Concentric Spatial Maps for Neural Network Based Navigation

Gerald Chao and Michael G. Dyer

Computer Science Department,
University of California, Los Angeles
Los Angeles, California 90095, U.S.A.
gerald@cs.ucla.edu, dyer@cs.ucla.edu

Abstract

A model for navigation based on artificial neural networks is proposed and tested for the task of planning and reaching a goal location in a continuous, dynamic environment. The model is based on the concept of concentric spatial maps, which is a cognitive representation of the surrounding environment in the central nervous system. With the construction and maintenance of these cognitive maps by a specialized neural network architecture, initial results demonstrate the system's ability to reach the goal by continuously fusing sensory inputs from the environment with the cognitive maps in memory to dynamically plan a path to the goal.

Introduction

Most mobile organisms are able to navigate within their ever-changing environment. Navigation, as described by Gallistel, is the process of determining and maintaining a course or trajectory from one place to another (Gallistel, 1990). Whether it is simple chemotaxis or high-level planning of paths, animals must first have the ability to determine their current location, either by dead-reckoning or estimating the distance to known landmarks. Once the current location is determined, the animal receives sensory inputs about spatial relationships of the objects in the current environment. This information is then presumably stored on some kind of map. Finally, based on the information gathered, animals can make movements toward desired targets. As stated by Levitt and Lawton (1990), navigation can be defined by the following three questions: 1) "Where am I?"; 2) "Where are other places relative to me?"; and 3) "How do I get to other places from here?".

Based on this formulation of navigation, maps are needed to store and maintain spatial relationships between the animal and its surroundings. Coined by Tolman (1948) as "cognitive maps", these internal maps are described by Gallistel (1990) as:

> "A record in the central nervous system of macroscopic geometric relations among surfaces in the environment used to plan movements through the environment."

This hypothesis is supported by biological experiments such as Morris' water maze experiments (Morris, 1981), where rats are able to reach a goal location without direct cue stimuli. In these experiments, animals demonstrated the ability to reach a goal without directly sensing it. Therefore, cognitive maps are hypothesized to maintain previously detected objects and their locations instead of relying solely on immediate sensory inputs.

Although there is currently little physical evidence of biological cognitive maps, many models with different types of maps have been proposed. Since there has not yet been any consensus on the representation and mechanism of cognitive maps, there is currently a wide range of systems that perform specific tasks. For a comprehensive review, see Trullier et al. (1997). Unfortunately, the global mechanism which fuses multi-sensory inputs with prior knowledge while maintaining, updating, and utilizing this information for navigation is currently not well understood (Truiller et al., 1997).

In this article, a new navigational model, based on concentric spatial maps (CSMs), is proposed. It attempts to address some of the attributes and behaviors missing from current models. The CSM will first be presented, then the navigational model, and finally, some preliminary results from computer simulations.

Concentric Spatial Map (CSM)

A CSM is a cognitive map that builds and maintains an egocentric, topological map of the obstructions in the surrounding environment by fusing two sensory inputs: distances to surfaces and dead-reckoning. The sensor that interfaces the external world can be of any type that can sample the environment and extract the distance/depth of surfaces relative to the animat. For example, it can be optical (e.g., stereoscopic eyes), tactile (e.g., antennae), or sound (e.g., sonar-like). In addition, the sensor could detect additional information, such as color, intensity, tactile, etc. For simplicity, visual inputs of a two-dimensional environment will be used in describing this model. The other input the CSM receives is the movements made by the animat, namely, rotations and translations. This is used to update the map with respect to the movements by maintaining the

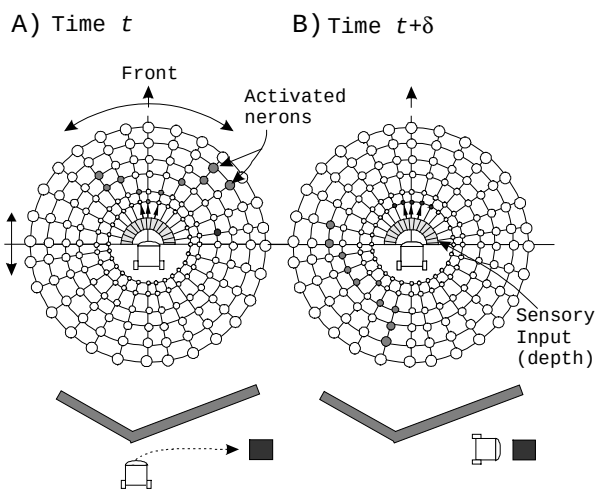

Figure 1. A schematic of how obstacles (two walls) are represented on the CSM. The activated (shaded) neurons on the map represent the detected walls, which are in front of the animat in part A. As the animat turns, each neuron on the map also rotates the map opposite to the turn by transferring the activations to neighboring neurons. As the animat moves forward, the map shifts downward, and the opposite occurs when moving backwards. When the animat reaches the square in part B, most of the walls are behind the animat and thus undetectable by the front sensors. However, the CSM still maintains the egocentric distances and angles to the walls.

egocentric distance and angle to previously detected obstructions.

The CSM is a network of neurons arranged in concentric circles with activations on the network representing detected obstructions (figure 1). The center of the CSM represents the animat itself, and the rings of neurons maintain a topological map of the surrounding environment. The inner-most ring contains obstructions closest to the animat, whereas the outer-most ring stores the most distant obstructions within the map's range. For example, a detected obstruction at three feet might activate neurons on the third ring, whereas obstructions 20 feet away might activate neurons on the tenth ring. The angle of detected obstructions with respect to the animat is provided by the sensors and is preserved by the CSM. By sending sensory inputs radially outwards to the columns of neurons, the angle of an obstruction with respect to the animat is saved by the neurons that are at the same angle with respect to the centerline of the CSM. For example, an obstruction 45° to the right at five feet will activate the neuron at the intersection between the radial column at 45° clockwise from the centerline and the concentric ring that is receptive to the distance of five feet.

Representation

With the information provided by the sensors, a firing neuron on the CSM represents an obstruction that has been detected at a particular distance and angle from the animat. With an animat that has a 180° visual sensor, one set of sensory inputs will activate the corresponding neurons on the CSM and form a topological map to the front of the animat (see figure 1). Because of the sampling done by the sensors, the CSM is an approximation of the surroundings proportional to the resolution of the sensors. Also, since the number of neurons is finite, the CSM will lose an increasing amount of information as the distance to an obstruction increases. That is, the farther the obstructions are, the higher probability that multiple obstructions will activate the same neuron on the CSM, resulting in the map being unable to distinguish between distant obstructions. Therefore, depending on the quality of the sensors and the amount of neural resources allocated, the CSM learns and maintains an egocentric topological map of decreasing accuracy from the inner towards the outer rings of the map.

Mechanics

Each neuron on the CSM performs two functions: first to receive new information that falls within its receptive range and second, to update the information on the map by firing periodically. As described above, neurons on the CSM are activated as a result of an intersection of the corresponding ring and the angle of detected obstructions. This is achieved by the depth sensory information gated by an interneuron at each ring of the map, while the angle information is preserved from the sensor inputs by the radially outward projections (figure 2A). Each interneuron filters out irrelevant distances and each CSM neuron does not receive irrelevant angle information. As a result, the topological map of the environment is easily built from the sensory inputs by the method of intersection.

The persistence or memory of the map is achieved by periodic firing and continual sensory input, such as from saccade or antennae movements. When the animat remains stationary, each neuron refreshes its own contents by a self-recurrent connection since each neuron is receptive to its own range of distances and angles (figure 2A). This ensures that any previously

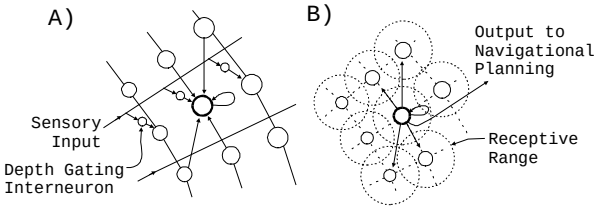

Figure 2. Part A shows all the efferent connections a neuron on a CSM receives from its neighbors, and part B illustrates the afferent connections this same neuron makes to the neighboring neurons. Note that the neurons are not fully connected to each other because horizontal connections are unnecessary since lateral movements are omitted in this model. The dashed circles are a representation of the receptive range of angle and distance for that neuron.

seen obstruction that is no longer detectable by the sensor, such as in the case of occlusion, will remain on the map. At the same time, a portion of the map is receiving sensory input and is refreshed with any changes in the environment. To ensure the map contains the most updated information, this portion of the older map (old activations that decays at the recurrent connections) is replaced with the new sensory input. As a result, the map will eventually lose, or forget, outdated information. Depending on the decay rate, a stationary animat will eventually forget obstructions behind the animat or surfaces obstructed by other surfaces, resulting in the map retaining only the geometric relationships to obstructions immediately detectable from the current location. This allows for the animat to be blindly relocated to a new environment yet still be able to quickly construct a new map and replace the old.

Movements

As the animat moves, the activations on the map have to be updated with respect to the movement. The movement information is provided by the inputs from dead-reckoning in the form of distance and angle of the movement. As the animat turns or rotates, the concentric rings allow for very easy updating since all activations on the map are simply shifted opposite to the rotation. This is achieved by each activated neuron on the map increasing or decreasing the stored angle by the rotation angle. This value is then propagated to its neighbors (and to itself through the recurrent connection) (figure 2B). Whichever neuron that is receptive to the range of angles the new value falls into will become activated. Therefore, if the angle of the turn is slight, the same neuron might be activated whereas a larger turn might activate a neighbor. However, this does not imply that each neuron needs to be fully connected to all of the other neurons on the same ring. The number of neighboring neurons to synapse with is determined by the sensor's limit of maximum detectable angular velocity, i.e., there is a maximum amount of change in angles the sensor can detect per sample, limited by its sampling rate. Therefore, as long as the CSM is updating at the same rate as the sensor is sampling the environment, the neurons do not need to be fully connected but only to a small number of nearest neighbors.

The same principle applies to translations, but some more calculations are needed to update the map. For simplicity, this discussion will be limited to only forward and backward movements. As the animat moves, the displacements will be provided by the dead-reckoning inputs. All of the active neurons compute the new distance d' and the new angle θ' by the following equations:

$$d' = \sqrt{(x^2 + y'^2)} \approx |x| + |y'|$$
$$\theta' = \tan^{-1}(y' / x)$$
$$y' = \sin(\theta)*d+\Delta y$$
$$x = \cos(\theta)*d$$

where Δy is the amount of the forward or backward movement. Although neural networks can be trained to perform these functions, this is assumed and not tested by the current implementation of the model. The updated information is then broadcasted to the neighbors and itself in the same fashion as the rotational movements. Whichever neuron that is receptive to the range of the distance and the angle will be activated. This is how the CSM maintains the geometric relationships between obstructions and the animat as it moves around in the environment.

By having all of the activated neurons updating the map continuously (with or without direct sensory input), the animat can lose sight of obstructions but still have a sense of where they are with respect to the animat. For example, if a wall is first detected and the animat turns around 180° and can no longer detect it, the CSM still maintains that there is a wall behind the animat (figure 1B). Notice that the environment's geometry is learned without weight adjustments; information is temporarily stored and manipulated for the duration when obstructions are within the distance range of the CSM. Due to the finite range and the decay of the map, obstructions that are far away or not regularly refreshed will eventually be forgotten. The map is not designed to be a permanent record but a temporary representation of the immediate environment for the purpose of navigation.

Goals

Information about the location of obstructions is useful for obstacle avoidance and path planning. Path planning requires knowing the associated navigational goals to be achieved, such as food goals for hunger and avoidance goals for survival. To store these goals and to maintain the topology of where they are with respect to the animat, a duplicate of the CSM is used. Instead of receiving depth information, the goal CSM (G-CSM) receives sensory inputs from "edible" sensors and "harmful" sensors. This map is then maintained in synchrony with the obstructions CSM (O-CSM) as described earlier, updating angles and distances as the animat moves. Therefore, the animat can have a general sense of the direction and distance to goals, even if the animat is not receiving direct stimuli from them.

Navigation

Once a goal has been determined, a path is planned from the current location to the goal. This is achieved by a navigational CSM (N-CSM) incorporating the O-CSM and the G-CSM to compute a path to the goal.

The neurons on the N-CSM have the same structure, i.e., in concentric rings, but the task they perform is different. Initially on the N-CSM, the neuron that is at the same location as the goal is activated, indicating the destination of the path (figure 3). This neuron then initiates a spreading activation to the neighbors and throughout the N-CSM. At the same time, this map receives inhibitory inputs from the O-CSM. That is, since each active neuron on the O-CSM represents an obstacle, these neurons send activations through inhibitory interneurons to disable the corresponding neurons on the N-CSM. Due to these deactivated neurons, the N-CSM will not spread the activation to neurons representing obstructions and thus does not compute a path through obstructions, but around them.

As the activation spreads across the map and around obstacles, it should eventually reach the inner-most ring of the N-CSM. If any of the neurons on this ring is activated, then there is a path from the animat's current location to the goal. Although the path from this neuron to the goal can be calculated by backtracking to the goal location, the animat only needs to know which initial direction to take to reach the goal. Thus, the angle of the neuron on the inner-most ring that receives the activation first is the next heading for the animat to turn towards. As long as the animat rotates toward this direction and moves forward, the N-CSM will be able to re-compute the remaining path and the direction to the goal.

The obstructions and goal CSMs are an integral part of this navigation system. Having a topological map of the environment makes computing the path around obstacles possible. Having the goal map maintaining the heading and distance to the goal provides the navigational system with the destination. Therefore, even without external stimuli, the N-CSM can still plan the path with the information from the other CSMs. Moreover, this path is constantly re-computed as the animat moves towards the goal. So in a dynamic environment, if new obstructions or previously undetected obstructions appear, the N-CSM can plan a new path to the goal immediately.

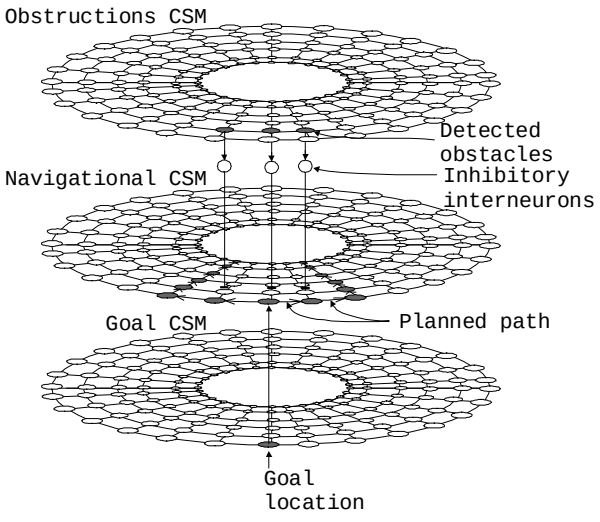

Figure 3. A schematic of the connections the navigational CSM (N-CSM) receives from both the obstructions CSM and the goal CSM. Only a small portion of the connections is shown for clarity. The shaded neurons on the N-CSM illustrate the spreading activation initiated by the goal neuron to the inner-most ring, which then determines the animat's next heading to reach the goal.

Experimental Results

A continuous, non-toroidal, confined (by walls) world is used in simulating an animat sensing and navigating within its environment. Static rectangles of arbitrary sizes are placed in the world as obstructions. The sensor samples the front 180° of the animat at 10° resolution and detects the distances from the animat to obstacles. The obstructions, goal, and navigational CSMs are each composed of 12 rings with 37 neurons per ring for a total of 444 neurons per map. The neurons on the maps have non-overlapping distance and angle receptive ranges, thus only one neuron can be activated by each sensory input. The depth-gating internerons use the following equation to determine its receptive range:

$$\text{Ring number} = \ln(\text{distance}+1)*3$$

For instance, if an obstacle is at distance 2, then ring number 3 is the receptive one. The log term allocates more neural resources to the closer obstructions and improves the distance range of the map, while the factor of three is used to improve the map's distance resolution. These constants reflect the amount of neural resource allocated in this particular implementation and are not critical to the generalized model.

The animat can rotate in both directions by an arbitrary angle, as well as move forward and backwards by distances within an upper limit. A single goal is activated manually on the goal map, and the animat is placed in the environment without any other *a priori* knowledge (figure 4A). For demonstrative purposes, a red square is placed at the goal location but does not participate in the navigational process. Also, the animat is not allowed to explore the environment. Rather, it is placed in the world with the only purpose of reaching the goal. The animat is then monitored for its ability to reach the goal while avoiding obstacles.

In the simplest scenarios, if no obstruction is placed between the animat and the goal, the animat turns toward the goal until it is facing the goal and moves in a straight path to the goal. If a horizontal wall is placed between the animat and the goal, the animat turns and moves toward a detected opening and proceeds to the

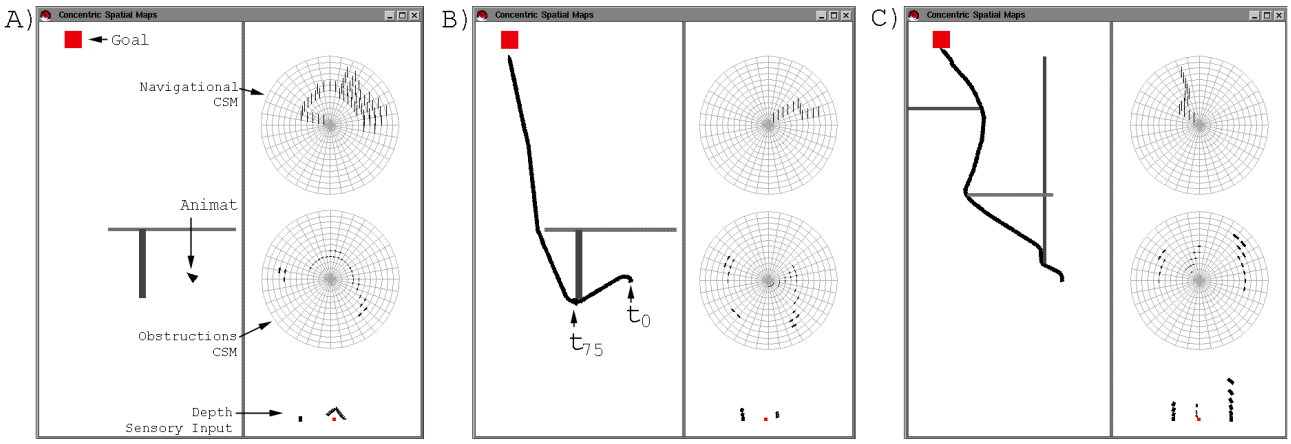

Figure 4. Screen shots of the simulation environment. In part A, the animat has turned 45° to the left. The sensory input shows the corner in front of the animat, and the obstructions CSM shows what the animat has learned from the environment thus far. The vertical bars on the navigational CSM indicate the spreading activation from the goal node to the inner-most ring. Part B is the trace of the path the animat took to the goal, as well as the state of the maps at time t_{75}. The N-CSM at t_{75} indicates that the animat should turn 50° to the right, go straight to the left corner of the horizontal wall, then turn right some more and head for the goal. Part C is the trace of another test scenario, with the state of the maps at time t_0 shown on the right.

goal once it clears the wall's corner (this behavior not shown here).

With a vertical wall to the left of the animat in addition to the horizontal wall (see figure 4A,B), the N-CSM calculates the shorter path to be around the bottom of the vertical wall at t_0. This is because the left half of the horizontal wall is blocked by the vertical wall, thus making it appear that the left side of the environment is clear of obstructions. Once the animat detects the left part of the horizontal wall at t_{75}, it still proceeds to go around the left corner of the horizontal wall since this path is shorter than turning around and going through the opening on the right side.

However, it is possible that there is no opening on the left side and the animat must turn around. This situation is tested and shown in more detail in figure 5. Initially, the animat again turns left because of the vertical wall blocking the view (figure 5A). As soon as the first corner is cleared, the animat begins to detect the horizontal wall. The animat continues forward because there still appears to be gaps in the wall (figure 5B). Once the animat determines that the path is completely blocked by the horizontal wall, it then turns left in hopes of moving around the left wall that surrounds the world (figure 5C). As the animat turns, the enclosed environment is finally completely mapped onto the O-CSM and the animat is forced to turn around (figure 5D) and then proceeds to clear the vertical wall and the right corner of the horizontal wall. Once the animat clears the right corner, it heads almost directly towards the goal (figure 5E). Figure 5F is the complete trace of the path the animat took to the goal.

Conclusion

The model presented here is able to guide the animat to the goal location, even using maps with a very coarse resolution. The model not only guides the animat to a goal while avoiding obstacles but also dynamically evaluates the path to the goal along the way. This specialized neural network architecture is able to construct and update the cognitive maps that represent geometric relations to surrounding surfaces by fusing internal information with sensory inputs. These maps are then integrated into a navigational system that is able to plan a path to the goal. The adaptation to the changing environment is achieved by the maps continually updating and re-evaluating the current surroundings, which allows for a new path to be computed as the world changes. The resulting system is one that performs similar navigational tasks seen in biological systems.

This model is simple yet powerful, able to scale easily because the complexity of the system is both linear in time and space except for simulating the spreading activation on the N-CSM, which can be done in polynomial time. The computations performed by this model are proportional to the number of neurons on the maps, not the complexity of the environment. Therefore, by simply allocating more neurons to the maps and increasing the sensor resolution, this model is able to navigate within more complex environments by forming more detailed maps. Also, this model is amicable to parallel implementations, since most computations can be performed simultaneously by the individual neurons.

Future Work

Due to the limited range of the maps, it is possible that a goal would be located beyond the range of the goal map and thus the N-CSM would be unable to calculate a path to the goal. A possible solution is a special outermost ring on the goal map that stores any

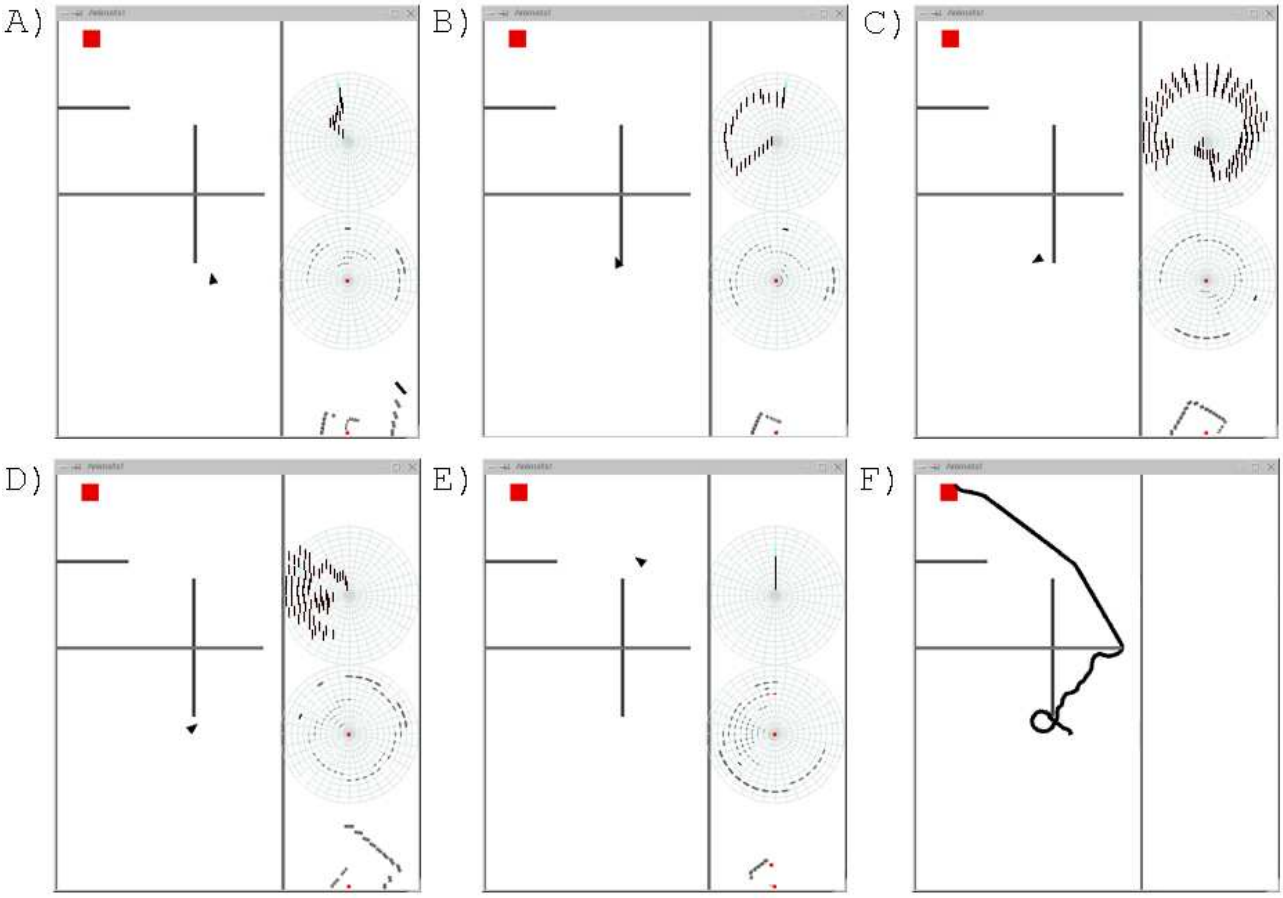

Figure 5. A sequence of movements demonstrating the model's ability to incorporate newly sensed obstructions and plan a new path to the goal. Please refer to text for more details.

out-of-range goals. A better way might be to extend the range of the current system with a hierarchy of CSMs that cover increasingly larger areas (figure 6). For example, a system can contain a short-range CSM of 10 feet for maintaining the immediate surroundings, a medium-range CSM for objects within 100 feet, and a long-range CSM for 1000 feet, etc. All of these maps are identical except for the depth-gating interneurons, with each map updating and maintaining egocentric distances and angles to surrounding surfaces that fall within their distance range. Therefore, the range of the system is extended by simply allocating more maps.

Another possible adjustment to the above model is to alter the decay rate at the recurrent connections so that long-range maps retain activations longer than the short-range maps. With time, the long-range maps will retain large and permanent objects over small and changing objects. These maps are in essence topological landmark maps, providing spatial information on large and permanent objects that can serve as landmarks. For example, the distances and angles between a goal and the landmarks can be calculated easily and be stored by an associative memory. When an animat comes upon a landmark in the future, it can retrieve any possible matches and recall where the goal might be.

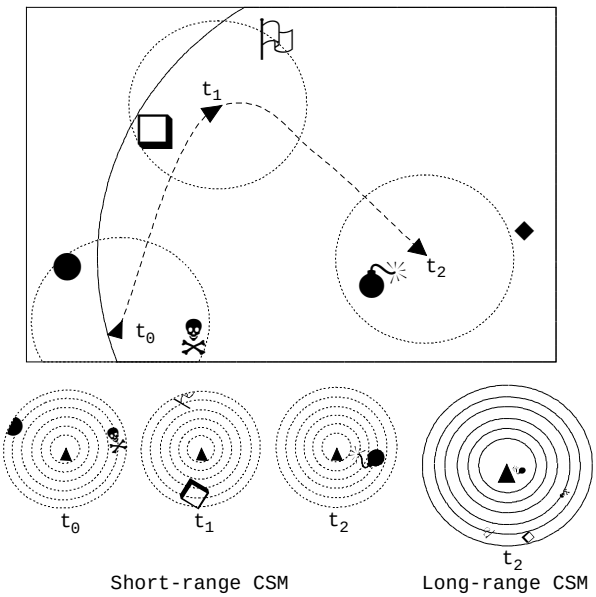

Figure 6. An illustration of a two-level hierarchy of CSMs. The short-range CSM maintains local information as it moves around within the environment and the long-range CSM retains most of the objects encountered up to time t_2.

With the goal activated on the goal CSM, the navigational CSM will then guide the animat to the possible goal location.

Lastly, the CSMs presented here can be extended to represent 3-D spatial relationships by the same principle. Instead of rings, neurons can form spheres and be activated by a 2-D matrix of distance sensory inputs. So instead of just planning a path around obstacles, the animat could also entertain paths over or under obstacles, as well as through holes and tunnels.

The model presented in this article is a general, simple, robust, and extensible navigational system. Whether this model resembles biological systems remains to be seen, but the demonstrated capabilities of the model should be useful for applications in generalized navigation within a continuous and dynamic environment.

Acknowledgment

The second author is supported in part by an Intel University Research Program grant.

References

Gallistel, C. R. (1990) *The Organization of Learning*. Cambridge, Massachusetts: The MIT Press.

Levitt, T. S., Lawton, D. T. (1990) Qualitative Navigation for Mobile Robots. *Artificial Intelligence*, 44, pp. 305-360.

Morris, R. G. M. (1981) Spatial Localization Does Not Require the Presence of Local Cues. *Learning and Motivation*, 12, pp. 239-260.

Tolman, E. C. (1948) Cognitive maps in rats and men. *Psychological Review*, 55, pp. 189-208.

Trullier, O., Wiener, S. I., Berthoz, A., Meyer, J. (1997) Biologically Based Artificial Navigation Systems: Review and Prospects. *Progress in Neurobiology*, 5, pp. 483-544.