THE SKELETON IN THE COGNITIVE MAP
A Computational and Empirical Exploration

BENJAMIN KUIPERS *is a professor in the Department of Computer Science at the University of Texas, Austin. Dr. Kuipers received his Ph.D. in mathematics from the Massachusetts Institute of Technology in 1977. His research focuses on the representation of common-sense and expert knowledge, with particular emphasis on the effective use of incomplete knowledge.*

DAN G. TECUCI *is currently a Ph.D. student in computer sciences at the University of Texas, Austin. He received a master of science in computer sciences from University of Texas, Austin in 2001. His main research interest is artificial intelligence, with focus on knowledge representation, development of knowledge bases through composition, cognitive modeling, and machine learning.*

BRIAN J. STANKIEWICZ *is an assistant professor in the Department of Psychology at the University of Texas, Austin. Dr. Stankiewicz received his Ph.D. in cognitive psychology from the University of California, Los Angeles in 1997. His research focuses on computational models of spatial navigation and object recognition along with testing the predictions of these models.*

ABSTRACT: Experts seem to find routes in complex environments by finding a connection from the source to a "skeleton" of major paths, then moving within the skeleton to the neighborhood of the destination, making a final connection to the destination. The authors present a computational hypothesis that describes the skeleton as emerging from the interaction of three factors: (a) The topological map is represented as a bipartite graph of places and paths, where a path is a one-dimensional ordered set of places; (b) a traveler incrementally accumulates topological relationships, including the relation of a place to a path serving as a dividing boundary separating two regions; and (c) the wayfinding algorithm prefers paths rich in boundary relations so they are likely to acquire more boundary relations. This positive-feedback loop leads to an oligarchy of paths rich in boundary relations. The authors present

ENVIRONMENT AND BEHAVIOR, Vol. 35 No. 1, January 2003 81-106
DOI: 10.1177/0013916502238866

preliminary computational and empirical tests for this hypothesis, and provide initial results.

A THEORY TO EXPLAIN THE SKELETON

Expert wayfinders in a complex large-scale environment use a “skeleton” of important paths and places to guide their problem solving (Chase, 1982; Lynch, 1960; Pailhous as cited in Chase, 1982). How is this skeleton represented? How is it acquired? And how does it help in wayfinding?

In this article, we describe a preliminary computational hypothesis to explain the phenomenon of the skeleton based on the concepts in the Spatial Semantic Hierarchy (Kuipers, 2000) (which extends the TOUR model; Kuipers, 1978, 1982). We also describe a set of computational and empirical tests that can be applied to this hypothesis. Our preliminary results suggest directions for further investigation.

THE SKELETON

Researchers who have studied expert wayfinders such as experienced taxi drivers (Chase, 1982; Golledge, 1999; Pailhous as cited in Chase, 1982; Timpf, Volta, Pollock, & Egenhofer, 1992) have observed a common strategy. Such an expert knows a large number of places and paths, but much of their travel occurs within a small subset of “major” paths, which is sometimes called the skeleton (Figure 1). When given a wayfinding problem, the expert first finds a route from the initial point to the nearest point on the skeleton, then finds a route within the skeleton to a point near the destination, and finally finds a route from that point to the destination itself.

This sketch raises several questions. How are the paths and places in the skeleton selected from the larger set the expert knows about? Is there a qualitative difference between the skeleton and the rest of the map, or is the role of the skeleton an emergent behavior of some uniform mechanism applied to the entire cognitive map?

The hypothesis presented here is that the skeleton is an emergent phenomenon arising from the interaction between:

1. the topological representation for places and paths;
2. the incremental, opportunistic learning of “boundary relations” during travel; and
3. the use of boundary relations to provide subgoals during wayfinding.

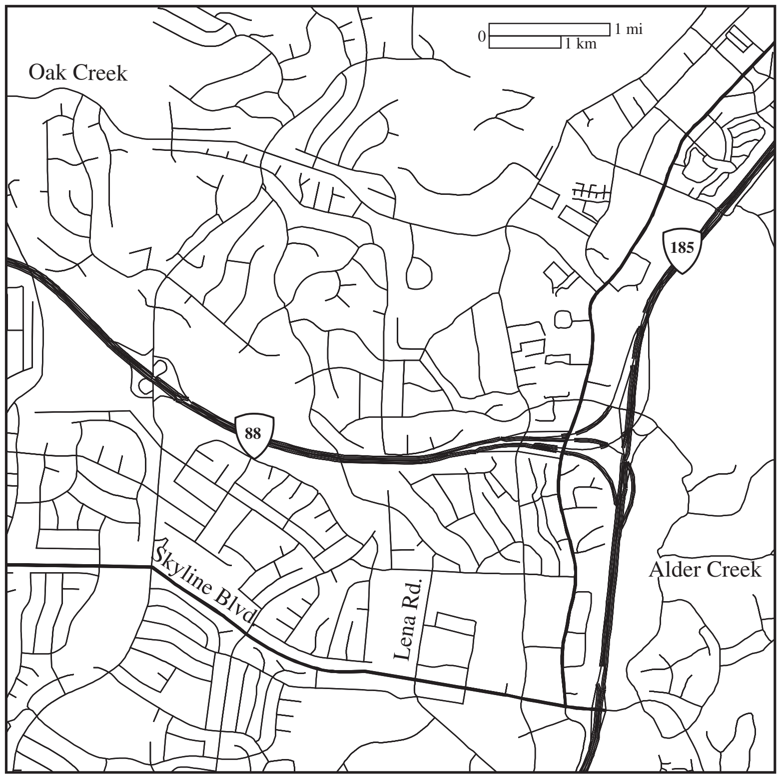

Figure 1: A Large-Scale Cognitive Map Has a Skeleton of Major Paths; the Graphical Conventions for Emphasizing Major Streets and Highways on a Printed Map Are Related but Not Identical to the Skeleton in the Cognitive Map

THE TOPOLOGICAL MAP

It is widely accepted (Lynch, 1960; Siegel & White, 1975) that the "cognitive map" includes a topological level of description, in which places (0-D), paths (1-D), and regions (2-D) are symbolically described and linked by relations such as connectivity, order, and containment. Metrical relationships such as distance and direction may also be associated with the topological map, but there is typically no single global frame of reference, and metrical errors in a variety of tasks are much more common than topological errors.

The Spatial Semantic Hierarchy (SSH) (Kuipers, 2000) is a computational model of knowledge of large-scale space, consisting of a lattice of different but related representations for space. The control level consists of knowledge

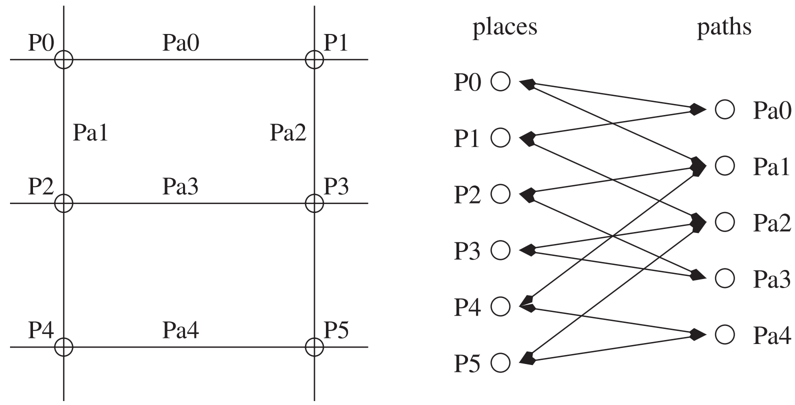

Figure 2: The Rectangular Block Environment (Left) Is Described in the Topological Map (Right) as a Bipartite Graph of Places and Paths

of control laws for taking the agent from one "distinctive state" within the environment to another. A state (position plus orientation) is distinctive if it is the stable point of a local "hill-climbing" control law that eliminates moderate amounts of accumulated error by bringing the agent to a particular state from anywhere in its local neighborhood. The causal level of the SSH abstracts the control laws to actions and represents behavior in the environment as a set of discrete causal schemas (*S*, *A*, *S'*), describing the relation between a state, an action, and the resulting state. The topological level posits places, paths, and regions to account for the experienced regularities in the causal description. Local pieces of metrical information can be used throughout the other levels, but a global metrical model with a single frame of reference can only be created after all the other descriptions exist. The TOUR model (Kuipers, 1978, 1982) has been incorporated into the causal and topological levels of the SSH.

In the SSH topological map, a path describes an extended one-dimensional structure such as a street. The topological map is a bipartite graph (Figure 2), with nodes corresponding to places and paths, and arcs corresponding to the assertion that a particular place is on a particular path. Other relations included in the topological map but not represented explicitly by the bipartite graph include (a) the linear ordering of places along a path, (b) the circular ordering of paths intersecting at a place, and (c) the boundary relations discussed in the following. The benefit of the bipartite graph of places and paths is that physically distant places on the same path may be close in the topological map, making wayfinding easier.

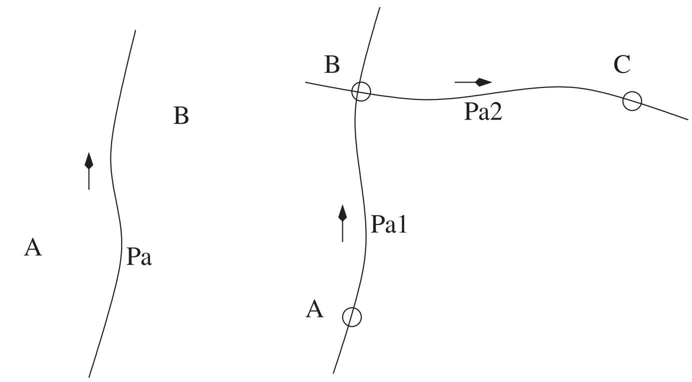

Figure 3: Boundary Relations. A path serves as a boundary separating places on the left from those on the right. Boundary relations can be inferred from local travel patterns.

BOUNDARY RELATIONS

The SSH topological map can represent more than connectivity and order. A path is a one-dimensional subset of the environment, with a direction implied by the order on its places. A directed path is described by (*Pa*, *dir*), where *dir* is either *pos* or *neg*, and –*dir* is the other one. If a directed path extends to infinity, it divides the places in the environment into three subsets: those on the path, those on the right, and those on the left. Note that *right* and *left* are used here as topological terms. If the path curves, a place that is topologically to the right may occasionally be visible to the traveler's egocentric left.[1]

The assertion that a place *B* lies to the right of a directed path (*Pa*, *dir*) is called a *boundary relation*: *right_of*(*Pa*, *dir*, *B*). (See Figure 3, left.) We define *left_of* similarly and provide the axiom

$$right\_of(Pa, dir, P) \equiv left\_of(Pa, -dir, P).$$

Boundary relations can be acquired incrementally during travel by simple local rules. For example (Figure 3, right):

If the traveler moves along a path (*Pa*1, *dir*1) from place *A* to place *B*
and takes a right turn at *B* onto (*Pa*2, *dir*2) and
travels along (*Pa*2, *dir*2) to reach place *C*,
then we can conclude *right_of*(*Pa*1, *dir*1, *C*).

We call this the *(AB)C rule*. Under the same conditions,

> We can also conclude that *right_of*(*Pa*2, *dir*2, *A*).

This is called the *(BC)A rule*.

We are making some relatively weak and plausible assumptions. We assume that the finite length of the boundary *Pa*1 does not lead us astray. We assume that path *Pa*2 does not intersect or cross *Pa*1 except at *B*. These inferences are implemented as *default rules*, so that if there is contrary evidence, no conclusion is drawn (Remolina & Kuipers, 1998). The aforementioned rule applies only when there is a direct connection from *Pa*1 to *C*, but it is straightforward to handle more complex connections. Meanwhile, we are not assuming that *Pa*1 or *Pa*2 are straight. We are not assuming that a right turn is a 90-degree turn. We are not assuming that *C* is close to *Pa*1, as the path *Pa*2 can be quite extended.

Using local rules such as the ones mentioned earlier, any experience traveling through the environment will lead the topological map to accumulate boundary relations among places and paths experienced during travel. This is the first link in a positive-feedback system to ensure that paths that are used frequently tend to be used more frequently.

WAYFINDING USING THE BOUNDARY HEURISTIC

Wayfinding is the process of finding a route from an initial place to a destination place. At the SSH topological level, a route is an alternating sequence of places and paths, each connected to its neighbors. Once a topological route is found, it can be refined for execution, first to an alternating sequence of states and actions at the causal level and finally as a sequence of control laws at the control level.

There are a number of graph search algorithms that can find paths in a topological map (Elliot & Lesk, 1982). Metrical information such as estimates of path segment lengths can be used to guide heuristic search in the A* and Dijkstra algorithms. However, the boundary relation can be used as the basis for a purely qualitative heuristic to guide wayfinding search (Figure 3).

> If we are searching for a route from place *A* to place *B*, and
> if there is a path *Pa* such that *left_of*(*Pa*, *dir*, *A*) and *right_of*(*Pa*, *dir*, *B*),
> then consider *Pa* a subgoal and search for routes from *A* to *Pa* and from *Pa* to *B*.

When places *A* and *B* are on opposite sides of path *Pa*, the route connecting them must necessarily cross the boundary. The heuristic can also be useful in

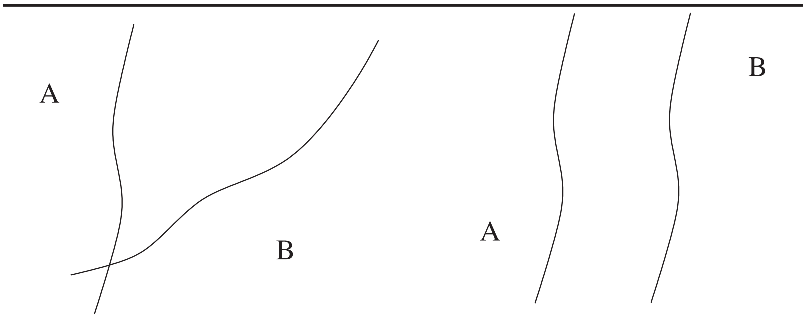

Figure 4: Wayfinding Using Boundary Relations

case both places are on the same side of the boundary, although of course the route could be inefficient. Unfortunately, we cannot in general expect the same boundary to be related to both endpoints of the desired route. A more general form of the boundary heuristic is (Figure 4):

> If we are searching for a route from place *A* to place *B*, and
> if there are paths *Pa*1 and *Pa*2 such that *A* shares a boundary relation with *Pa*1 and *B* with *Pa*2,
> then propose the subgoal of finding a connection from *Pa*1 to *Pa*2.

In the simplest case, *Pa*1 and *Pa*2 can be connected by sharing a place (Figure 4, left). Because paths are extended 1-D subsets of the environment, this will not be uncommon. In more complex cases (Figure 4, right), we can search for a connection from *Pa*1 to *Pa*2 using the same heuristics.

It is possible to extend the concept of boundary relation to describe "topological grids," building on the relation between "topologically parallel" paths (Kuipers, 1978, 2001). These larger-scale structures can be very useful for wayfinding.

Two places *A* and *B* may have multiple boundary relations with different paths. If there are several possible boundaries, order them according to the number of boundary relations they have with other places. This will increase the probability of finding a useful connection earlier in the search. It is the second link in the positive-feedback cycle that leads to the emergence of the skeleton.

A POSITIVE-FEEDBACK CYCLE

There is a positive-feedback cycle between the inference of boundary relations and the effect of boundary heuristics on wayfinding search.

- Travel along a path *Pa* makes it likely that a boundary relation, say *left_of*(*Pa*, *dir*, *A*), will be observed and inferred.
- The existence of a boundary relation *left_of*(*Pa*, *dir*, *A*) increases the probability that path *Pa* will be used in the solution to a wayfinding problem even if place *A* is not involved in the route but of course more so if it is.
- Following the newly found route, travel along the path *Pa* increases the probability that a new boundary relation, say *right_of*(*Pa*, *dir*, *B*) will be observed and inferred.

This is a self-reinforcing, "rich-get-richer" process, leading to an oligarchy of paths (the skeleton) rich in boundary relations. The skeleton perpetuates itself because wayfinding most easily finds routes using paths within that subset. Note that there is no qualitative distinction between paths within the skeleton and those outside. There is simply a distribution of boundary relations among the paths in the cognitive map.

The boundary relation hypothesis is:

Hypothesis 1: The empirical phenomenon of the skeleton—that expert wayfinders in an environment preferentially use a small set of important paths—is explained computationally by the positive-feedback cycle between inference of boundary relations during travel and the use of the boundary heuristic during wayfinding.

RESEARCH QUESTIONS

The boundary relation hypothesis suggests a number of computational experiments that can be carried out on a simulated model of a real or artificial urban street network.

Implement a simulated agent that travels from place to place in the simulated environment model. As it travels, it builds its own topological map of the environment, including both connectivity relations between places and paths and boundary relations. The route followed by the agent as it explores the environment can be determined in a number of different ways.

- The agent could explore randomly.
- The agent could follow a route specified by an "oracle" (e.g., the experimenter).
- The agent could attempt to plan a path to a given destination, calling the oracle or exploring randomly when planning fails. Over time, as the cognitive map improves, planning should succeed more and more often.
- The oracle could be the observed behavior of a human participant in a related experiment in the same environment.

Once the agent has learned a cognitive map, we can test for the presence of a skeleton by measuring the frequencies of use of different paths. If there is a highly skewed frequency distribution, with a small set of paths used much more frequently than others, then we conclude that this small set is the skeleton.

A first testable prediction of the boundary relation hypothesis is that paths with larger numbers of boundary relations will be more frequently used in routes. In the following, we present the results of computational and empirical tests of this prediction.

A second more detailed type of prediction is that when faced with a decision among alternate routes, the agent's selection can be predicted from the distribution of boundary relations in the cognitive map. We present a preliminary assessment of this type in the following section.

Another question that can be explored by these methods is the dependence of the structure of the skeleton on the specific geographical properties of the environment (e.g., lengths of streets, density of destination places, presence of bottleneck places that many paths must pass through, etc.) and on the distribution of destinations that determine the agent's experience.

COMPUTATIONAL TEST OF THE HYPOTHESIS

INTRODUCTION

In this section we describe the evaluation of the boundary relation hypothesis from a computational perspective. We report on the results obtained performing the computational experiments suggested in the previous section.

We have implemented the wayfinding algorithm that uses the boundary heuristic and tested it on a virtual environment. We have performed two kinds of experiments; in one we test how well the usage frequency correlates with the number of boundary relations for each path, and in the other we test the model prediction accuracy of how humans find their way in the same environment.

METHOD

To test the boundary relation hypothesis, we have implemented a wayfinding algorithm that can drive a simulated agent through a virtual environment. The wayfinding algorithm uses the boundary heuristic, presented in the previous section. It is given a pair of locations—source and target—and outputs a route[2] through the environment that takes the agent to the target location. While executing this travel, the agent builds a cognitive map of the environment based on the Spatial Semantic Hierarchy (Kuipers, 2000).

The algorithm's performance at the wayfinding task clearly depends on how well it knows the environment, therefore we have adopted a train-and-test approach, with alternating train and test sessions. During a training session, the agent travels through the environment building its cognitive map, while during a testing session, it first tries to build a route from source to target, and after it finds one, it follows it.

Environment and Data Sets. All experiments were conducted in virtual grid-like environments (see Figure 9) that consist of a set of intersections connected by corridors. In a complete cognitive map of these environments, each intersection would correspond to an SSH place and each corridor to a path. From the point of view of the simulated robot, every such place has four views, facing north, east, south, and west.[3] Data were collected by observing the way humans navigate in the same environments.[4]

Experiments. In the first experiment the robot follows the human participant while he or she is randomly exploring the environment. When he or she is presented with a target, the robot tries to find a route from the current position to the destination and follows that route. If none is found, the robot jumps to the next trial. For this experiment, we collect statistics that relate the number of boundary relations per path to their usage frequency in routes that the robot finds from source to target locations.

In the second experiment, the simulated robot takes the passenger's seat. As before, it follows the human in his or her random explorations, but during a test session (i.e., when presented with a target), the robot computes its move corresponding to the first action in a route from the current position to the given target and compares it to what the human does. No matter what its move is, the robot executes the move the human makes. In this way, at every step during the wayfinding task, we are able to measure how well the boundary relation algorithm can predict the human's next move. We also collect data concerning null moves (the algorithm did not find a way to the target).

We have performed the two experiments on four data sets, corresponding to the four human participants in the experiments described in the next section.

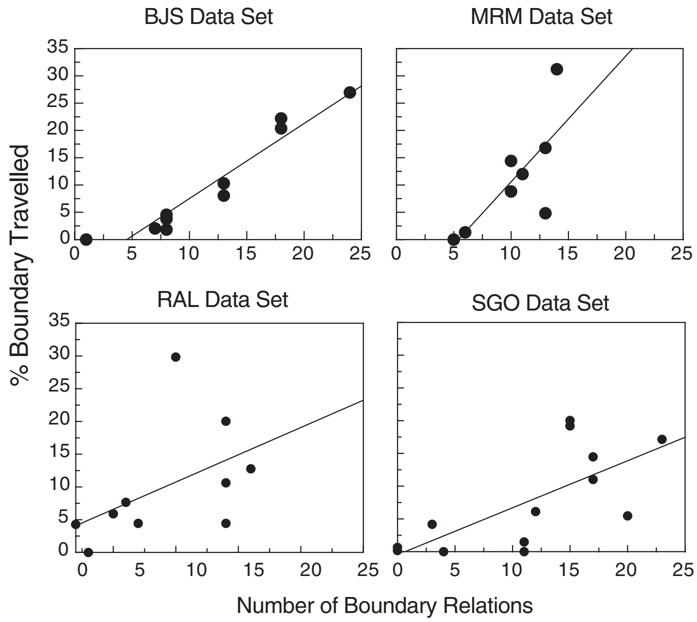

Figure 5: Frequency of Use of Paths in Routes Selected by the Simulated Agent as a Function of the Number of Boundary Relations Acquired by Each Path During Each Participant's Training Phase (Compare With Figure 10 in Which the Routes Are Selected by the Human Participant)

RESULTS

The graphs in Figure 5 draw the usage frequency of paths in travels from source to destination versus their number of boundary relations. The straight line was obtained by fitting a first-degree polynomial in a least squares sense to the data. We trained the robot on the human data obtained while he or she randomly explored the environment and tested on the set of targets provided to the human.

In Figure 6, we draw the accuracy of the model in predicting the participant's next move. At each step of the travel, the model computes a route that would take the agent from the current location to the target location and compares its action to the one the human made. The algorithm has several choices for its actions depending on the particular location the robot is at. There are two classes of actions: translations—move forward one or several steps—and rotations—right, left, and around.

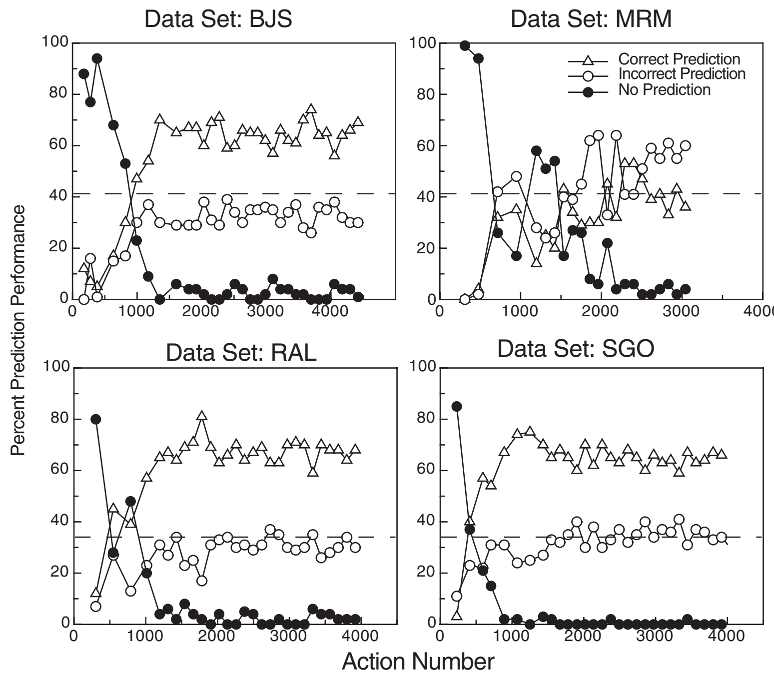

Figure 6: Model Prediction Accuracy of Human Behavior (Computed on Participants BJS, MRM, RAL, and SGO Data Sets)

NOTE: The accuracy of the model prediction as well as the performance of the model at the wayfinding task improves as the number of moves and implicitly the number of boundary relations increases. The dotted line represents the average probability that the model predicts the right move by chance.

The graph has a curve for each of the three possible outcomes: the model predicted the correct move, the model predicted a wrong move, or the model did not predict any move at all.[5]

The dotted line is the average probability that a move chosen randomly by the model is the correct one. This number was computed taking into account all possible states (location in the environment and orientation) and considering the number of correct moves versus the number of possible moves from that state. For the first environment (Figure 9, left) on which the Participant BJS and Participant MRM data sets were collected, this probability is 40.21%, whereas for the second environment (Figure 9, right), used in collecting data for the Participant RAL and Participant SGO data sets, it is 34.75%. This number depends on the complexity of the environment, one factor that directly influences it being the number of +–junctions.

The x-axis represents the number of moves the human/agent made up to that point, whereas the y-axis represents the percentage of the last 100 moves that were correct, respectively, wrong and null. One such measurement is made every 100 moves.[6]

The prediction accuracy is a conservative estimate of the model's real prediction power, as a move suggested by the algorithm is only considered correct if it totally agrees with the move the human made. For example, if from some location and orientation the model suggests a three-step translation[7] but the human decides to do a two-step translation, this is still considered an incorrect prediction.

DISCUSSION

In the computational experiments described here, there are fewer points than in the similar experiments performed on humans (e.g., 10 vs. 13 for the BJS data set and 9 vs. 13 for the MRM data set) (Figure 10). This is due to the fact that the simulated agent did not acquire all the paths that were present in the environment.

For all data sets, there seems to be strong evidence for the fact that there is a positive correlation between the frequency of usage and number of boundary relations for path in the environment.

For the BJS, RAL, and SGO data sets, the accuracy of the model's prediction improves with the number of moves the agent makes through the environment, reaching a 65% to 70% plateau. After 1,500 actions, the level of accuracy is well above the random chance level.

For all four data sets, the performance of the wayfinding algorithm using boundary relation heuristic improves with the number of moves/boundary relations acquired, quickly reaching a 95% to 100% level. In all cases, except the MRM data set, the jump in both prediction accuracy and performance seems to happen between 500 and 1,500 moves.[8] We attribute this behavior to the fact that between these points a lot of random explorations of the environment happen, therefore a lot of boundary relations are acquired and unknown parts of the environment are discovered.

For the MRM data set, the improvement in both prediction accuracy and performance seems to rise much more slowly, and although the performance reaches 95%, the accuracy stays at around 40%. The low accuracy is not the result of null predictions, as this number converges to 5%, but to the generation of predictions that disagree with the participant's moves. This can be attributed to the fact that this participant rarely selected the topologically shortest routes (only 15% of the times).

EMPIRICAL TEST OF THE HYPOTHESIS

INTRODUCTION

Humans provide an important existence proof for the ability to navigate robustly through large-scale spaces. We are able to navigate under a number of different conditions and in a variety of environments. We are also able to learn about a novel environment and easily navigate through that environment at a later time.

An important goal for developing an autonomous robot is to provide an algorithm that allows for robust navigation through a familiar environment along with an algorithm for acquiring useful knowledge about a novel environment. In Section 1, we described a specific theory of how a system might acquire knowledge about a novel environment in addition to how this information may be used for robust navigation. In Section 2, we implemented a computational interpretation of this model to better understand some of the predictions made by this theory. As evidenced in Figure 5, one of the predictions made by this theory is that paths with more relations are going to be selected over paths that have fewer relations.

In the current section we test whether human observers show this same pattern of effects as demonstrated by the computational model. Specifically, we will investigate whether there is a positive correlation between the number of times that an observer travels along a path and the number of boundary relations for that path. In addition, we will investigate whether there is a bias toward taking a route that contains paths with more boundary relations over those with fewer boundary relations. This bias is a strong prediction made by the theory and computational model.[9]

The current study used desktop virtual reality to train and test our observers in a novel indoor environment. Observers navigated through the environment by making key presses. Three different keys were used. One corresponded to translating the participant one hallway forward, and the other two rotated the participant by 90° clockwise or counterclockwise.

The study was conducted using a training-and-test procedure. The training procedure was designed to allow the participants to explore the environment without any goals or expectations. During a training session, participants were allowed to navigate freely through the environment. While navigating, observers heard an auditory signal informing them that they were at a particular target position (e.g., “Position 4”). During the testing session, participants moved from one target location to another. To do this, participants were instructed with an auditory signal about which target to move to (e.g., “Go to Position 3”). When they arrived at the goal location they were

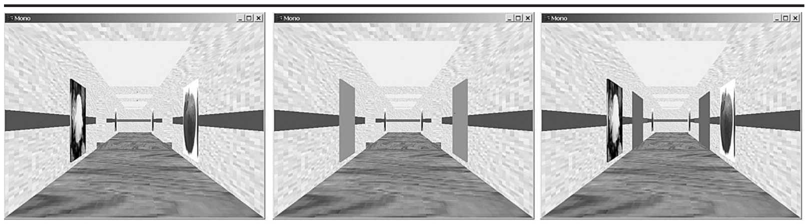

Figure 7: A Sample View of a Hallway.
NOTE: In addition to the two walls and a ceiling in each hallway, there was also a "railing" and "ceiling lights" to give the observer perspective. On the left and right wall of each hallway was a picture. The figure on the left shows an intact image during spatial navigation. During the experiment, the participants were queried about the specific properties of the environment. This query consisted of covering up either the pictures (center) or the hallways at the next intersection (right).

informed of this (e.g.,"Position 3") and then given a new goal location (e.g., "Go to Position 1"). Using the data from the training and testing sessions, we computed an estimated cognitive map including a set of boundary relations for each participant. We also calculated how often the participant chose a particular path when traveling from one target location to another.

METHODS

Participants. Two male and two female participants volunteered for this experiment. Participant BJS is a 34-year-old man who is an author on the current article.[10] MRM is a 22-year-old male, and SGO and RAL are both 22-year-old female students at the University of Texas. All participants had normal or corrected-to-normal vision. Participant SGO was paid $8/hour for her participation.

Materials. The study was conducted on a Dell Presario computer with an Intel IV 2.0 GHz processor. The display was a Dell P1130 21-inch CRT monitor with a NVIDIA GeForce MX video graphics card that had 32 MB of memory.

The environments were specified using virtual reality modeling language (VRML). These environments were rendered using Vizard software (WorldViz). Participants viewed the environments from a first-person perspective (see Figure 7) and moved through the environment using key presses. Each hallway segment contained two pictures, one on each hallway wall. The left panel of Figure 7 shows a sample rendering of the environment from the observer's perspective. Although there were 80 pictures within the

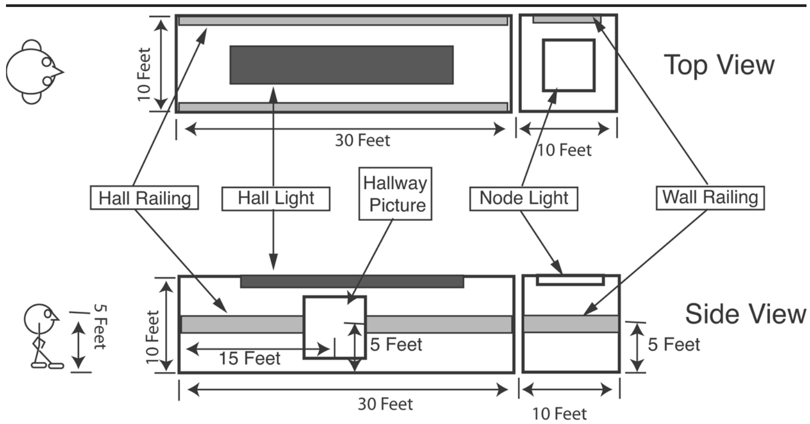

Figure 8: An Illustration of a Hallway Segment, Hallway Node, and Node Wall
NOTE: Each hallway and node in the environment used these positions and sizes.

environment (2 pictures per hallway × 40 hallways), there were only two unique photos used. Each picture was either a picture of a fish or of an apple.

Each hallway in the environment was composed of intersecting hallway segments and hallway nodes. With the exception of the pictures on the walls, each hallway segment was identical to all of the other hallway segments. The hallway nodes were used to combine multiple intersecting hallway segments. The specification of the hallway segments, hallway nodes, and node wall can be seen in Figure 8.

The configuration of hallway segments in the environment was generated by the computer using a random-layout generating program. This program produces layouts with hallways that are at right angles to one another. To generate an environment, the program uses three parameters: the number of hallway segments and the maximum extent of the environment in the X and Z dimensions. The environment used in the current study was constrained to 40 hallway segments, the maximum X extent was 25 hallways, and the maximum Z extent was 25 hallways.

To generate the environment, the program used the following algorithm:

1. Select a random hallway from the 25 × 25 grid.
2. Calculate the set of potential hallways that can connect to the initial hallway.
3. Randomly select one of these potential hallways.
4. Compute the set of hallways that can connect with the previously selected hallway. Add these hallways to the potential hallway list (each hallway is listed only once).
5. Continue with Steps 3 and 4 until the number of hallways selected is 40.

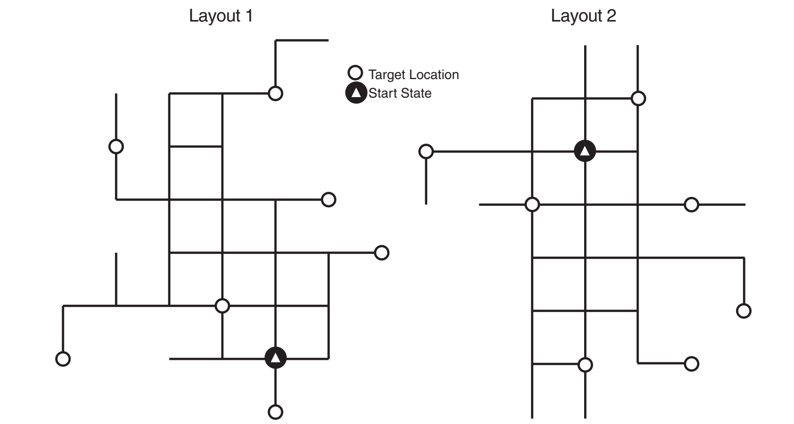

Figure 9: The Study Used Two Environments Shown Here.
NOTE: The circles indicate the target locations in each environment. The participants started each session from the same start state indicated by the start state symbol.

The maps used for the current study using this algorithm are shown in Figure 9.

Seven target locations were selected by the experimenters. These locations were selected using the following general principles: (a) The set of target locations were reasonably distributed throughout the environment, and (b) the target places would have at least one T-junction, one dead end, and one +–junction. The layout and the location of the seven target locations are shown in Figure 9 as circles.

Each target location was assigned a number from 1 to 7. The participant received an auditory signal by the computer when they walked over any of the target locations (e.g.,"Position 5"). There were no visual cues indicating which positions were the target locations and which were nontarget locations. The auditory signal provided a local cue for identifying the target locations.

Participants navigated through the environment by making key presses. The participant had three different actions to choose from: translate forward one hallway segment, rotate clockwise 90°, and rotate counterclockwise 90°. Both the rotation and translation produced the appropriate optic flow for the action. All three actions took less than one second to complete.

Design and procedure. The study was conducted over multiple days. Participants navigated through the environment by making key presses. Participants were instructed on how to navigate through the environment and were given a series of practice trials on another novel environment to familiarize themselves with the procedure.

Before beginning the first session, the participants were told that the purpose of the study was to understand how we acquire knowledge about an environment and how that knowledge is retained over time. Participants were informed that they were going to navigate through an unfamiliar virtual environment and that there were seven target locations in the environment. The participants were told to remember the positions of these target locations in the environment because later they would be required to move to the different target locations.

On the first day, participants engaged in 10 training sessions and 10 test sessions. These sessions alternated starting with a training session.

During a training session, participants explored the environment freely. The participants were instructed to learn the locations of the seven target locations in the environment because later they would be expected to move from one target location to another. The exploration session terminated after the participant made 100 translations within the environment.

During the test phase, participants started at the start state (see Figure 9). Standing at the start state, the computer gave an auditory signal indicating which target location was the current goal ("Go to position X," where X was a randomly chosen number between 1 and 7). When the participant arrived at the goal location they received an auditory signal indicating that they had arrived at the location ("Position X") followed by an auditory signal indicating the next goal location ("Go to position Y," where Y was a number from one of the six other target locations). If the participant walked over one of the other target locations on the way to the current goal location the computer provided an auditory cue indicating that they were at that target location.

While traveling to the goal location, the computer queried the participant about the hallway structure at the next intersection or the set of pictures in the current hallway (see center and right panels in Figure 7). These queries occurred randomly during the test sessions. The following things happened when the participant was queried about the pictures or the hallway structure:

1. An auditory signal was given (the sound of a bell).
2. Either the pictures in the hallway or the next intersection was covered with a virtual tarp (see the center and right panels of Figure 7).
3. Participants entered their response into the computer using the keyboard about what they thought the pictures were under the tarps or the next intersection.
4. After entering the response, the tarps were removed and the participant continued to the goal location.

These queries were done to address an empirical question outside the scope of the current article and will be published in the future.

The exploration phase terminated after 25 hallway and 25 picture queries were completed.

RESULTS

One of the predictions made by the boundary relation hypothesis is that participants should choose paths that have more boundary relations over paths with fewer boundary relations. To test this prediction, we computed the frequency that participants traveled a particular path as a function of the number of boundary relations that the path possessed.

Computing the boundary relations. Boundary relations are specified by a set of three distinct places (A, B, and C). These three locations have distinct roles in the boundary relation calculation. *PlaceA* is the starting location. *PlaceB* is the decision point where a rotation occurred (in the current experiment, this is a 90° rotation to the right or left). *PlaceC* is the end point.

The current analysis used this ABC function to define the set of boundary relations for the two participants. The ABC relations were defined using the following algorithm:

Consider all triples (A, B, C) of places visited by the participant, where the participant

- started or made a 90-degree turn at A,
- traveled from A to B with an unbroken sequence of translations,
- made a 90-degree turn at B,
- traveled from B to C with an unbroken sequence of translations,
- made a 90-degree turn or terminated at C.

(AB)C rule: Store PlaceC with a boundary relation to Path AB.
If C had not already been stored, increment the number of boundary relations of Path AB by 1.
(BC)A rule: Store PlaceA with a boundary relation to Path BC.
If A had not already been stored, increment the number of boundary relations of Path BC by 1.

Figure 10 illustrates the percentage of times that each participant traveled a path as a function of the number of boundary relations.[11] This graph illustrates the basic effect that participants selected paths that had more boundary relations over those with fewer boundary relations while traveling in these environments. The results are very similar to those found in the computational modeling section of this article (see Figure 5).

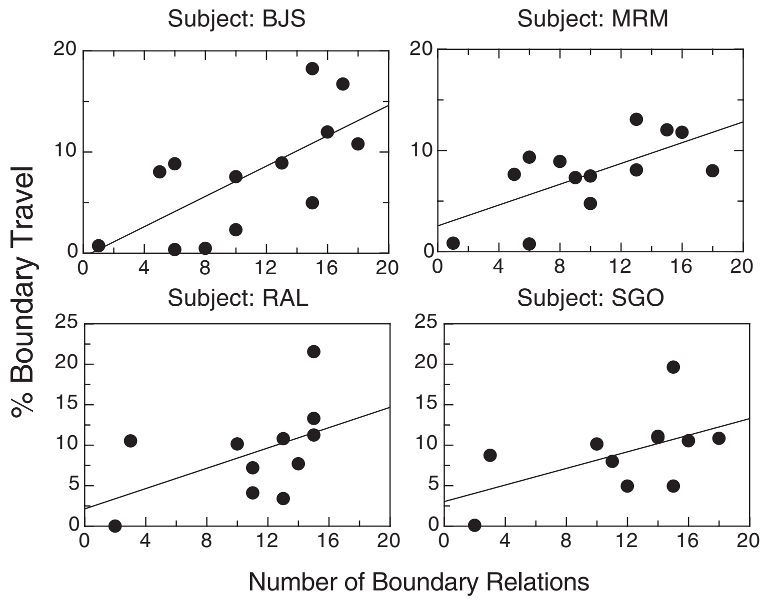

Figure 10: Frequency of Use of Paths in Routes Selected by the Human Participant as a Function of the Number of Boundary Relations Acquired by Each Path During Each Participant's Training Phase (Compare With Figure 5 in Which the Routes Are Selected by the Simulated Agent)

Path selection bias. A strong prediction made by the boundary relation hypothesis is that an observer should be biased toward selecting routes that have paths with more boundary relations over those with fewer boundary relations. To investigate this question, we started by computing the set of topologically shortest routes between each of the target locations. For many of these tasks (starting location to goal location), there were multiple routes that were topologically shortest. The left side of Figure 11 shows the percentage of times that each participant took one of the topologically shortest routes.[12] For three of the observers, a significant proportion of the selected routes were one of the topologically shortest routes (approximately 75% of the routes). One participant (MRM), however, rarely selected a topologically shortest route (only 15% of the time) and therefore will not be included in the following analysis.[13]

For many of the source-to-goal tasks there were multiple routes that were shortest and were topologically shortest. Typically these different routes differed by one path. We were interested in understanding whether participants

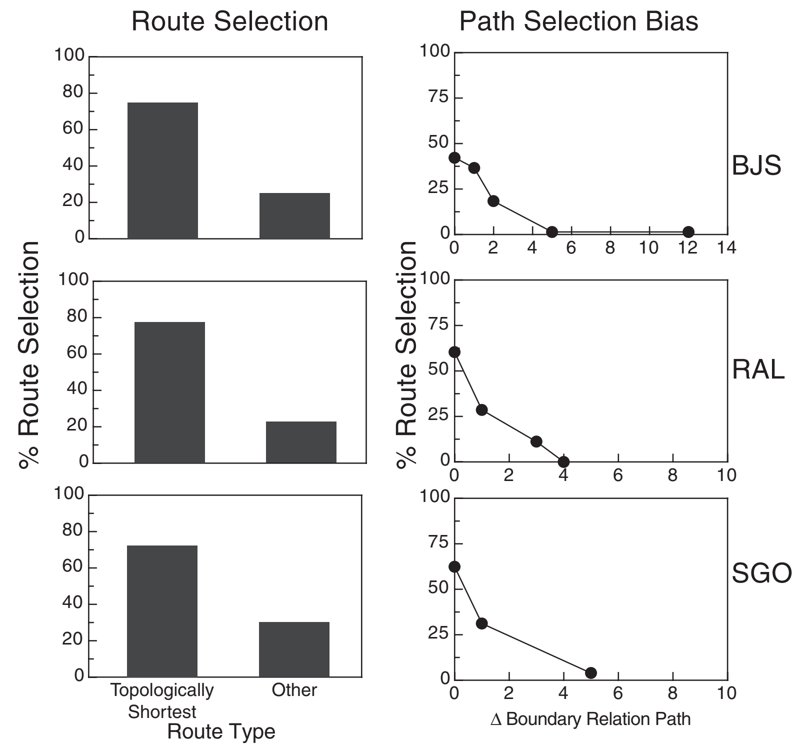

Figure 11: Participants' Route Choices

NOTE: The percentage of times that each participant chose the topologically shortest route between two places and other routes is shown on the left. For many destinations there existed more than one topologically shortest route. On the right is a graph illustrating when the participant selected the route with the most boundary relations (Δ boundary relations path = 0) versus the routes with 1 less than the maximum boundary relation path (Δ boundary relations = 1), up to 12 fewer than the maximum number of boundary relations.

(a) randomly selected one path over another or (b) were biased toward selecting paths with more boundary relations.

To answer this question we computed the set of boundary relations for each path based on the participants' experiences in the environment. For the paths that possessed more than one topologically shortest route we found the set of paths that differentiated these different routes (as mentioned before, these different routes typically differed by one path). We then classified these routes by the relative number of boundary relations that the differing paths possessed. For each set of paths there was one path that had the maximum number of boundary relations. This route was given a Δ value of 0. Some paths differed from the maximum boundary relation path by 1. We assigned this route a Δ value of 1 (1 less than the maximum). We computed the Δ value for each route from each source location to each target location that had more than one topologically shortest route.

If participants are biased to routes that have paths with more boundary relations we should find that participants choose the routes with a Δ value of 0 (maximum number of boundary relations). The right side of Figure 11 provides an illustration of the data. Participants selected the route with the maximal number of boundary relations most often (between 42% and 62%). They selected the maximal route or a route with one fewer boundary relations between 79% and 93% of the time. These data suggest that when participants are faced with multiple routes of the same length (or cost), they will select paths that have more boundary relations over those with fewer boundary relations.

The results in Figures 10 and 11 show that participants are biased toward selecting paths with more boundary relations over those with fewer boundary relations. Figure 10 shows a positive correlation between the number of boundary relations that a path possesses and the frequency that a participant travels that path. Figure 11 shows that when there are multiple routes with an equivalent cost associated with them, participants are biased toward selecting routes that contain paths with more boundary relations over those with fewer boundary relations. Both of these findings are consistent with the boundary relation hypothesis.

DISCUSSION

The current empirical test evaluates an important prediction of the boundary relations hypothesis. Namely, a human observer will be biased toward selecting paths that have more boundary relations over those paths with fewer boundary relations.

We tested this prediction by training 4 participants to navigate through a simple virtual indoor environment. For each path in the environment we calculated the set of boundary relations that the participant might have inferred given their experience with the environment. We found that there was a positive correlation between the frequency with which a participant used a path and the number of boundary relations that the participant could have inferred about the path (see Figure 10).

These data support the use of boundary relations in selecting a path while navigating through a large-scale space. However, it is possible that the routes with the most boundary relations are also those routes that provide the best, or most efficient, paths between the multiple goal states in the environment. To evaluate whether participants were biased toward selecting paths with more boundary relations we evaluated trials in which there was more than one topologically shortest path between a start state and a goal state. We hypothesized that if participants are influenced by the number of boundary relations that a path had, they should be biased toward selecting paths with more boundary relations over paths with fewer boundary relations. We found that participants selected the paths with the most boundary relations between 42% and 62% of the time.

GENERAL DISCUSSION

The boundary relation hypothesis says that the empirical phenomenon of the skeleton—that expert wayfinders in an environment preferentially use a small set of important paths—is explained computationally by a positive-feedback cycle between inference of boundary relations during travel and the use of the boundary heuristic during wayfinding. Basically, the rich get richer: A path that is already rich in boundary relations is more likely to be used in routes and will thus get richer by acquiring new boundary relations.

We can test this hypothesis computationally by building a simulated agent to explore a simulated environment. Because we can examine its internal state, we can inspect the cognitive map that it creates as a result of its exploration, and we can count the number of boundary relations acquired by each path. Then we can observe its behavior in response to wayfinding problems and determine how the frequency of use of a particular path is related to the number of boundary relations it has.

Applying the same analysis to human participants is more difficult because we have limited ability to examine their internal state. We forge ahead nonetheless, with a three-stage method:

1. Collect data from human participants as they learn about and navigate through a novel environment.
2. Use the participant's sequence of actions to estimate the structure of the cognitive map and the set of boundary relations learned by this individual participant.
3. Based on the estimated cognitive map and its set of boundary relations, predict individual wayfinding decisions as the participant travels to specified places.

Using this method, we can estimate each participant's cognitive map, including the set of boundary relations acquired during the experiment. Based on this estimate of each participant's cognitive map, the model predicts the participant's behavior at each decision point when finding and following a route to a goal state. We then compare the predicted action with the action the participant actually took in the environment. Random guessing at each decision point would yield 35% to 40% accuracy overall, depending on the specifics of each environment. The model was able to predict individual action decisions with 65% to 70% accuracy (Figure 6).

It is important to note that the computational implementation we tested is a relatively simple instance of the boundary relation hypothesis. For example, the estimated cognitive map depends on a particularly simple model of learning of boundary relations. The model only draws inferences from ABC triplets of adjacent decision points along a route. This restriction would tend to underestimate the set of boundary relations in the cognitive map. On the other hand, the model also assumes that every ABC triplet is stored in memory and successfully retrieved when needed. Research on human memory clearly shows that not every experience is stored in memory or accurately retrieved. This suggests that the model will tend to overestimate the set of boundary relations in the cognitive map. We expect that further investigation will improve the quality of the estimate of the cognitive map and the set of boundary relations.

At this stage of investigation, with plenty of opportunities to refine the elements of the model, we consider 65% to 70% accuracy at predicting individual decisions to be very promising.

PATH SELECTION BIAS

In addition to the rich-get-richer prediction of the boundary relations hypothesis, there is a second important prediction: When a participant is faced with multiple routes with the same travel cost, the routes including paths with more boundary relations will be selected.

We can use the estimated cognitive map and set of boundary relations to test this prediction. For each start and goal state in the environment, we computed the set of topologically shortest routes. For some start/goal state pairs, there existed multiple topologically shortest routes. According to the hypothesis, participants should be biased toward those routes that have paths with more boundary relations. We found that participants selected the routes that had the maximal number of boundary relations 42% to 62% of the time, and they selected the route with the maximum number of boundary relations or a route that had one fewer than the maximum number of boundary relations 79% to 93% of the time (Figure 11). This supports the hypothesis that boundary relations play an important role in influencing path choice during wayfinding.

SUMMARY AND CONCLUSIONS

Even our current simple model gives us a surprisingly good ability to predict individual decisions by human participants. We have provided strong evidence for the use of boundary relations in spatial navigation by using a computational model that estimates the cognitive map of a human participant based on behavioral data and then uses the estimated map to predict behavioral choices. We have also shown strong evidence for a bias during wayfinding toward selecting paths with more boundary relations over those with fewer boundary relations.

It still remains to be seen, through future work, whether the emergence of the skeleton is explained by the rich-get-richer positive-feedback cycle between boundary relation acquisition and the use of the boundary heuristic in wayfinding or by some other mechanism.

We plan to extend the model with (a) boundary relation inference rules not restricted to triplets of adjacent decision points; (b) a probabilistic model of successful inference, storage, and retrieval; (c) an improved wayfinding algorithm; and (d) explicit models of individual variation separate from variation among participants in training experience.

One contribution of this work is a computational model and preliminary evaluation explaining the emergence of the skeleton in the cognitive map, a significant phenomenon in spatial cognition. A second contribution is a demonstration of a method for testing a fine-grained cognitive hypothesis against human behavior by creating an estimated model of the cognitive state of each individual participant. We believe that both contributions warrant further investigation.

NOTES

1. Example: When traveling east along the Charles River separating Boston and Cambridge, Boston is topologically to the right. However, because of the curve of the river, the highly visible John Hancock Tower in Boston can sometimes be seen in the distance to the traveler's left. The Boston area is a treasure trove of spatial paradoxes for the cognitive map theorist.

2. Sequence of translations and rotations.

3. Even if a place has only one, two, or three neighbors, it still has four views, some of them possibly facing walls.

4. For a detailed description of the experiment involving human participants from which data was collected, see the next section.

5. This happens when no route to the target was found.

6. The data points are not evenly spaced because these statistics are only collected during testing (i.e., travels that have a target).

7. Although in the experiment the participants were doing one-step translations, for the purpose of this analysis, successive one-step translations were merged into one long translation.

8. This number is dependent on the data set and on the structure of the environment.

9. The data set used here is from an experiment investigating another issue that will be described in another research article. However, after collecting the data we realized that we could begin to address the use of boundary relations with this data set.

10. Although Participant BJS is an author on this study, he did not know the layout before starting the study. A research assistant generated the environment and tested Participant BJS.

11. This graph shows this relationship using both the (AB)C rule and the (BC)A rules. There was no significant difference between using just the (AB)C rule and combining the two rules. The rest of the empirical data uses boundary relations using the (AB)C and (BC)A rules.

12. These were defined as the routes that would get the observer from the source state to the goal state in the fewest number of translations and rotations

13. The reason for excluding Participant MRM from this analysis is because there were only a small number of trials in which he took the topologically shortest route. With only a few data points it is difficult to show support or refute the existence of a selection bias.

REFERENCES

Chase, W. G. (1982). Spatial representations of taxi drivers. In D. R. Rogers & J. A. Sloboda (Eds.), *Acquisition of symbolic skills*. New York: Plenum.

Elliot, R. J., & Lesk, M. E. (1982). Route finding in street maps by computers and people. In *Proceedings of the Second National Conference on Artificial Intelligence* (AAAI-82, pp. 258-261). Cambridge, MA: AAAI Press/MIT Press.

Golledge, R. G. (1999). *Wayfinding behavior: Cognitive mapping and other spatial processes*. Baltimore: Johns Hopkins University Press.

Kuipers, B. J. (1978). Modeling spatial knowledge. *Cognitive Science*, *2*, 129-153.

Kuipers, B. J. (1982). The "map in the head" metaphor. *Environment & Behavior*, *14*, 202-220.

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

Kuipers, B. (2001). The skeleton in the cognitive map: A computational hypothesis. In J. Peponis, J. Wineman, & S. Bafna (Eds.), *Space syntax: Proceedings of the Third International Symposium* (pp. 10.1-10.7). Ann Arbor: University of Michigan.

Lynch, K. (1960). *The image of the city*. Cambridge, MA: MIT Press.

Remolina, E., & Kuipers, B. (1998). Boundary region relations. In *Cognitive robotics, papers from the 1998 AAAI fall symposium* (AAAI Technical Report FS-98-02, pp. 117-124). Menlo Park, CA: AAAI Press.

Siegel, A. W., & White, S. H. (1975). The development of spatial representations of large-scale environments. In H. W. Reese (Ed.), *Advances in child development and behavior* (Vol. 10, pp. 9-55). San Diego: Academic Press.

Timpf, S., Volta, G. S., Pollock, D. W., & Egenhofer, M. J. (1992). A conceptual model of wayfinding using multiple levels of abstraction. In A. U. Frank, I. Campari, & U. Formentini (Eds.), *Theories and methods of spatio-temporal reasoning in geographic space* (Vol. 639, pp. 348-367). New York/Berlin: Springer-Verlag.