One that contains places E1, E4, E2, and E5 and another that contains places E3 and E6.

By using place-region connections place E1 (see 1 in Figure E2) activates region E (see 2 in Figure E2) and the activation spreads in the network using the region-place connections (see 3 in Figure E2). Thus, places E1, E2, E4, and E5 become active. Next, the activation spreads from places E1 and E4, using the place-place connections in the LLM, to places D6, D3, A6, B4, and B5 (see 4 in Figure E2). Then, the activation spreads back at the ULM using the place-region connections, and regions D, A, and B become active (see 2 in Figure E2). The connections between the activated regions (D, A and B) and the current region E are reinforced in the ULM. Because region F is not active the connection between region E and region F is weakened. The connections for the new region E' are updated in a similar fashion. Region E' activates places E3 and E6 at the LLM (see 3 Figure E2). Next, places B6, C4, F1, and F4 become active (see 4 in Figure E2). Next, regions B, C and F are activated in the ULM (see 2 in Figure E2). The activated regions and the new region E' are connected.

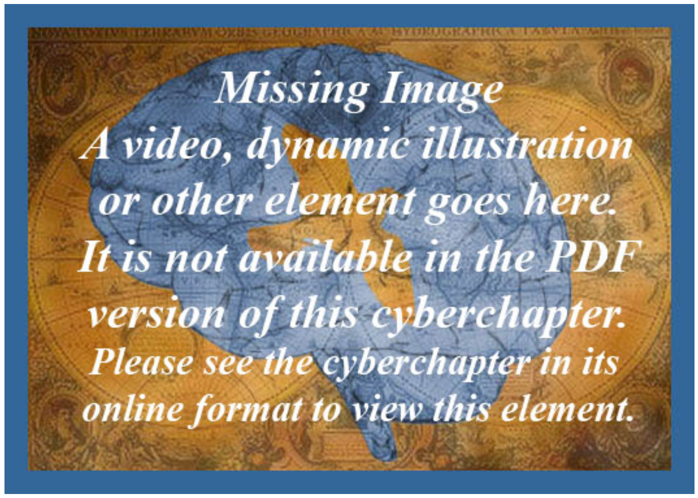

Figure E2. Update of region-region connections. Under the supervision of the control system, the current place 1 activates the following representations: current region 2, all places in current region 3, places in neighboring regions 4 and through the place-region connections the neighboring regions 2. Thick lines represent flow of information.

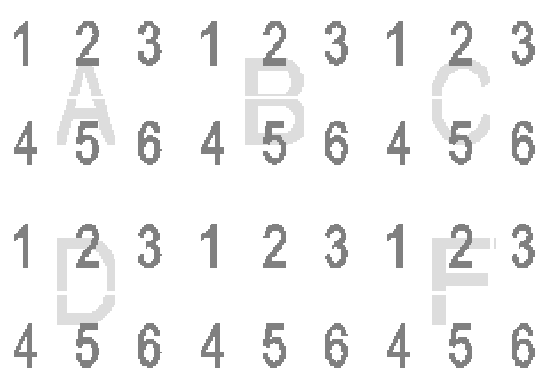

Figure E3. Region division. The obstacle that occupies places 2 and 5 divides the region in two subregions E and E'.

In sum, as a result of the ULM update, region E is connected to regions D, A and B, and region E' is connected to regions B, C, and F. Figure E3 shows how region E is split in two subregions.