Journal of Experimental Psychology:
Human Perception and Performance
2006, Vol. 32, No. 3, 688–704

0096-1523/06/$12.00 DOI: 10.1037/0096-1523.32.3.688

Lost in Virtual Space: Studies in Human and Ideal Spatial Navigation

Brian J. Stankiewicz
University of Texas at Austin

Gordon E. Legge
University of Minnesota, Twin Cities Campus

J. Stephen Mansfield
State University of New York College at Plattsburgh

Erik J. Schlicht
University of Minnesota, Twin Cities Campus

The authors describe 3 human spatial navigation experiments that investigate how limitations of perception, memory, uncertainty, and decision strategy affect human spatial navigation performance. To better understand the effect of these variables on human navigation performance, the authors developed an ideal-navigator model for indoor navigation whose optimizing algorithm uses a partially observable Markov decision process. The model minimizes the number of actions (translations and rotations) required to move from an unknown starting state to a specific goal state in indoor environments that have perceptual ambiguity. The authors compared the model's performance with that of the human observer to measure human navigation efficiency. Experiment 1 investigated the effect of increasing the layout size on spatial way-finding efficiency and found that participants' efficiencies decreased as layout size increased. The authors investigated whether this reduction in navigation efficiency was due to visual perception (Experiment 2), memory, spatial updating strategy, or decision strategy (Experiment 3).

Keywords: spatial navigation, ideal observer, spatial memory, decision making, Bayesian

After starting a new job in a new building, you need to develop an understanding of frequently visited locations within your new environment. Locations like the closest restroom to your office, the stairwell, the mailroom, the photocopy room, and other locations are all important places in your everyday activities. Rather than studying a map, someone may show you where some of these locations are, or you may simply find them through your own exploration. After a period of time (perhaps a few weeks) you typically have woven these experiences into an internal representation of the space that allows you to move effectively from one location to another. This internal representation is usually referred to as a *cognitive map* (Hirtle & Heidorn, 1993; Kuipers, 2001; O'Keefe & Nadel, 1978; Tolman, 1948).

The concept of a cognitive map has influenced how researchers describe the ultimate representation that people obtain after extensive exploration of a large-scale space. However, Tolman's (1948) introduction of the concept did not provide an explicit description of how one generates this cognitive map from experience or what information is made explicit in the map. Furthermore, it is not clear how one accesses this information while navigating through a familiar environment.

Later research by Siegel and White (1975) provided a description of the metamorphosis of a cognitive map as a person becomes familiar with a specific environment. According to Siegel and White, a cognitive map begins as a set of landmarks. This is followed by a representation that makes explicit a sequence of actions needed to get from one location to another, called *routes*. Finally, these sequences of actions (or routes) are united into an internal representation of the environment that is more maplike, called a *survey representation*. This survey representation is useful because it can be used to generate novel routes between two locations within the environment.

Although Siegel and White (1975) proposed that there is a sequence of stages in the development of a cognitive map, a clear understanding of the specific processes underlying the acquisition and access of that map remains uncertain. It should be noted that there has been a great deal of research into what information is stored in the cognitive map (Cohen & Schuepfer, 1980; Franz, Schölkopf, Mallot, & Bülthoff, 1998; Gillner & Mallot, 1998; Hirtle & Hudson, 1991; Hirtle & Jonides, 1985; Hirtle & Kallman, 1988; Hirtle & Mascolo, 1986; Jacobs, Thomas, Laurance, & Nadel, 1998; Klatzky et al., 1990; Kuipers, 2000, 2001; Mallot, Franz, Schölkopf, & Bülthoff, 1997; Mallot & Gillner, 2000; McNamara, Hardy, & Hirtle, 1989; O'Keefe & Nadel, 1978; Ruddle, Payne, & Jones, 1997; Schölkopf & Mallot, 1995).

To develop a better understanding of human spatial navigation, it is important to understand how the cognitive map is developed

Brian J. Stankiewicz, Department of Psychology, University of Texas at Austin; Gordon E. Legge and Erik J. Schlicht, Department of Psychology, University of Minnesota, Twin Cities Campus; J. Stephen Mansfield, Department of Psychology, State University of New York College at Plattsburgh.

This research was supported by National Institutes of Health Grant EY02857 to Gordon E. Legge, Air Force Office of Scientific Research Grant FA9550-04-1-0236 to Brian J. Stankiewicz, National Institutes of Health Grant EY016089 to Brian J. Stankiewicz, and a University of Texas Summer Research Grant to Brian J. Stankiewicz. We thank Joshua Gefroh for his help in collecting data in Experiments 1 and 2, Matthew McCabe for his help in collecting data in Experiment 3, and Erin Dies for helping collect additional data in Experiment 2. We also thank Kim Pleticha and Anthony R. Cassandra for their comments on an early draft of the article.

Correspondence concerning this article should be addressed to Brian J. Stankiewicz, Department of Psychology, University of Texas, 1 University Station A8000, Austin, TX 78712-0187. E-mail: bstankie@mail.utexas.edu

and the nature of its representation; it is also important to understand the processes and strategies used in accessing and utilizing the information within the cognitive map. The present studies investigate issues associated with accessing information from the cognitive map. Specifically, these studies investigate how effectively participants use the information available in goal-directed navigation through familiar environments. In these studies, human performance is compared with the performance of a computer model, which makes optimal use of the available information. We use a measure termed *efficiency*, which is an index for comparing human performance with optimal performance. The computer model, implementing an optimizing principle, is in the tradition of ideal-observer models in perception. In our case, we refer to it as an *ideal-navigator model*.

The findings from these studies provide us with important information about how effectively people use information in the cognitive map. In Experiment 1 we measured human navigation efficiency in environments that vary in their size. The primary finding is that human navigation efficiency declines as the environments become larger. We hypothesize that this size-dependent inefficiency may arise from one of four important processes involved in spatial navigation: perception, accessing the cognitive map, spatial updating, or decision strategy. Experiment 2 investigated whether perception limited the participant's ability to choose an efficient route through the environment. The findings suggest that perception was not the limiting factor. Experiment 3 investigated whether the participant's inefficient behavior might be due to limitations in their ability to access their cognitive map, update their current state under conditions of uncertainty, or make decisions with uncertainty. By providing the participant with different types of information in different conditions, Experiment 3 found that most of the inefficiency found in Experiment 1 appears to be due to an inefficiency in the participant's spatial updating procedure.

It is important to appreciate that these studies investigated a human observer's navigation efficiencies—that is, how well did the participant move relative to an ideal observer with perfect perceptual processing, perfect map memory, and the ideal decision strategy. By comparing human performance with that of the ideal observer, we are able to take into account variations in human performance that are due to *task difficulty*. This is because the ideal observer gives us the very best performance in the task. Because of this, the ideal observer is only sensitive to factors associated with task demands; any change in performance above and beyond these task demands can be attributed to the cognitive and perceptual processing limitations of the human observer. In the following section, we describe the ideal navigator used in these studies.

Formalization of a Spatial Navigation Task

This section describes the key properties of our model's spatial navigation task. Later we build on the formalism to describe an ideal navigator that uses principles from *partially observable Markov decision processes* (POMDP; Cassandra, Kaelbling, & Littman, 1994; Chung, 1960; Kaelbling, Littman, & Cassandra, 1998; Sondik, 1971).

Perception

An observer navigating through a complex environment receives perceptual input in many different forms, including two-dimensional visual images (and possibly stereo depth if the observer has two eyes), auditory cues, tactile cues from the terrain, vestibular input, kinesthetic feedback from joints and muscles, and processed forms of these perceptions that allow for path integration. At any given moment during a navigation task, one can specify all of this information in a high-dimensional perceptual vector (**P**). The perceptual vector is generated by specifying a specific state in the environment (s_i) and the observation function (Γ) that converts the physical properties of the environment into a perceptual vector. The perceptual vector makes explicit all of the information that is available to the observer at any given moment during a navigation task:

$$p_i = \Gamma(s_i). \tag{1}$$

In addition to the perceptual vector, we can also specify a high-dimensional observation vector (**O**) that dictates which aspects of the perceptual vector are used during spatial navigation. This observational vector may store image-based data, a visual "snapshot" image taken at a particular state or an abstract representation, such as a list of recognized observable objects from a particular state (Gillner & Mallot, 1998; Mallot et al., 1997; Mallot & Gillner, 2000). One important aspect of comprehending human spatial navigation behavior is to understand what information is made explicit or encoded given a specific perceptual vector.[1] That is, what people encode and what is perceptually available may not be the same. Ψ is the function that converts the perceptual vector into the stored observation vector given by Equation 2:

$$O_i = \psi(p_i). \tag{2}$$

Spatial Updating

Spatial updating is one's ability to determine one's location and heading in a large-scale space given one's knowledge about the environment (cognitive map) and the sequence of specific observations and actions while navigating. Spatial updating may rely on a number of perceptual inputs, including proprioception, vision, and audition. One form of spatial updating might simply include integrating the vestibular and proprioceptive cues associated with movements through the environment (perhaps for path integration). In addition to these nonvisual inputs, vision seems to play an integral role in one's ability to update and move through an environment. The objects within the environment and the structure of the environment can give important cues as to one's current location and the set of actions that one needs to take to reach one's destination.

[1] An interesting example of this is the work in change blindness (Simons, 2000). Here, participants observe an image for a brief period and are able to "perceive" the entire image. Following a brief, blank screen, a second display is shown in which something significant has been changed between the two images (e.g., person changes, car color changes, building is moved, etc.). This sequence of Display A, blank, Display A′ continues until the participant identifies the change. It typically takes many of these displays to finally "observe" or encode the information despite it being perceptually available.

Decision Strategy

A navigation system has available to it a set of actions (A) that can move it from one state in the environment to another. An action might be as complex as "follow the center of a hallway until it ends" or as simple as "rotate clockwise by 90°." In defining a navigation task, one needs to specify the set of actions that are available to the navigator and how these actions are selected.

Navigation Goal

To formalize a spatial navigation task, one needs to specify the goal. Different tasks generally require different goals. For example, the most common spatial navigation goal is to travel from one known location (like one's office) to another known location (perhaps the mailroom). Other times, the goal might be to figure out where one is located in the environment after getting lost. Another goal might be to determine which large-scale space one is located in out of a set of possible environments.[2] Each of these goals may produce different behaviors given identical perceptual input, actions, decision strategy, and spatial updating procedures.

Ideal-Observer Modeling

An *ideal-observer model* provides optimal performance given the information available in the task. Typically, ideal observers are not proposed as models of human perception or cognition. Instead, the ideal observer provides a benchmark by which to compare human performance. More specifically, these models illustrate what optimal performance should look like. When human performance matches that of the ideal-observer model, one can conclude that the human is making use of all of the information in the task. When the human underperforms the ideal observer, specific discrepancies between the human data and the ideal data may illuminate the constraints imposed by the human information-processing system.

Ideal-observer analysis has been used to understand perceptual functions from the quantum limits of light detection (Hecht, Shlaer, & Pirenne, 1942) to many forms of visual pattern detection and discrimination (Geisler, 1989), to reading (Legge, Hooven, Klitz, Mansfield, & Tjan, 2002; Legge, Klitz, & Tjan, 1997), to object recognition (Liu, Knill, & Kersten, 1995; Tjan, Braje, Legge, & Kersten, 1995; Tjan & Legge, 1998), eye movements (Najemnik & Geisler, 2005), and also in reaching tasks (Trommershäuser, Gepshtein, Maloney, Landy, & Banks, 2005). In the present studies we were interested in understanding the cognitive limitations of human spatial navigation. To do this, we developed an ideal-navigator model. This model uses an optimal algorithm to solve a spatial navigation task. We can compare the performance of this model with the performance of human participants on the identical task. The model has available the same information that is available to the human participants.

In the following section, we formalize an ideal-observer model for a specific spatial navigation task. The model uses principles from POMDP theory (Cassandra et al., 1994; Kaelbling et al., 1998) to navigate through familiar indoor environments. The goal of the model is to travel from an unknown starting point to a known target location using, on average, the fewest number of actions. We use the model to measure the expected change in performance due to task demands.

Ideal Navigator

We describe the ideal navigator in two sections. The first section provides an intuitive appreciation of what the navigator is doing, and the second provides a more formal description of the model.

Intuitive Explanation

The model has perfect knowledge of the environment. This knowledge comes in two forms. First, the model knows what it expects to see from every state within the environment. Second, the model knows the connection matrix for the entire environment. In other words, the model knows exactly where it will end up if it executes a particular action at a particular place and orientation in the environment. Given that it knows what it expects to see from every state in the environment, it will also be able to determine what it expects to see in the new state following the action.

The model starts from an unknown location (i.e., placed at a random location within the environment) and is instructed to move to a known target location in as few actions as possible. Here are the steps that the model takes to choose the next action:

1. *State elimination.* Compare the current view with the views that the observer would expect to see in the set of potential states.[3] Eliminate the states from the set of potential states that are not consistent with the current view. A set of candidate states remains. If there is only one candidate state, there is no remaining state uncertainty.

2. *Route generation.* From each of the candidate states compute the shortest route that starts with each action from the set of available actions (i.e., rotate-left, rotate-right, and forward). The route has to reach the goal state in the fewest number of actions with no remaining uncertainty.

3. *Action cost.* Step 2 provides a list of routes from each candidate state that starts with a particular action. For each action, compute the average number of moves that it would take to reach the goal state from each of the states. This will provide up to three averages (one for each action).

4. *Action selection.* Choose the action that has the minimum number of moves on average (from Step 3).

5. *Candidate state updating.* Update the set of potential states by computing what state the observer would be in if it executed the selected action from Step 4.

[2] This can happen when one exits an elevator on the wrong floor or emerges from a train station at an unknown stop.

[3] Initially, the set of potential states is all of the states in the environment. However, after the initial view and set of actions, the model eliminates certain states from consideration. Those that remain are the "potential states."

6. *Action execution.* Execute the action in the real environment.

7. *Encode observation.* Encode the new observation.

8. *Check if done.* Check if at target location with no remaining uncertainty: "Yes" indicates done; "no" indicates return to Step 1.

This algorithm minimizes the set of actions *on average* to reach the goal state. Given that there is uncertainty about the observer's position in the environment, the model does not always take the direct route from the starting state to the target state. Typically, this situation occurs when the best move from most of the candidate states is one move (e.g., rotate-left), but the best move from one of the states is different (e.g., move forward). Given the uncertainty, the best move is rotate-left. However, if the observer is starting from this atypical state, the model may need to backtrack after a few moves.

Simple example. We would like to provide a simple illustration of how the underlying algorithm actually works. To do this, we work through a few steps in the environment illustrated in Figure 1. The left side of this illustration shows the *ground truth* of the current state of the problem. The observer is at location E facing south (E-South). This state generates the observation shown in the center of Figure 1. Given this observation, the model computes all of the states that it could be in given that observation (Step 1 above). The model's current belief is shown graphically on the right side of Figure 1. More specifically, the model computes that there is a .5 probability of being in state E_S (at location E facing south) and a probability of .5 of being in state D_N (position D facing north).

Computing the optimal action. When computing actions, the model only computes the next best action, assuming that all of the actions following this action will be optimal. The current model can execute three different actions: rotate-left 90°, rotate-right 90°, and move forward to the next gray square in Figure 1. Table 1 illustrates how the model would compute the optimal action given the two states that it could be in (E-South and D-North). For each action, the model computes the number of actions assuming that it is at each state. For each of these assumptions, it also computes the set of actions required to reach the goal when this assumption is made and is wrong. The model computes these routes by computing a breadth-first search through the state-transition matrix until it finds a route that will reach the goal starting with the specified action. A more elaborate description of how these values were calculated is available in the Appendix.

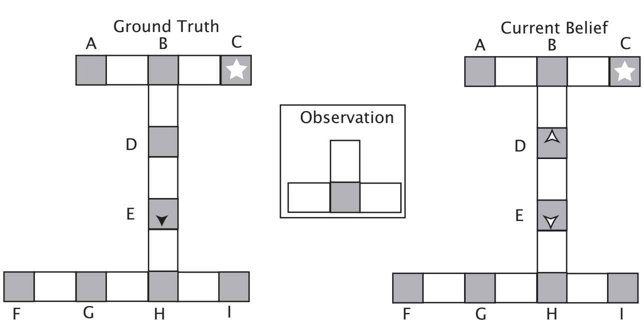

Figure 1. An illustration of a simple environment. The star at Position C is the goal. The illustration on the left indicates the ground truth of the current state of the environment. The center graphic shows the current observation made by the observer, and on the right is the observer's belief given the initial observation, $p_{ES} = p_{DN} = 0.5$.

Table 1 shows the most efficient routes assuming that the model starts in each of the possible states given the previous observations and actions (Step 2 above). Using these routes, the model computes the expected cost for starting with each action (see the far-right column of Table 1; Step 3 above). The action with the lowest expected cost is "Forward" (5.5 actions on average; Step 4 above). After computing that the optimal action is "Forward," the model then generates that action (Step 6), makes a new observation (Step 7), checks to see if it has reached the goal state with no remaining uncertainty (Step 8), and then updates its belief about where it is in the environment (Step 1). The model then continues through this cycle until it reaches the goal state with no remaining uncertainty.

After computing that the optimal action is "Forward," the model then generates that action, gets a new observation, and updates its belief about where it is based on the new observation and the action that was generated. The model then repeats the entire process again until it ultimately reaches the goal with no remaining uncertainty.

Formal Description

The model has perfect knowledge of the environment. This knowledge comes in the form of knowing the observation vector for each state in the environment (see Equation 2). Furthermore, the model can invert the function such that it can produce the state or set of states that are consistent with a particular view. It should be noted that for any given state in the environment there is a single observation vector,[4] but for a particular observation vector (o_i), there may be more than one state that could have produced that vector.

Given a single observation, there may be some amount of state uncertainty (i.e., more than one location may generate the same observation). We can represent this uncertainty by specifying a belief vector (**B**) that indicates the probability that the observer assigns to being in each of the states of the environment (**S**). We can write $b(s)$ as the probability assigned to state s when the observer's belief state is b.

In addition to the belief state vector and the observation vector, there is also an action vector A. The action vector is the set of movements or actions the observer can make within an environment.[5] Given a belief vector, observation vector, and the action vector, the model requires a function for estimating the observer's state within the environment. This state estimator (*SE* in Equation 3) takes as input the current belief vector, the previous action, and the current observation and returns an updated belief state. To

[4] There is no random noise in the perceptual or observation vector in the present version of the model. The uncertainty occurs because multiple states may generate the same observation vector.

[5] In the present studies the ideal observer and the human observer have three actions available to them: rotate-right 90°, rotate-left 90°, and move forward one hallway unit.

Table 1
An Example of How the Ideal Navigator Computes the Optimal Action Given the Scenario Shown in Figure 1

	D-North (assumed)		E-South (assumed)		
	D-North (true state)	E-South (true state)	E-South (true state)	D-North (true state)	Expected cost
Left	*L*RF(R*)F = 5	*L*RF(R*)RFFFRF = 10	*L*LFF(R*)F = 6	*L*LFF(R*)RFFFRF = 11	.5 × (.5 × 5 + .5 × 10) + .5 × (.5 × 6 + .5 × 11) = 8
Right	*R*LF(R*)F = 5	*R*LF(R*)RFFFRF = 10	*R*RFF(R*)F = 6	*R*RFF(R*)RFFFRF = 11	.5 × (.5 × 5 + .5 × 10) + .5 × (.5 × 6 + .5 × 11) = 8
Forward	*F*(R*)F = 3	*F*(R*) RFFFRF = 8	*F*(R*)RFFFRF = 8	*F*(R*)F = 3	.5 × (.5 × 3 + .5 × 8) + .5 × (.5 × 8 + .5 × 3) = 5.5

Note. The observer can be in two states: D-North and E-South. The column titled D-North provides a description of the estimated cost (in actions) assuming that the observer is at D-North. Below that is the cost when the observer is at D-North versus when the observer's assumption is wrong (E-South). In each of these columns is the set of actions (L = rotate-left; R = rotate-right; F = forward). The asterisk indicates when the observer would have gotten a different observation depending on whether the observer is in the assumed state or the other state. The actions up to the asterisk are the same. In this case, given the situation in Figure 1, the observer would choose "move forward" because this action minimizes the expected cost.

update the current belief state, one simply needs to apply Bayes' rule for estimating the likelihood that the observer is in a given state s':

$$SE_{s'} = p(s'|a,o,b)$$

$$SE_{s'} = \frac{p(o|s',a,b)p(s'|a,b)}{p(o|a,b)}. \tag{3}$$

Optimal Move Decisions

Route Generation

The goal of the system is to reach the goal state with no uncertainty, $b(s_{\mathrm{Target}}) = 1.0$, using the fewest number of actions (i.e., the lowest cost). To compute these routes (or cost), the model does a breadth-first search through the *state-transition matrix*. The state-transition matrix makes explicit the resulting state (s') that the observer would be in if the model in one state (s) and generated a specific action (a; see Equation 3). The model searches through the state-transition matrix using a breadth-first search algorithm until it finds a route in which the observer would reach the goal state with no remaining uncertainty.

Action Cost

In the model, each move has an associated cost (C_a). For each state (s) we can compute a set of routes (R) that will move the observer from a state (s) to the goal state (s_{Target}) with no remaining uncertainty. For each state there are multiple routes to the target state (s_{Target}). For each of these routes we can compute a cost associated for that route, $Cost(r,s,b)$, as specified in Equation 4:

$$Cost(r,s,b) = b(s)\sum_{a\epsilon r} C_a. \tag{4}$$

Note that r is the vector of actions that can move the observer from state s to the target state. Note also that the sum of the costs is multiplied by the likelihood that the observer is located at a particular state in the environment. That is, the observer weights the expected cost for reaching the goal state by the likelihood that they are in that particular state, b(s). In this way the observer optimally integrates the costs of reaching the goal based on the certainty that it is in each state.

Because the model may be in a state of uncertainty, the model only plans an optimal move one move ahead. The reason for this is that after the move, the observer will collect new perceptual information that will modify the belief vector (see Equation 2). For each state s the model computes the least expensive route that starts with each action ($r_1 = a$):

$$StateActionCost(s,a\epsilon A) = \min[Cost(r_1 = a,s,b)]. \tag{5}$$

After computing the cost of making each action from each state, the model then selects the action that minimizes the cost across all of the states in the environment (see Equation 6). If there are two moves that have the same minimum cost, the model randomly chooses one of the moves:

$$Move = \min_{a\epsilon A}\left[\sum_{s\epsilon S} StateActionCost(s,a)\right]. \tag{6}$$

After choosing the action selected by Equation 6 and making the action, the model receives a new observation vector (**o**). The model then updates its belief vector using the new observation vector (**o**) and the selected action (a) using Equation 3.

Experiment 1: Effect of Layout Size on Spatial Navigation Performance

In Experiment 1 we were interested in understanding the effect of increasing layout size on human spatial navigation performance. The purpose of this study was twofold. First, we wanted to determine if participants were inefficient at navigating with state uncertainty, and second, we wanted to investigate how these inefficiencies change as a function of the layout size. Increasing the layout size should increase the demands placed on an observer's perceptual processing (i.e., gathering perceptual data from short hallways vs. longer hallways), accessing the cognitive map (the larger environments have more hallway structures to remember than smaller environments), belief vector updating (given an observation, there are more places to consider in larger environments than smaller), and decision strategy (because the participant will have to travel longer to reach the goal, the participant will have to make more decisions in larger environments than smaller environments). We hypothesized that if participants have a cognitive limitation in one of these functions, it should become evident as a decline in efficiency as the environment becomes larger. To investigate this issue, we used four different layout sizes defined by the number of hallway segments in a layout. The four levels of layout size were 10, 20, 40, and 80 hallway units (see Figures 2 and 3).[6]

In this experiment, we trained participants to navigate through these virtual environments until they attained a specified learning criterion. Once proficiency was established, we then started participants from a random state within the environment and instructed them to move to a target location and resolve any remaining uncertainty they might have, using the fewest possible actions. Because the environments were perceptually sparse (no object landmarks), the participant could start a test session with considerable state uncertainty. We also ran the ideal navigator through the equivalent environments with the same set of actions and the same goal.

Method

Apparatus

The experiment was run on a Dell computer with a 19-in. (~48-cm) color monitor. The participant moved through the environment by making keypresses that corresponded to a 90° clockwise rotation, a 90° counterclockwise rotation, or a forward translation of one hallway unit.

After the participant made a keypress, the computer would rotate or translate the virtual "camera" in the virtual space. The camera would produce the appropriate optic flow for the action indicated by the keypress. A rotation was completed in 750 ms, and a translation was completed in 900 ms.

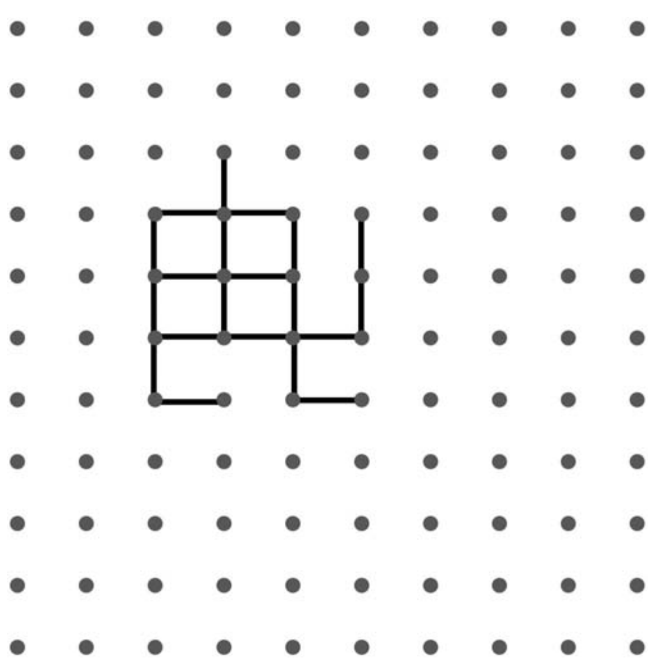

Figure 2. An example of the types of maps generated by the random layout generator. This environment is a 20-hallway environment. Each line segment connecting adjacent grid points in the figure represents a hallway. The end of the lower-left L-junction served as the starting point for each of the exploration sessions.

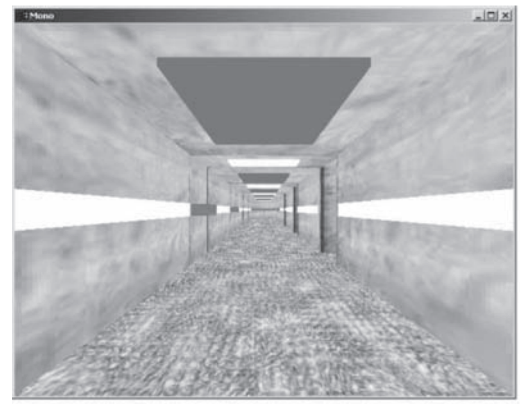

Figure 3. An example of a first-person rendering of an environment used in Experiment 1: a view from the perspective of the observer in the environment. Each hallway segment had the same structure in terms of the textures on the walls, floor, and ceiling. The objects (railings and ceiling lights) in each hallway were also the same for each hallway.

Stimuli

For each layout size (10, 20, 40, and 80 corridors), two different environments were produced for a total of eight different environments. Figure 2 provides an illustration of an environment of size 20. Each layout had one common feature. They all had an L-junction on the exterior part of the layout (see lower left-hand portion of the hallway structure in Figure 2). The end of this L-junction served as the starting point for each trial of the exploration phase of the training session.

The environments were randomly generated. We generated the environments by specifying three parameters: the number of hallway units comprising the environment, the maximum number of vertical hallway units, and the maximum number of horizontal hallway units that the environment could have (plat size). The plat size for the present studies was 20 hallways by 20 hallways.

To generate an environment, the computer began by randomly selecting one of the 400 potential hallways in the plat. After selecting a hallway, the computer then identified the set of potential hallways. Potential hallways were all of the hallways that connected to the current set of selected hallways. The computer then selected one of the potential hallways and removed the hallway from the list of potential hallways. It then recomputed the new set of potential hallways. The computer continued with this process until it reached the layout size minus 2 hallways. At the end, two hallways were added to the environment (an L-junction) that served as the starting point during the exploration phase of the experiment.

The environments were rendered from a first-person perspective with the eye height of the camera (in the virtual environment) placed at 5 ft. Figure 3 provides a sample view of the environment from the participant's perspective. To increase the participant's ability to differentiate between an intersecting hallway and a wall, we placed red railings at junctions on walls where there was no intersecting hallway (see Figure 3, left).

[6] A single hallway unit is a 30-ft-long hallway that was used to generate the random environments used in the present study. These hallway units were randomly placed on a grid to generate environments similar to the map shown in Figure 2.

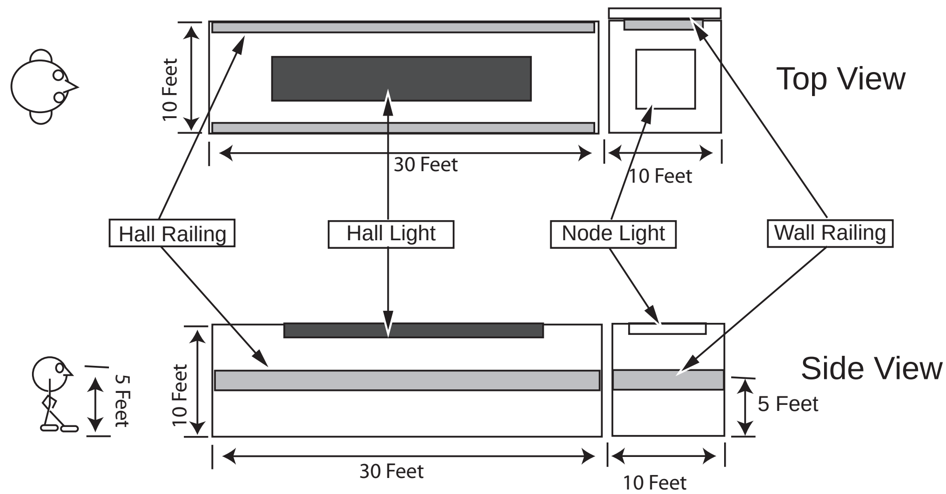

Figure 4. To generate the environments, we used a "parts kit" that consisted of a hallway, a node, and a wall. All of the structures were generated using these three fundamental parts. Some of the measurements are given in the graphic here; more of the details are listed in Table 2. The upper graphic shows the measurements when looking at each part from above, and the lower graphic illustrates the layout when looking at the environment from the side.

To ensure that all of the hallways were identical, we used an "environment parts kit" to generate the environments. This parts kit consisted of three basic parts: a hallway, an intersection node, and a wall. To generate a virtual environment, we placed these parts in specific configurations to give the appropriate environment layout. Figure 4 illustrates the parameters and properties of the parts kit, and Table 2 lists details about the parameters of each part, including the rendering color or texture.

Procedure

Participants took part in a training session and in a test session for each condition of the experiment. The training session was designed to provide the participant with a detailed representation of the environment. The training session consisted of two phases: an exploratory phase and a drawing phase. In the exploratory phase, the participant started from an external L-junction. The participant explored the environment for 3 min by making keypresses on the number pad to initiate the desired movement. The "8" corresponded to a forward movement, and the "4" and "6" corresponded to rotate counterclockwise and clockwise rotations, respectively. During the training phase, participants learned both the layout of the environment and a target location. The target location was specified by an auditory signal (the sound of a bell) each time the participant walked over the target location. Participants were told that later in the experiment they would start from a random place in the environment and would need to move to the target location making as few actions as possible.

After exploring the environment for 3 min, the participant took part in the drawing phase. In the drawing phase, the participant was given a grid pattern that had a single L-junction placed near the center of the grid that corresponded to the starting location in the exploration phase of the study. The participant was told to "connect the dots" to recreate the environment just explored. Participants were informed that each dot could be thought of as a node or the stopping location when they made a forward movement through the environment.[7] If the participant's map drawing did not perfectly reproduce the grid layout of the environment, then the participant took part in another 3-min exploration session followed by another drawing test. Participants continued in the 3-min exploration phase followed by the drawing phase until they drew the environment correctly twice in a row.

After reaching criterion in the training phase, the participants entered the test phase of the experiment. In the test phase, participants started from a random state (i.e., a random location and orientation) in the environment. Participants were instructed to move to the target location using as few actions (keypresses) as possible. They were informed that a rotation and a translation were both considered an action. When participants reached the target location, there was no auditory signal indicating they were there. Instead, participants were required to indicate when they believed they had reached the target location by pressing the spacebar on the computer. After the spacebar was pressed, the screen went white; when the participant was ready to begin the next trial, the participant pressed the spacebar a second time to reveal the new starting view.

Participants started from each possible position in the environment an equal number of times (e.g., each of the 17 nodes in Figure 2). At each starting position, the computer randomly selected one of the four possible starting orientations. Notice that this means that a participant sometimes started facing a wall. Each participant participated in 320 trials in each of

Table 2
Properties of the Environment Parts Kit Used to Generate the Random Environments for Experiments 1 and 2

Item	Height (ft.)	Width (ft.)	Length (ft.)	Color/texture
Hall floor		10	30	Burlap
Hall ceiling		10	30	Cement
Hall wall	10		30	Cement
Hall railing	1	0.5	30	White
Hall light	0.5	8	28	Purple
Wall	10		10	Cement
Wall railing	1	0.5	4	Red
Node floor		10	10	Burlap
Node ceiling		10	10	Cement
Node light	0.5	6	6	White

[7] To be certain that participants understood this process, we first trained them on a small (10-hallway) layout environment before starting the experiment.

the environments. Participants were tested in each of the eight environments in a random order.

Participants

Three female participants were tested in the experiment. Each participant had normal or corrected-to-normal vision. The ages of the participants ranged from 21 to 24 years, and all were students at the University of Minnesota. The participants were paid $8/hr to participate in the study.

Results

The primary dependent measure in Experiment 1 was the number of moves required to reach the target location. For participants to solve this task, they had to determine both their location within the environment (because they were starting from a random location) and the shortest path (fewest actions) to the target location. The upper graph of Figure 5 illustrates the average number of moves participants took as a function of layout size. Trials in which participants indicated the target location at the wrong position or trials in which it would have been impossible for the humans to have reached the goal state with no remaining uncertainty were excluded from the analysis (combined, these trials were less than 1% of the trials). Figure 5 also illustrates the number of moves, averaged across layouts of a given size for the ideal navigator.[8] The number of actions for the ideal observer increases initially (from 10 to 20 hallways) with the effect of layout size having a diminishing effect between 20 and 80 hallways.

The lower plot of Figure 5 illustrates an efficiency measure that allows us to factor out task difficulty in human behavior. Efficiency is computed as the number of actions required by the ideal observer divided by the number of actions made by the human. The lower plot of Figure 5 illustrates the efficiency functions for the 3 participants. A one-way analysis of variance (ANOVA) revealed a significant effect of layout size on navigation efficiency for each of the participants: Sub1, $F(3, 908) = 5.422$, $p < .01$; Sub2, $F(3, 908) = 32.11$, $p < .01$; and Sub3, $F(3, 908) = 22.03$, $p < .01$. Human efficiency dropped from approximately 80% for the smallest environment (10 hallway units) to approximately 50% in the largest environment (80 hallway unit environments).

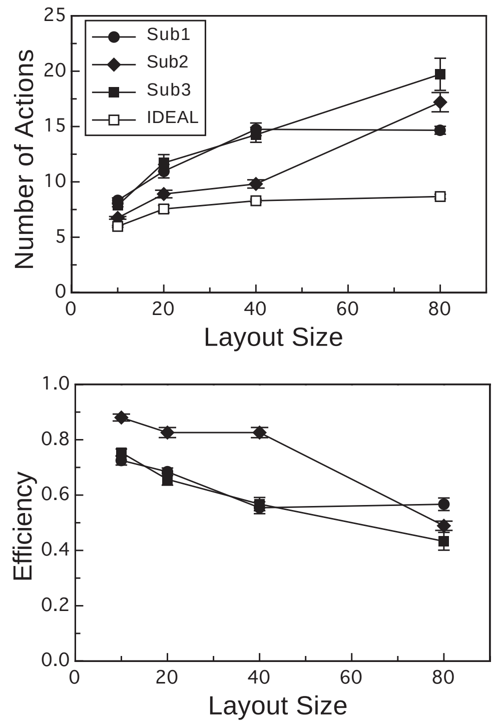

Figure 5. The effect of layout size on human navigation performance. The upper graph shows the average number of moves for the 3 participants (Sub1, Sub2, and Sub3) and the ideal navigator. The lower graph shows navigation efficiency of the participants, computed as the average number of actions of the ideal navigator divided by the number of actions made by the participant. Error bars represent 1 standard error of the mean.

Discussion

The results from Experiment 1 show that participants' navigation performance is suboptimal. This is demonstrated by the fact that the efficiency is below 1.0 across each of the layout sizes. The results also show that participants become less efficient as the size of the environment becomes larger. The sources of inefficiency might have several causes, which can be broken down into four primary categories: perception, accessing the cognitive map, spatial updating, and decision strategy.

Source of the Cognitive Limitation

Using the ideal navigator, we can begin to investigate where the cognitive limitations might exist in spatial navigation. In the following four subsections, we outline different causes of this suboptimal behavior. In addition to this, we also consider Schölkopf and Mallot's (1995) view-graph model of spatial navigation to determine if this model can explain the suboptimal behavior.

Limited perceptual processing. One function that might be limiting human navigation performance is the visual information used by the human observer. Although the environments used in Experiment 1 were very sparse, the observers may not have processed all of the hallway information while navigating. For example, participants may simply have identified the structural information up to the next hallway unit, and then ignored the structural information beyond that intersection. Of course this perceptual limitation could be extended so that participants might only process *N* closest hallways and ignore the perceptual information beyond that. A perceptual limitation like this would affect navigation in larger environments more than in smaller environments because the average length of a corridor in the larger environments is longer than in the smaller environments.

Accessing the cognitive map. Participants may have difficulty accessing the entire cognitive map while they are navigating. Although participants are required to draw the environment correctly twice in a row before they can participate in the testing phase, they might not be able to access the entire cognitive map with sufficient precision to make the relevant comparisons implicit

[8] The ideal-observer performance was computed for each participant in each environment. That is, the model started from the same states as the human observers. Because there was no significant variation across the participants, we plotted the average performance across the trials, participants, and environments.

in ideal performance. Facility in accessing the global cognitive map may also be adversely affected by the concurrent task of navigating through the environment. Therefore, it is plausible that working memory capacity is preventing participants from loading the entire cognitive map while they are navigating. If participants have a limited working memory capacity that only allows them to load a limited amount of the layout information, then as the environment becomes larger, the proportion of the entire layout that is being considered drops, with a corresponding reduction in navigation efficiency.

Spatial updating. A third cognitive limitation may arise from the participants' ability to update their position within the environment. In the present experiment the participant starts from an unspecified location within the environment. When participants make their first observation, there may be many locations consistent with this observation; in other words, there may be substantial ambiguity. We quantify this ambiguity with the *belief vector*. As the participants make actions within the environment and obtain new observations, the belief vector should change, and the uncertainty should decrease until ultimately the participants know exactly where they are located. One possible source of suboptimal human performance might be due to the participant's inability to accurately update this belief vector while moving through the environment. Participants may be limited in the number of locations they can consider (e.g., 7 ± 2 places; Miller, 1956), or they may not possess an optimal updating strategy. Because the larger environments consist of more locations, the number of states in which the observer could be while navigating also increases. Thus, if participants have a suboptimal spatial updating strategy or limited belief vector memory, their performance will be negatively affected as the environment becomes larger.

Suboptimal decision strategy. Finally, the strategy used by the participants may be suboptimal. For example, one suboptimal strategy is an inappropriate cost function placed on each action. The participants in the experiment were explicitly instructed to make as few actions as possible to reach the goal state. This would place an equal cost on translations and rotations. It is possible that, despite these instructions, participants did not treat each action as having an equivalent cost. For instance, they may have associated a larger cost with a translation than a rotation. Another example of a suboptimal decision strategy would be if participants first set the goal of unambiguously localizing themselves within the environment, and only then focused on moving to the target location with the minimum number of actions.

View-Graph Strategy

One model that does not use an ideal decision-making strategy is the *view-graph model* of spatial navigation proposed by Schölkopf and Mallot (1995). In this model, spatial navigation is accomplished by storing a series of views in memory and, for each view, assigning a specific action for reaching a particular destination. This model has been used in robotics for learning and navigating through complex environments. Research by Gillner and Mallot (1998) found evidence that human participants may use this type of strategy while navigating through complex environments. It is possible that participants were engaging in this type of suboptimal decision strategy in Experiment 1. To test whether participants are engaging in this strategy, we investigated how well one can predict the human action given a particular view. If participants are using this strategy for navigating through these environments, we would expect to be able to predict their selected action for each view. To investigate this question, we measured the conditional entropy between the action selected (A) and each view in the environment (see Equation 7). High conditional entropy indicates that the observer's responses were highly variable for each view, whereas low conditional entropy indicates that for a given view the observer typically made the same response. According to the view-graph model of spatial navigation, humans should choose the same action (A) given a specific view. That is, the view-action entropy should be relatively low (or possibly 0):

$$H(A|V) = \sum_{v \in V} p(v) \left[\sum_{a \in A} p(a|v) \log \frac{1}{p(a|v)} \right]. \tag{7}$$

As a comparison, we also computed the conditional entropy for the ideal navigator. The navigator is not in any way constrained by a view-action association. The model instead selects the optimal action given its state of uncertainty and knowledge of the environment. The model serves as a baseline by which to compare the human data. It provides a lower bound on the action entropy if a system is making optimal use of the perceptual information and is not limited in its strategy.

Figure 6 shows that the ideal navigator's conditional entropy is lower than that for each of the participants. That is, it is more difficult to predict the human action (given the view) than the ideal navigator's action. This finding suggests that if participants are limited by their decision strategy in Experiment 1, they do not seem to be limited in the way suggested by the view-graph model.

Interim Summary

Experiment 1 investigated the effect of increasing layout size on human navigation performance. We measured the number of actions a participant used to move from an unspecified location in a familiar virtual reality environment to a target location. Four layout sizes were used: 10, 20, 40, and 80 hallways. The raw data showed that participants required more moves to reach the goal as

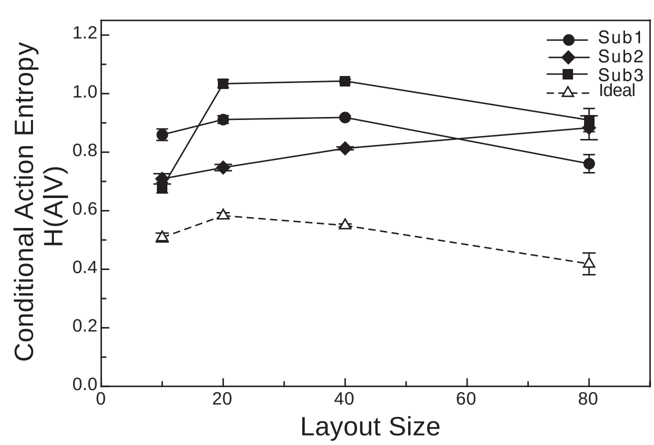

Figure 6. Illustration of the conditional action entropy—$H(Action|View)$, or $H(A|V)$—for the 3 participants (Sub1, Sub2, and Sub3) and for the ideal navigator. Error bars represent 1 standard error of the mean.

the layout size increased. These findings are similar to those reported by O'Neill (1991) in which participants backtracked more often and took more wrong turns as the environment complexity increased. However, in the present study and those reported by O'Neill, the increase in wrong turns or more actions could be due to cognitive limitations or task demands. To determine whether a cognitive limitation contributed to the increase in the number of actions in the present study, we developed an ideal navigator with no perceptual, memory, or strategy limitations (cognitive limitations). The model's performance was compared with human performance. We found that the increase in the number of human actions outpaced that of the ideal navigator, suggesting there is a cognitive limitation that is affecting human spatial navigation.

We hypothesized that the inefficiency found in this experiment might be due to inefficiencies in perceptual processing, accessing the cognitive map, spatial updating, or decision strategy. In Experiment 2 we investigated whether the inefficiency can be explained by an inefficient perceptual processing system. In Experiment 3 we investigated whether the inefficiency is in accessing the cognitive map, spatial updating, or decision strategy.

Experiment 2: The Effect of Limited Visual Information on Spatial Navigation Performance

Experiment 2 examined whether human navigation performance is limited by visual information-processing constraints. For this experiment, the participants performed the same navigation task as in Experiment 1. However, rather than manipulating the size of the layout, we manipulated how far the observer could see along each corridor (the view depth). This was accomplished by adding "virtual fog" to the environment (see Figure 7). The ideal-navigator model was confronted with the same limitation, by restricting how much of the observation was available from each state. As in Experiment 1, we compared the navigation performance of human observers with the performance of the ideal navigator to obtain navigation efficiency.

If human navigation inefficiency in Experiment 1 is due to a failure to use information beyond some view depth, then restricting the view depth for both humans and the ideal navigator should bring performance into closer agreement. Accordingly, we expect to see an improvement in efficiency when the view depth is limited.

Method

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli

View depth was manipulated by adding virtual fog to the environments to limit how far down the hallway an observer could see. Figure 7 illustrates this effect by showing a view without fog (unlimited-view condition) and a view with fog (fog depth = 1 hallway). The fog provided a method for manipulating viewing depth that felt relatively natural.

Figure 8 provides an analysis of the view information available for the four environments used in Experiment 2. One concern in running this study was whether the fog manipulation in fact imposed a substantial perceptual limitation. Suppose the fog restricted visibility to a distance of one hallway unit. The effect would be substantial if there were many states in the layout in which an unrestricted view extends more than one hallway unit. The question is whether the environments we used included a substantial proportion of states with these longer views. Figure 8 plots the frequency of states as a function of the number of collinear hallways for each environment (i.e., hallway length). In the present study, the fog (limited view) condition would limit the perceptual information available in the states in which the hallway length is greater than one. Across the four environments, this constitutes 36% of the states.

Procedure

Four different environments were used in Experiment 2. Each environment consisted of 40 hallway units. Participants were tested in all four environments: two with limited view depth and two with unlimited view depth. Each environment was used in both viewing conditions across the participants.

The training procedure was the same as in Experiment 1. When participants trained for the limited-view condition, they completed the training session with limited visual input.

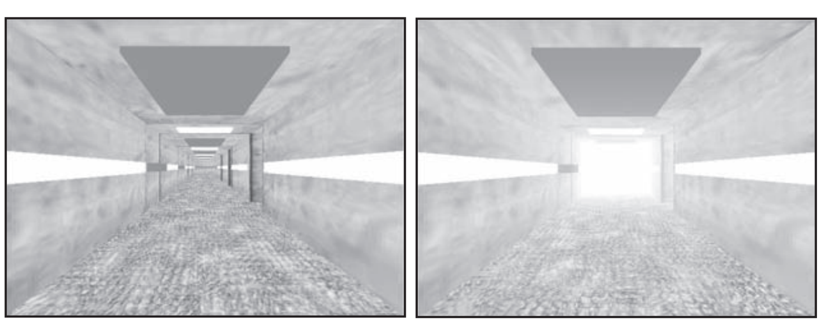

Figure 7. Sample views from Experiment 2. The image on the left illustrates the view in the unlimited-view condition, and the image on the right illustrates the same perspective in the limited-view condition. Note that in the image on the left the participant can determine the hallway structure for the next node, but beyond that node the participant has no perceptual information.

Participants

Eight participants were tested in Experiment 2: 4 women and 4 men. Their ages ranged from 20 to 22 years. All participants had normal or corrected-to-normal vision.[9]

Results

Figure 9 shows the average efficiency for the 8 participants in each condition. The primary question of interest is whether humans became more efficient when view depth decreased (*limited-view* vs. *unlimited-view* conditions). Figure 9 shows that there was no improvement in the limited-view condition over the unlimited-view condition, and in fact, the trend is in the opposite direction (the unlimited-view condition is better than the limited-view condition). A two-tailed paired t test showed that there was no significant difference between these two conditions, $t(7) = 2.03$, $p = .082$.

Discussion

This experiment investigated the effect of limiting visual input on spatial navigation performance. The purpose of this experiment was to determine whether the cognitive limitation found in Experiment 1 was due to inefficient visual processing. For example, in Experiment 1, participants might have processed information only one or two hallway units down a corridor view. If participants' performances were limited by visual processing, then the participants should be less efficient in environments that have more long corridors[10] than those with fewer long corridors. The smaller environments in Experiment 1 have fewer long corridors than the larger environments, and thus limited visual processing might account for the data.

We predicted that if participants were limited in the amount of visual information that they process, then we should find an increase in efficiency when we redefine the task to have a reduced visual demand. In other words, the fog manipulation reduced the useful information for both the human participants and the ideal-navigator model, potentially putting the human and model on a more even footing. If this were the case, we would have found higher efficiency in the fog condition. Because we did not find higher efficiency, we conclude that human inefficiency in Experiment 1 is not due to a failure to encode and use visual information beyond a restricted view depth.

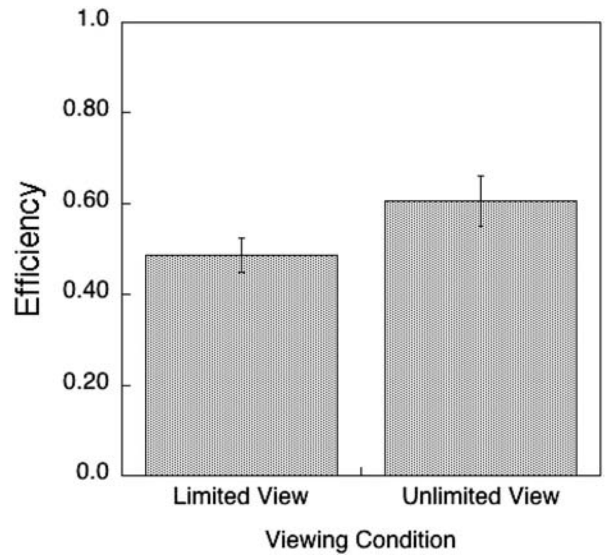

Figure 9. The mean navigation efficiency when navigating in the unlimited and limited viewing condition in Experiment 2. In the limited-view condition, visual information was available as far as the next intersection (further details were obscured by "fog"). In the unlimited-view condition, visual information was available to the end of the corridor. Error bars represent 1 standard error of the mean.

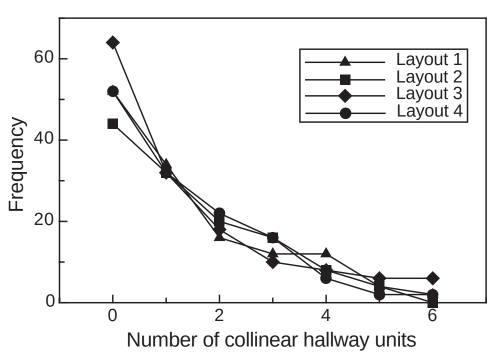

Figure 8. View-depth analysis of the four environments used in Experiment 2. The graph shows the number of states (Frequency) that have a specific number of collinear hallway units that can be observed from that state. For example, when the participant is looking at a wall, the number of collinear hallway units is 0. Observers in the unlimited-view condition would be able to see all of the hallway units. Observers in the limited-view condition would be able to see all of the information in the 0 and 1 view-depth states, but in the states with a view depth greater than 1 they would lose access to the more distant information (beyond a view depth of 1).

Experiment 3

Experiment 3 investigated whether the inefficiency found in Experiment 1 was due to inefficiencies in (a) accessing the cognitive map, (b) spatial updating, or (c) decision strategy. In Experiment 3 we did not manipulate the layout size (we used only one layout size: 40 hallways) but instead provided participants with supplementary map information while they navigated to the target location from an unspecified starting location. The experiment had three conditions differing in the type of supplementary information made available to the participant: *no-map, map,* and *map + belief vector* conditions. In each of these conditions, the information that is made explicit (map or belief vector) is implicitly available to the human observer and to the ideal-navigator model during the experiment. That is, we are not changing the task at all (from an information perspective), we are simply making

[9] We ran more participants in Experiment 2 because it took much less time to collect a complete data set in this study than it did in Experiment 1 (approximately 18 hr/participant in Experiment 1 and only 8 hr/participant in Experiment 2).

[10] A corridor is a series of collinear hallways.

some information explicit that might be difficult for participants to generate independently.

The no-map condition allowed us to measure baseline efficiency performance for the task and was almost identical to the conditions in Experiment 1. The only difference was that the participants simply saw a blank square on the lower-left corner of the screen while they were navigating through the environment (see Figure 10). In the map condition, we provided a map of the environment on the computer monitor while the participant was navigating. Because Experiment 3 used the same map-drawing technique as in Experiments 1 and 2, we know that participants' cognitive representations are sufficient to generate a global map of the environment. But it is possible that generating the map is computationally burdensome, and that during navigation it is not easy to maintain access to a global cognitive representation of this type. With the use of a map display, the participants were not required to access a cognitive map during navigation. Instead they could reference the visual map presented on the display. Thus, we hypothesized that if human inefficiency is due to problems accessing the cognitive map, then we should find a significant improvement in performance from the no-map condition to the map condition.

Alternatively, participants might be inefficient in their spatial updating ability. Remember that in this experiment participants are starting from an unspecified location in which the initial view can leave the observer with state uncertainty. As participants move through the environment, they have to use the perceptual information and the action selections to update where they believe they are in the environment. Or in terms of the model, they have to update their *belief vector.* In the map + belief vector condition, we presented the observers with the map and the target location. Superimposed over the map was an accurate belief vector, assuming that all of the perceptual information was being used (see lower-left and lower-right panels of Figure 10). This belief vector was updated after every action made by the observer. The computer updated the current belief vector by factoring in the current and previous observations and the previous actions. We hypothesized that if the navigation inefficiencies found in Experiment 1 were due to spatial updating, adding the belief vector information should show an increase in movement efficiencies from the map condition to the map + belief vector condition.

Finally, participants might simply have an inefficient movement strategy. It is possible that participants might be accessing the cognitive map perfectly and also generating an accurate belief vector, but their action selection strategy is inefficient. For example, participants may decide to choose the action that minimizes the overall distance between their current position and the goal state. Although under some conditions this might be a good strategy, in others it might lead participants down dead-end hallways,

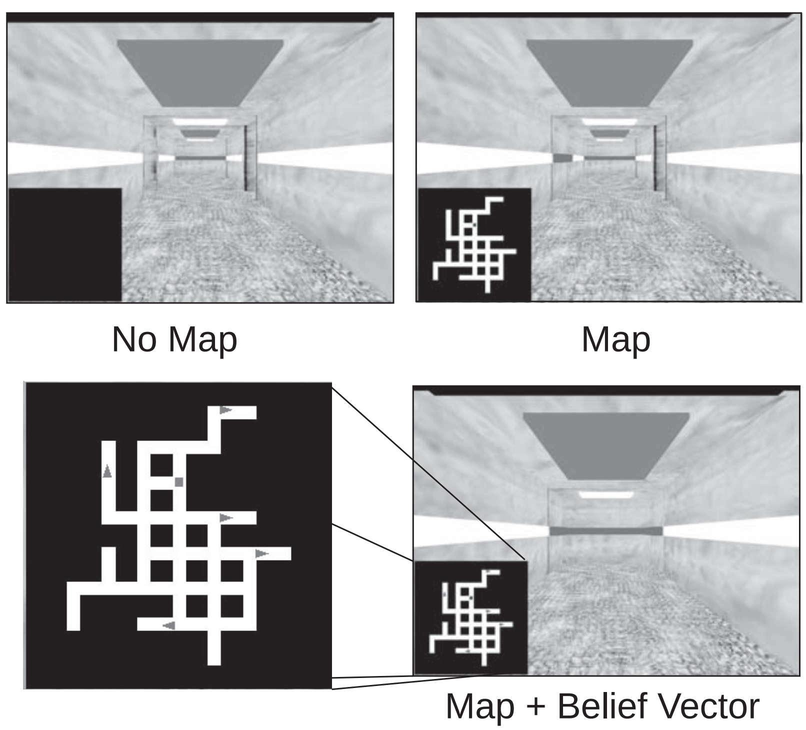

Figure 10. Illustration of the three conditions as viewed by the observer. The upper-left illustration shows a view when there is no map information provided. The upper-right panel shows a view when the map is provided. In addition to the map, the location of the target position was also shown by a small square at that position. The lower-right panel shows a view when the observer sees the map with an accurate belief vector superimposed on top of it. The lower-left panel shows an enlarged version of the belief vector map. The arrows on the map show where the participant could be given the previous views and actions along with the current view. Note that not all of the dead ends are shown as possible locations. This is because this image was generated after a sequence of actions.

which would be inefficient. If this is true, then adding the supplemental information will not increase the participant's efficiencies at all. Thus, if the inefficiencies are in the action selection strategy, then we predict there will be no difference in performance across the three conditions.

Method

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli

Experiment 3 manipulated the supplemental map information available to the participant. There were three types of supplemental information: no map, map, and map + belief vector. Figure 10 provides a sample illustration of how these conditions appeared to the participant.

In each condition there was a black square that was presented in the lower-left quadrant of the display. The location of this display was carefully chosen so that it would not block any of the informative information during the experiment. In the no-map condition this square was blank; in the map condition a static map image was superimposed over the black square. In addition to the map, there was a blue square indicating where the target location was in the environment. In the map + belief vector there was the map plus the target location symbol superimposed on the black square. In addition to this, there was a collection of red pointers indicating where the participant could be located in the environment given the participant's current view, previous views, and actions. These pointers were updated after each of the participant's actions.[11]

Procedure

Experiment 3 used one environment that consisted of 40 hallway units. The training procedure was the same as in Experiment 1. No supplementary map information was present during the training phase. After reaching criterion (i.e., drawing the environment correctly twice in a row), participants started the testing phase.

During the testing phase, participants started from each state in the environment three times, once for each of the three conditions. The order of state + supplemental information conditions was randomized for each of the participants.

Participants

Four participants were tested in Experiment 3: 2 women and 2 men ranging in age from 20 to 22 years. All participants had normal or corrected-to-normal vision. Participants 6, 7, and 9 were all volunteers who worked in the lab. Participant 8 was a paid participant who received $10/hr for his participation in the study.

Results

Experiment 3 manipulated the supplemental information available to the participant while navigating from an unspecified location to a target location. For each participant we computed their movement efficiency by taking the ratio of the number of moves required by the ideal observer to the number of moves required by the human observer.

Figure 11 illustrates these data for the three supplementary map conditions. Planned comparison *t* tests found that there was no significant difference between the no-map and the map conditions for all 4 participants (see Table 3), but there was a significant effect between the map and the map + belief vector condition for all 4 participants.

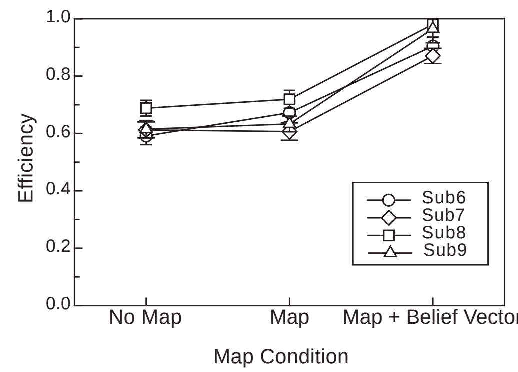

Figure 11. Action selection efficiency plotted for the 4 participants (Sub6, Sub7, Sub8, and Sub9) in Experiment 3 as a function of the supplementary information provided to the participant while navigating. Error bars represent 1 standard error of the mean.

Discussion

The findings from Experiment 3 are very clear. There was no significant increase in performance from the no-map to the map condition. This result suggests that participants did not have difficulty in accessing their cognitive map or that the global information afforded by the supplementary map was not useful to them. By contrast, there was a significant improvement in performance from the map to the map + belief vector condition (from approximately 60% efficiency to 95% efficiency). This large increase suggests that most of the inefficiency in Experiments 1 and 2 is due to the processes involved in updating their belief vector. The belief vector is a list of states that the observer (human or ideal) could be in given their previous observations and actions. After the participant makes an action and makes a new observation within the environment, he or she will update his or her belief vector (i.e., update the set of states the observer believes that they could be in). Our results indicate the participants have difficulty integrating the set of observations and actions with their cognitive map to generate an accurate list of states.

The results from the present study provide important information about sources of human inefficiency while navigating, but they do not tell us *why* humans are inefficient. Understanding why participants are inefficient at updating their belief vector will be addressed in future research. However, we speculate that there are four subprocesses involved in generating an accurate belief vector. These processes include the following:

1. Remembering the belief vector prior to the action you just made.

[11] In the present studies there was no noise in the actions, nor was there any uncertainty in the observation vector given to the model. Because of this, the probability that a participant could be in a particular state was either 0 or $1/N$ (where N is the number of potential states in the belief vector).

Table 3
Planned, Paired t-Tests for the 4 Participants in Experiment 3

	No map vs. map			Map vs. map + belief vector		
Participant	*df*	*t*	*p*	*df*	*t*	*p*
Sub6	126	−1.881	.062	127	−5.791	<.001
Sub7	127	0.142	.887	127	−6.764	<.001
Sub8	124	−0.769	.443	124	−6.796	<.001
Sub9	125	0.463	.644	127	−5.848	<.001

2. Computing the states that you could be in given your previous belief vector and your most recent action (e.g., if I turned left, and I were in these three locations, where could I be now?).
3. Considering all of the states that are consistent with your current view.
4. Eliminating the set of possible states given the observed view (or only considering the intersection of states from the second and third process).

General Discussion

Spatial navigation is composed of multiple processes that include perception, memory, spatial updating, and decision making. A breakdown in any one of these processes can have detrimental effects on human spatial navigation performance. The series of studies described here investigated navigation efficiencies as a function of the size of the environment, with reduced visual input and with supplemental layout information. The goal of these studies was to understand how efficient human spatial navigation was under conditions of state uncertainty (i.e., when one does not know one's exact position and orientation within a familiar environment). Each of these studies compared human performance with that of the ideal navigator, based on principles from POMDP (Cassandra et al., 1994; Chung, 1960; Kaelbling et al., 1998). The findings from the present study provide insight into the cognitive limitations in navigating through indoor environments. First, Experiment 1 shows that participants become less efficient in their ability to navigate through large-scale spaces as the environments become larger (in terms of hallway units). This suggests that there is some sort of cognitive limitation that is correlated with layout size.

One possible limitation might be in the analysis of the visual information that is available from each view. Experiment 2 was designed to address this issue. By limiting the visual information for the human and ideal observer, we were able to see whether human efficiency changed under these two conditions. There was no change in efficiency, which suggests that the results from Experiment 1 were not due to a limitation in processing perceptual information.

Experiment 3 was designed to address whether the limitation might be due to accessing the cognitive map from memory, considering multiple states simultaneously, or the participant's decision strategy. The results from this study clearly showed that a large part of the limitation is due to ineffective use of preceding views and actions in resolving where they could be in the environment.

Navigating With Uncertainty

The present studies suggest that participants have difficulty navigating when there is state uncertainty. The three experiments show that the major factor in participants' inefficient navigation behavior lies in the methods they use to update their belief vector. Within this updating procedure, there are four specific procedures; any one of them, or a combination of inefficient processing, can lead to the inefficient behaviors found in the present experiments. These subprocesses include the following: (a) generate all of the states that are consistent with the current observation, (b) remember the candidate states, (c) update this collection of states on the basis of the most recent action; and (d) eliminate the candidate states that are inconsistent with the current view. Any or all of these processes may have produced the inefficient behavior found in these studies.

Lack of Proprioceptive Information

The present studies have investigated way-finding behavior using desktop virtual reality. This technology allows one to generate environments of arbitrary size and to manipulate and control the visual information available to the observer. One may question whether these results will actually generalize to navigation under more realistic conditions. Specifically, would efficiency be higher if participants physically navigated through comparable real environments? Previous studies by Klatzky, Loomis, Beall, Chance, and Golledge (1998) showed that participants perform worse with path integration when information for turning is conveyed only by visual cues (optic flow) in the absence of vestibular or kinesthetic cues. Others have shown that when immersive environments are used, participants have a difficult time making accurate judgments of distance relative to when they are in a real environment (Thompson et al., 2004). These results suggest that the lack of proprioceptive information or the use of virtual environments may limit the generalization of the present results to natural way-finding tasks.

We have addressed the issue associated with the lack of proprioceptive information in a study to be reported in detail in a later publication. In brief, we have conducted an experiment in which we compare efficiencies in three different conditions: keypress condition (identical to the condition used in the current study), joystick condition, and immersive condition. The joystick condition allowed the participant to move through a desktop environment using a joystick. The joystick movement allowed continuous movements through the environment rather than the quantized movements provided by the keypresses. In the immersive condition we tracked the participant's movements in a virtual reality arena. In each of the conditions we measured the participant's efficiency for reaching the goal state. In this study, we found no significant differences between the three conditions. These results suggest that providing proprioceptive information does not help improve efficiency in this specific task.

The work by Thompson et al. (2004) showing that participants make inaccurate absolute distance judgments using immersive virtual reality could still pose a problem. However, the tasks that participants completed in our studies do not rely on accurate, absolute judgments. Recall that the layouts were constructed on a Cartesian grid with standard unit lengths. The only requirement for distance estimation would be to compute the integer number of hallway units in the forward-facing view. Moreover, the results of Experiment 2 imply that difficulty in gathering visual information at a distance is not a limitation on performance. For these reasons, we doubt that an underestimate of metric distance judgments in virtual environments is a cause of human inefficiency in our study.

Spatial Navigation and Low Vision

Although we tested participants who had normal vision, we may speculate on how these results might apply to low-vision navigation. People with low vision refers to people who have uncorrectable visual impairment that affects their everyday activities. The present studies investigate human navigation in environments in which there is very little visual information and, correspondingly, increased spatial ambiguity. These sparse environments may simulate some of the characteristics of low-vision navigation in a real environment. For instance, individuals with low vision may have a difficult time detecting or identifying specific landmarks within the environment and therefore, functionally, the environments are visually sparse with heightened spatial ambiguity. Our results indicate that participants with normal vision have trouble generating an accurate belief vector under such circumstances. If this is the case, one challenge faced by individuals with low vision might be related to their difficulty in generating an accurate belief vector within an environment after becoming disoriented. Normally sighted participants can choose from a rich array of landmarks to provide important information about their state within the environment for reorientation; however, individuals with low vision have reduced access to these perceptual landmarks. These results suggest that one of the challenges that someone with low vision might have is being able to consider multiple places simultaneously when they become disoriented within an environment.

Summary and Conclusions

Spatial navigation is a fundamental human cognitive process that is used hundreds of times each day by most people. From navigating from home to work to finding one's way to the mailroom or even navigating through one's office to the bookshelf to reference the most recent issue of the *Journal of Experimental Psychology: Human Perception and Performance*, one is using one's navigation skills. The present collection of studies investigates a person's navigating skills when there is uncertainty about his or her current state in the environment and there is very little visual information to reduce this ambiguity.

To investigate this issue, we developed an ideal-observer model of spatial navigation. The tasks for the human and for the model were comparable. This model provided us with optimal behavior for each environment and for each task. We used the ideal observer to compute how efficiently participants were able to accomplish each task. By varying the information available to the participant (i.e., Experiment 2 manipulated the visual information available) and making some information explicit rather than implicit (Experiment 3), we were able to narrow down the human inefficiencies in these tasks to properly updating the set of candidate states (or the participant's belief vector) given the prior observations and actions. We were able to eliminate inefficient perceptual processing, inefficient memory for the map, and an inefficient decision process.

In the *Discussion* section of Experiment 3, we provided four possible subprocesses that might lead to the inefficient belief vector updating process. These included accurately generating the candidate states given an observation, accurately remembering the belief vector, accurately updating a belief vector given an action, and accurately eliminating candidate states given the current observation. Future research will investigate the contributions that these four subprocesses have on a person's navigating behavior when navigating with uncertainty.

References

Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting optimally in partially observable stochastic domains. In *Proceedings of the 12th National Conference on Artificial Intelligence* (AAAI-94, Vol. 2, pp. 1023–1028). Seattle, WA: AAAI Press/MIT Press.

Chung, K. (1960). *Markov chains with stationary transition probabilities.* Berlin: Springer.

Cohen, R., & Schuepfer, T. (1980). The representation of landmarks and routes. *Child Development, 51,* 1065–1071.

Franz, M., Schölkopf, B., Mallot, H., & Bülthoff, H. H. (1998). Where did I take that snapshot. *Biological Cyberneticx, 79,* 191–202.

Geisler, W. S. (1989). Ideal-observer theory in psychophysics and physiology. *Physica Scripta, 39,* 153–160.

Gillner, S., & Mallot, H. (1998). Navigation and acquisition of spatial knowledge in a virtual maze. *Journal of Cognitive Neuroscience, 10,* 445–463.

Hecht, S., Shlaer, S., & Pirenne, M. H. (1942). Energy, quanta, and vision. *Journal of General Physiology, 25,* 819–840.

Hirtle, S. C., & Heidorn, P. B. (1993). The structure of cognitive maps: Representations and processes. In T. Gaerling & R. Golledge (Eds.). *Advances in psychology: Vol. 96. Behavior and environment: Psychological and geographical approaches* (pp. 170–192). Amsterdam: North-Holland.

Hirtle, S. C., & Hudson, J. (1991). Acquisition of spatial knowledge for routes. *Journal of Environmental Psychology, 11,* 335–345.

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

Hirtle, S. C., & Kallman, H. J. (1988). Memory for the locations of pictures: Evidence for hierarchical clustering. *American Journal of Psychology, 101,* 159–170.

Hirtle, S. C., & Mascolo, M. F. (1986). Effect of semantic clustering on the memory of spatial locations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 182–189.

Jacobs, W. J., Thomas, K. G. F., Laurance, H. E., & Nadel, L. (1998). Place learning in virtual space: II. Topographical relations as one dimension of stimulus control. *Learning and Motivation, 29,* 288–308.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence, 101,* 99–134.

Klatzky, R. L., Loomis, J. M., Beall, A. C., Chance, S. S., & Golledge, R. G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science, 9,* 293–298.

Klatzky, R. L., Loomis, J., Golledge, R., Cicinelli, J., Doherty, S., & Pellegrino, J. W. (1990). Acquisition of route and survey knowledge in the absence of vision. *Journal of Motor Behavior, 22,* 19–43.

Kuipers, B. (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.), *Proceedings of the Third International Symposium* (pp. 10.11–10.17). Ann Arbor: University of Michigan.

Legge, G. E., Hooven, T. A., Klitz, T. S., Mansfield, J. S., & Tjan, B. S. (2002). Mr. Chips 2002: New insights from an ideal-observer model of reading. *Vision Research, 42,* 2219–2234.

Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: An ideal-observer model of reading. *Psychological Review, 104,* 524–553.

Liu, Z., Knill, D. C., & Kersten, D. (1995). Object classification for human and ideal observers. *Vision Research, 35,* 549–568.

Mallot, H. A., Franz, M. O., Schölkopf, B., & Bülthoff, H. H. (1997). The view-graph approach to visual navigation and spatial memory. In *Proceedings of the Seventh International Conference on Artificial Neural Networks* (pp. 751–756). New York: Springer.

Mallot, H. A., & Gillner, S. (2000). Route navigating without place recognition: What is recognised in recognition-triggered responses? *Perception, 29,* 43–55.

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.

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

Najemnik, J., & Geisler, W. S. (2005, March 17). Optimal eye movement strategies in visual search. *Nature, 434,* 387–391.

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

O'Neill, M. J. (1991). Effects of signage and floor plan configuration on wayfinding accuracy. *Environment and Behavior, 23,* 553–574.

Ruddle, R. A., Payne, S. A., & Jones, D. M. (1997). Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied, 3,* 143–159.

Schölkopf, B., & Mallot, H. A. (1995). View-based cognitive mapping and path planning. *Adaptive Behavior, 3,* 311–348.

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). New York: Academic Press.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7,* 1–16.

Sondik, E. (1971). *The optimal control of partially observable Markov decision processes.* Unpublished doctoral dissertation, Stanford University.

Thompson, W. B., Willemsen, P., Gooch, A. A., Creem-Regehr, S. H., Loomis, J. M., & Beall, A. C. (2004). Does the quality of the computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments, 13,* 560–571.

Tjan, B. S., Braje, W. L., Legge, G. E., & Kersten, D. (1995). Human efficiency for recognizing 3-D objects in luminance noise. *Vision Research, 35,* 3053–3069.

Tjan, B. S., & Legge, G. E. (1998). The viewpoint complexity of an object recognition task. *Vision Research, 15/16,* 2335–2350.

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

Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience, 25,* 7169–7178.

(*Appendix follows*)

Appendix

Illustration of the Ideal Navigator Algorithm

Table A1

An Illustration of the Actions and Observations Used to Compute the Estimates in Table 1 When the Model Assumes That It Is Starting at D-North and Starts With a Rotate-Left Action

Action no.	Action		Resulting state		Observation		Belief state	
	D-North	E-South	D-North	E-South	D-North	E-South	D-North	E-South
1	Left	Left	D-West	E-East	Wall	Wall	D-W, E-E	D-W, E-E
2	Right	Right	D-North	E-South	C-LRW	C-LRW (T-junction)	D-N, E-S	D-N, E-S
3	Forward	Forward	B-North	H-South	Wall	Wall	B-N, H-S	B-N, H-S
4	Right	Right	B-East	H-West	C-W	C-C-Wall	B-E	H-W
5	Forward	Right	C-East	H-North	Wall	C-C-C-LRW	C-E	H-N
6	Done	Forward		E-North		C-C-LRW	C-E	E-N
7		Forward		D-North		C-LRW		D-N
8		Forward		B-North		Wall		B-N
9		Right		B-East		C		B-E
10		Forward		C-East		Wall		C-E
11		Done		C-East		Wall		C-E
Total	5	10						

Note. The first column on the left illustrates the action number. The second column illustrates the actions made when this assumption is correct (D-North columns), and the third column illustrates when this assumption is incorrect (i.e., observer is actually at E-South). The estimates of the number of actions (5 and 10) correspond to the expected cost when this assumption is correct (initial state really is D-North) and when it is not correct (initial state was actually E-South). A similar computation would be computed starting with a right turn and a forward action. In addition, the same computation is carried out assuming that the observer is at E-South. The observations can be a corridor (C) and the intersections on the left (L) and the intersections on the right (R) or a wall (W).

Table 1 illustrates how the model computes the cost for the next move in the situation illustrated in Figure 1. However, Table 1 does not provide an illustration for how those calculations were generated. Table A1 illustrates the actions, the expected observations at each step, and the current belief at each step when assuming the initial state is D-North and starting with the action rotate-left. A similar analysis would have to be done for all of the other actions (rotate-right, forward) and all of the other possible states. As shown in Table A1 in addition to Table 1, the model has to consider both the condition when its assumption is correct (i.e., the initial state is actually D-North) and when this assumption is incorrect (i.e., the initial state is actually E-South).

The second column in Table A1 illustrates the shortest route assuming that the initial state is D-North and this assumption is correct.[A1] The first line of this column (Action 1) illustrates the assumed initial action. Remember, the model computes the expected cost for starting with each available action. The first action is rotate-left. Column 4 illustrates the new state that the observer would be in if the model's initial state was actually D-North. The resulting state assuming that the initial state is D-North and the observer turned left is D-West. Column 5 illustrates the resulting state assuming the initial state is E-South. Column 6 illustrates the expected observation at the resulting state assuming the initial state is D-North, and column 7 shows the expected observation at the resulting state if the initial state is E-South. Notice that after the initial action, the expected observation is the same. Column 8 illustrates the model's current belief assuming that the initial state is D-North, and column 9 illustrates the model's belief vector if the initial state is E-South.

There are a number of important aspects of the model that are clearly illustrated in Table A1. The first point to note is that the model is able to reach the goal in five actions if the true state is D-North *and* the first action is a rotate-left action. The second aspect to note is that actions in Column 3 are the same (true state is E-South) for the first four actions. Remember, this table is computing the cost *assuming* that the initial state is D-North. After the fourth action, the model expects to receive a different observation depending on whether the true state is D-North or E-South. This is illustrated in Columns 6 and 7 of the table. More specifically, the expected observation when the model reaches B-East is different from the expected observation at H-West. Also note that in columns 8 and 9 the model's current belief also diverges at this point. If the model observes a single corridor dead-end (C-W), the model will know that it is at B-E and the initial state was D-North. But if the model observes two consecutive corridors and then a wall (C-C-W), then the model knows that its initial state was really E-South and its current state is H-West.

Row 5 (Action 5) is the point at which these two cases diverge. If the true state was actually D-North, then the optimal action is to move forward and reach the goal at C-East. However, if the model was actually at E-South, the model now has to "recover" from this incorrect assumption by turning around and moving toward the goal. Remember, Table A1 illustrates the routes when the model assumes that it is at D-North and the first action is rotate-left. The model will compute a similar table for all of the possible states, including E-South and all of the initial actions. The summary of each of these computations can be found in Table 1.

[A1] It may still be unclear how these specific actions were generated. As mentioned before, the model does a breadth-first search through the state-transition matrix until it finds a route that reaches the goal with no remaining uncertainty starting with a specific action.

Received December 29, 2004 ■