Lost in Virtual Space II: The Role of Proprioception and Discrete Actions when Navigating with Uncertainty

Brian J. Stankiewicz & Kyler Eastman
University of Texas, Austin
Department of Psychology & Center for Perceptual Systems

Under Review: *Environment & Behavior*

Abstract

Recent studies in spatial navigation by Stankiewicz, Legge, Mansfield and Schlicht (2006) investigated human navigation efficiency in indoor virtual reality (VR) environments when participants were uncertain about their current state in the environment (i.e., they were disoriented or "lost"). These studies used quantized actions (i.e., rotate-left or right by 90°and translate forward a specific distance) and lacked vestibular and proprioceptive cues as participants navigated through the virtual spaces. The current studies investigated whether the results from these studies generalize to more realistic conditions in which participants move continuously through space and have proprioceptive and vestibular information available to them. The results suggest that for large-scale way-finding tasks, desktop VR may provide a valid way to examine human navigation with uncertainty.

Introduction

Spatial navigation is a behavior that is used hundreds if not thousands of times each and every day of our lives. We use it when we are going from our living room to our kitchen in our home, and we use it when we are traveling across the city from our home to the grocery store. In addition to being able to readily and rapidly navigate through these large-scale spaces, we can also rapidly acquire useful knowledge about novel environments. Once we have acquired this internal representation of a large-scale space, we can use this representation to easily navigate from one location to another within that space. This internal representation of an external large-scale space is typically referred to as a *cognitive map* (Tolman, 1948) or a *cognitive collage* (Tversky, 1993).

These internal representations are useful when we know our current *state* (position and orientation) in the environment in that they allow us to move from one state (e.g., our

This research was partially supported by a grant from the Air Force office of Scientific Research (FA9550-05-1-0321 and FA9550-04-1-0236) a grant from the National Institutes of Health (EY16089-01) and a grant from the University XXI/Army Research Labs to BJS. Correspondence concerning this research should be addressed to Brian J. Stankiewicz, Ph.D. at bstankie@mail.utexas.edu.

office) to another (e.g., our home) with relative ease. Furthermore, they can be used to plan routes between two known locations in which we have never traveled directly between before (e.g., we may have traveled from our home to the grocery store, and from our home to our office, but we may have never travelled from our office to the grocery store).

In addition to planning and traveling between known states within the environment, this internal representation is also useful for localizing ourselves within a familiar environment after becoming disoriented or "lost". When one is lost in a familiar environment, one can represent this as a case of *state uncertainty*. That is, the observer has a certain amount of uncertainty about which position and orientation (state) that they are in within the large-scale space (such as a city or a complex building). By making observations within the environment (e.g., identifying buildings in a city, or seeing a painting in a building) and making actions (translating and rotating) observers can gather and integrate this information to reduce their uncertainty until there is none remaining.

Previous studies by Stankiewicz, Legge, Mansfield, and Schlicht (2006) have investigated human navigation *efficiency* when participants are lost in virtual environments. Stankiewicz et al. (2006) investigated human navigation efficiency under a number of conditions to try and determine how efficiently participants performed and determine the cognitive limitation that prevented participants from navigating optimally. To measure human navigation efficiency, Stankiewicz et al. (2006) developed an ideal navigator that was based upon a *Partially Observable Markov Decision Process* (POMDP) (Cassandra, Kaelbling, & Littman, 1994; Cassandra, 1998; Kaelbling, Cassandra, & Kurien, 1996; Kaelbling, Littman, & Cassandra, 1998; Sondik, 1971). The Stankiewicz et al. (2006) experiments were conducted using **desktop virtual reality**. In these studies, participants observed the environment from a first person perspective (see Figure 1 as an example of the type of display used in these studies), but remained sitting and made button presses to move through the environment. In these studies participants did not have access to any *egothetic* cues[1] that are typically available when navigating in the "real world". Furthermore, participants moved by using quantized actions in the Stankiewicz et al. (2006) studies. In these studies they rotated left by 90°or right by 90°or moved forward a prescribed distance by making specific button presses. These quantized actions are not typical in the "real world". Instead we move through space in a continuous manner. The studies described here investigated whether human efficiency changes as a function of whether participants have access to egothetic cues or not and whether they move discretely or continuously through the environment.

The results from Stankiewicz et al. (2006) suggest that the cognitive limitation in efficiently solving these tasks was in the participant's inability to accurately update and maintain the set of states (positions and orientations) that they would be within the large-scale space, given their previous actions and observations. The experiments described here investigated whether these effects generalize to more realistic conditions – specifically, whether performance will change with the addition of proprioceptive and vestibular information and the use of continuous movements instead of discrete movements (as used in the Stankiewicz et al. (2006) studies). To investigate this question we conducted a series of experiments that used the same paradigm as Stankiewicz et al. (2006) but had participants navigate through

[1]Egothetic cues are internal cues that provide information about the changes in position and orientation due to changes in the vestibular system and proprioceptive cues from movement of the legs.

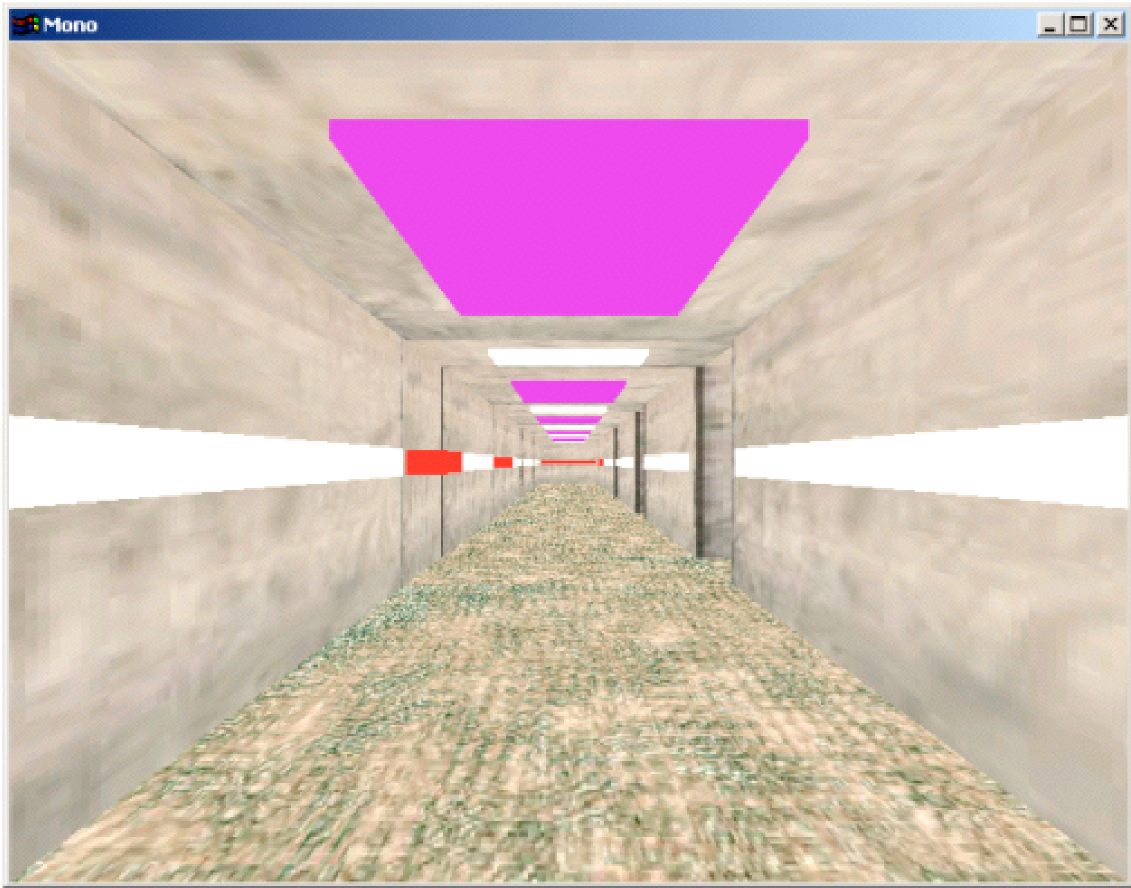

Figure 1. An image from the studies conducted by Stankiewiewicz et al. 2006. These studies used indoor virtual environments in which participants navigated using key presses to move through the environment. The environments were visually sparse to produce a certain amount of state uncertainty within the task.

the environments using immersive technology in which participants navigated through the virtual space using a head-mounted display in which their position and orientation was also tracked. This provided the participants with proprioceptive information that was not available to them in the desktop virtual reality environments.

Using VR to Study Human Navigation Ability

Over the last few decades, virtual reality (VR) has been increasingly popular in psychological applications. (e.g., see, Mallot & Gillner, 2000; Gillner & Mallot, 1998; R. Ruddle, Payne, & Jones, 1997; Steck & Mallot, 2000; Schölkopf & Mallot, 1995; Gillner & Mallot, 1998; Franz, Schölkopf, Mallot, & Bülthoff, 1998; Chance, Gaunet, Beall, & Loomis, 1998; Stankiewicz et al., 2006; Fajen & Warren, 2003). Psychologists have been attracted to the ease of natural interaction with an environment without actually having to construct it, which can save a great amount of time and money. In addition, the possibility of manipulating the environment precisely and systematically now exists in a way that was never possible before.

An issue that persists with virtual environments is whether the results obtained in VR environments are applicable to how people navigate in real environments. It appears that under some conditions, the results obtained in virtual environments do **not** generalize to real environments but, under other conditions, the results do generalize.

For example, one question that begins to address this issue is whether the knowledge acquired in a virtual environment transfers to the corresponding real environment. The argument is, that if the information transfers then this suggests that the underlying representation extracted in the VR environments simulates those that are generated in real environments. A number of researchers have investigated the transfer of knowledge from VR to real environments (e.g., see Darken & Banker, 1998; Koh, Wiegand, Garnett, Durlach, & Shinn-Cunningham, 1999; R. A. Ruddle & Jones, 2001; R. Ruddle et al., 1997; Thorndyke & Hayes-Roth, 1982; Waller, Hunt, & Knapp, 1998; Witmer, Bailey, & Kerr, 1996). For example, Witmer et al. (1996) investigated how well the knowledge acquired in virtual environments would transfer to real environments. In these studies, participants were trained in either the real building, the virtual environment, or by using verbal directions and photographs from the real building. The virtual environment was rendered with high fidelity (the researchers carefully modeled the structure of the building along with the collection of object landmarks within the building). participants were trained on a specific route in one of these conditions and then were instructed to take the same route in the real environment. They found that participants made fewer wrong turns and took less time to travel a complex route when they were trained in the real environment than when they were trained in the immersive virtual environment. However the number of wrong turns (immersive=3.3 versus Real=1.1) was small comparatively to the number of potential wrong turns (over 100).

In another study by Koh et al. (1999) participants were given a 10-minute training session in either the real building, an immersive VR display, or a desktop VR display. The immersive VR display consisted of a head-mounted display where movements through the environment were controlled by a joystick. The virtual environments were rendered with high fidelity (the structure of the building was generated from architectural plans of the building and the walls were texture wrapped with photographs taken from the real environment). Koh et al. (1999) had participants estimate the direction and distance from one location to another unobservable location within the environment. They found that participants performed just as well when they were trained in the VR conditions (either immersive or desktop) as they did in the real environment. This result suggests the possibility of generating an accurate spatial representation using virtual reality.

Waller et al. (1998) also conducted a series of studies investigating how well participants transferred spatial knowledge about an environment when trained in the real environment, a map, desktop virtual reality, immersive virtual reality (two minute exposures to the VR environment) and long-immersive virtual reality (five-minute exposures to the VR environment). In these studies, participants were trained by moving through the environment using one of these formats (or in the control condition participants were not given any training). After each training session, the participants were blindfolded and taken to a “real” version of the environment. Participants were instructed to touch a series of landmarks within the environment while blindfolded (these landmarks were shown in the training session). Waller, et al. recorded the time to complete the route blindfolded and the number of times the participants bumped into the walls. Participants participated in six training/test sessions. In the first session, participants who were trained in the real environment performed significantly better than the participants who were trained in the other conditions. By the sixth session, performance for participants who were trained in the long immersive condition was not significantly different than those participants who were trained in the real condition.

These studies suggest that under certain conditions the representation generated in a VR environment may transfer to real environments, suggesting that the representations used in the two different spaces may be similar. However, other studies have shown that under more specific conditions human performance is different in virtual environment than they do in real environments. For example, Thompson et al. (2004) had participants judge distances between two points in virtual and real environments. Participants consistently underestimated the distances in the virtual environments relative to their measurements in the real environments. These underestimations occured even when the displays were photorealistic images. Furthermore (and perhaps more pertinent to the current studies) Chance et al. (1998) studied the effects of using egothetic cues when doing a dead-reckoning task. In this task, participants would walk two edges of a triangle and then they were asked to walk to the starting point. In these studies they manipulated whether participants had access to their vestibular cues (physical turning) or proprioceptive cues (locomotion). These cues were redundant with the visual cues that were given in the display. They found that participants' performances declined significantly when the vestibular cues were not available to the participant. However, if the vestibular cues were available, but the proprioceptive cues were not, then participants performed just as well as if they had both egothetic cues.

Summary: Utility of Using Virtual Reality. In summary, VR can be a powerful tool for understanding issues associated with spatial navigation behavior. However, one needs to be careful when generalizing their results from tasks that are conducted in VR to spatial abilities in real environments. As is made clear by the previous examples, generalization may or may not be appropriate. The current studies extend the work conducted by Stankiewicz et al. (2006). The current studies investigate the role of egothetic cues in addition to quantized actions when re-orienting oneself in a familiar environment.

The effects of egothetic cues and continuous actions

Though there are many differences between the VR spaces used in the Stankiewicz et al. (2006) study and real environments the current paper focusses on two important aspects. The first is the fact that participants did not have access to any egothetic cues while wayfinding in the original studies. Participants sat at a desktop computer and moved through the virtual space by pressing keys. The visual cues provided sufficient cues to identify the amount of translation or rotation, but as has been shown in Chance et al. (1998), visual cues are not necessarily adequate for accurate spatial updating. To investigate the effects of egothetic cues, we will have participants navigate using immersive technology in which participants will walk through real space with visual information will be presented to the participant on a *head-mounted display* (HMD).

The immersive condition will add egothetic cues that are not available in the desktop environment. However, it also changes the types of actions that participants can make. That is, in the immersive condition participants can move freely through the environment, or more specifically, their actions are continuous rather than discrete as in Stankiewicz et al. (2006). To investigate the use of continuous rather than discrete actions, we will also run participants in a condition in which they will move through the environment using a joystick. In this condition, participants will not have access to their egothetic cues, but they will move continuously through the space. Table 1 illustrates the three conditions used in the current study and whether or not they have egothetic information or use continuous motion.

Table 1: An illustration of the three conditions used in this study and the specific manipulations provided by these conditions.

	Experiment Conditions		
	Key Press	Joystick	Immersive
Egothetic	NO	NO	YES
Continuous	NO	YES	YES

To address the effect of egothetic cues we will compare performance in the Joystick condition to that of the Immersive condition. The primary difference when running with the joystick versus the immersive environment is the egothetic cues. To investigate the effects of continuous versus discrete actions we will compare performance in the Joystick versus the Key Press conditions. If quantizing the actions has an effect on human navigation performance, there should be a change in efficiency in these two conditions.

The Ideal Navigator

In the studies conducted by Stankiewicz et al. (2006) and the current studies, participants navigated with a certain degree of uncertainty. In both of these studies, participants were familiarized with a novel virtual environment[2]. As illustrated in Figure 1, the environments are very sparse – that is, they did not contain any object landmarks (e.g., drinking fountains, pictures, etc.). It should be pointed out that even with perfect perception, most of the observations within the environment are not unique. Thus, given that the observers (human and ideal) are starting from a randomly selected state, in most trials they will start the task with a certain amount of state uncertainty due to the fact that multiple states can generate the same observation. Stankiewicz et al. (2006) used a POMDP to formalize the ideal navigator (also ,see Sondik, 1971; ?, ?; Cassandra, 1998; Kaelbling et al., 1996, 1998). To illustrate the POMDP approach in spatial navigation, we will provide a formal description of the underlying equations and provide an example using a simple environment. The purpose of this illustration with an example is to provide both an intuitive understanding of how the model works in addition to a formal understanding of the underlying mathematics.

[2]Participants explored the environment until they could draw the environment correctly twice in a row

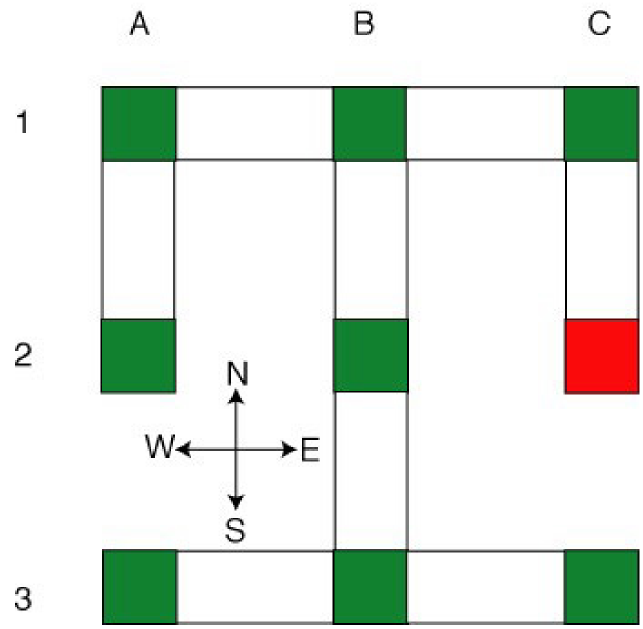

Figure 2. A simple map to demonstrate how the ideal navigator works. In this map the observer can be at any one of the green or red squares facing one of the four cardinal directions (N,S,E,W). The red square defines the "goal state" of that the observer wants to reach in the minimal number of actions.

For those who already understand the POMDP approach and/or are simply interested in the empirical results can skip the following section on the Ideal Navigator.

Defining the state and observation space. The ideal observer approach relies on converting a large continuous space into a quantized space. This is illustrated in Figure 2 as the set of states that the observer could be in. In Figure 2 the observer can only be on one of the red or green squares facing one of the four cardinal directions (North, South, East or West). In this environment there are a total of 36 different states that the observer can be in (9 positions $\times$ 4 orientations).

For each of these states we can define the *observation* that is expected from that state[3]. In Table 2 we have provided all of the observations and the states that could have generated that observation from the environment specified in Figure 2. As can be seen from this table, only four states (A1-E, C1-W, C3-W and A3-E) have unique views. All of the other states within the environment generate views that could be generated by at least one other state. Thus, if one were to be randomly placed within the environment only 4 times out of 36 (11% of the time) will the observer have no uncertainty about their location

[3]Although with the current studies and the studies by Stankiewicz et al. (2006) the observations were deterministic, one can actually model situations when the observations are noisy. That is, given a particular state the probability of an observation is 1.0 for one observation and 0.0 for all others. However, one might have observation noise such that the probabilities are not 1.0 and 0.0 (for more see, Cassandra et al., 1994; Cassandra, 1998; Kaelbling et al., 1996, 1998; Sondik, 1971)

Table 2: The set of observations and the states that could generate that observation in the environment shown in Figure 2.

Observation	States
Wall	A1-[N,S,W]; B1-N; C1-[N,E]; A2-[W,E,S]; B2-[E,W]; C2-[W,E,S]; A3-[N,S,W]; B3-S; C3-[N,S,E]
Hall-RightHall-Hall-RightHall	A1-E
Hall-Hall-RightHall/LeftHall	B1-S; B3-N
Hall (Dead End)	C1-S; B3-[E,W]
Hall-RightHall (Right L-junction)	B1-E,A2-N
Hall-LeftHall (Left L-junction)	B1-W;C2-N
Hall-LeftHall-Hall	C1-W; A3-E
Hall-RightHall/LeftHall (T-junction)	B2-[N,S]
Hall-RightHall-Hall	C3-W

within the environment. When deciding what action to generate the observer must take into account the likelihood that they are in a particular state within the environment.

Later we will use this table to generate the likelihoods that the observer is in a given state given its previous actions and observations (i.e., Belief Updating).

Defining the transition matrix. In the original Lost in Virtual Space experiments (Stankiewicz et al., 2006) and the *Keyboard* condition in the current study, participants moved through the environment by making one of three actions: **Rotate-Left** 90°, **Rotate-Right** 90°and move **Forward** one hallway unit. The transition matrix defines the probability of the resulting state if the observer generates a particular action in a given state (i.e., $p(s'|s,a)$). An illustration of the transition matrix is shown in Table 3 and Figure 3. This is a graphical representation of the transition matrix (and only a portion of the transition matrix for the entire environment shown in Figure 2). As can be seen in Table 3 the transition matrix makes explicit the resulting state given an initial state and a specific action[4].

Belief updating. Given the specifications of the observations and the specifications of the transition matrix one can begin to generate hypotheses about their current state within the environment given prior actions and observations. Equation 1 provides the Bayesian updating rule that computes the likelihood of being in a specific state (s') given the prior belief, the current observation and the action just generated.

[4]In the current studies and the studies conducted by Stankiewicz et al. (2006) the actions were deterministic. That is, the probability of the resulting state given an action and an initial state was 1.0 for one state and 0.0 for all other states. One does not need to assume deterministic actions, but instead the resulting state may be noisy. Modeling noisy actions can be found in Cassandra et al. (1994) and the other POMDP papers listed.

Table 3: The complete transition matrix for the environment shown in Figure 2 with the three actions Rotate-Left, Rotate-Right and Forward. The state on the far left of the table is the starting state, and the three columns following the starting state specify the resulting state if the observer were to generate that action.

Initial State	Rotate-Left	Rotate-Right	Forward
A1-N	A1-W	A1-E	A1-N
A1-E	A1-N	A1-S	B1E
A1-S	A1-E	A1-W	A2-S
A1-W	A1-S	A1-N	A1-W
B1-N	B1-W	B1-E	B1-N
B1-E	B1-N	B1-S	C1-E
B1-S	B1-E	B1-W	B2-S
B1-W	B1-S	B1-N	A1-W
C1-N	C1-W	C1-E	C1-N
C1-E	C1-N	C1-S	C1-E
C1-S	C1-E	C1-W	C2-S
C1-W	C1-S	C1-N	B1-W
A2-N	A2-W	A2-E	A1-N
A2-E	A2-N	A2-S	A2-E
A2-S	A2-E	A2-W	A2-S
A2-W	A2-S	A2-N	A2-W
B2-N	B2-W	B2-E	B1-N
B2-E	B2-N	B2-S	B2-E
B2-S	B2-E	B2-W	B3-S
B2-W	B2-S	B2-N	B2-W
C2-N	C2-W	C2-E	C1-N
C2-E	C2-N	C2-S	C2-E
C2-S	C2-E	C2-W	C2-S
C2-W	C2-S	C2-N	C2-W
A3-N	A3-W	A3-E	A3-N
A3-E	A3-N	A3-S	B3-E
A3-S	A3-E	A3-W	A3-S
A3-W	A3-S	A3-N	A3-W
B3-N	B3-W	B3-E	B2-N
B3-E	B3-N	B3-S	C3-E
B3-S	B3-E	B3-W	B3-S
B3-W	B3-S	B3-N	A3-W
C3-N	C3-W	C3-E	C3-N
C3-E	C3-N	C3-S	C3-E
C3-S	C3-E	C3-W	c3-S
C3-W	C3-S	C3-N	B3-W

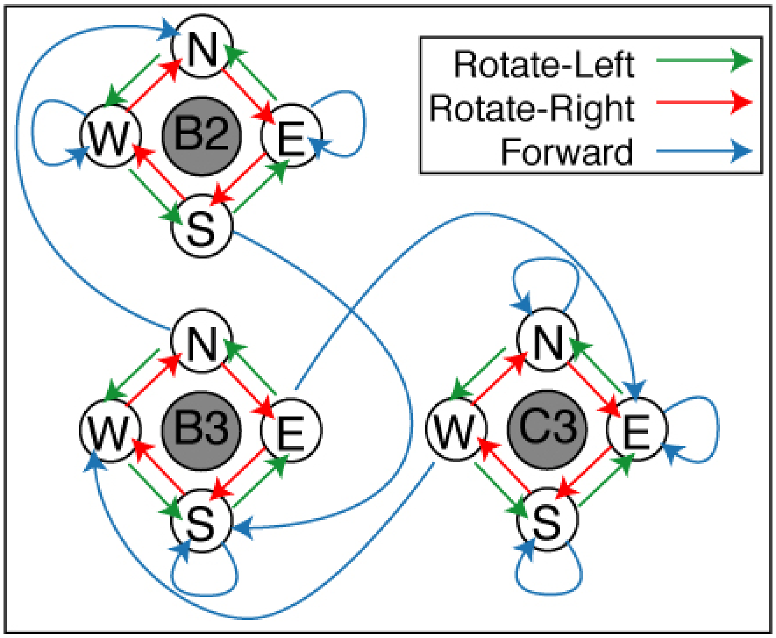

Figure 3. A graphical representation of a part of the transition matrix for the environment shown in Figure 2. The arrows specify the resulting state if the observer was in one of the states and generated one of the three actions (Rotate-Left, Rotate-Right and move Forward). This transition matrix is only for a portion of the environment shown in Figure 2.

$$p(s'|b,o,a) = \frac{p(o|s',b,a)p(s'|b,a))}{p(o|b,a)} \tag{1}$$

b is the probability that the observer is in a state given the prior observations and actions. For this example, the likelihood of being in any given state is initially uniform (i.e., $p(s') = 1/36 = 0.0277$).

Let us assume that the initial observation (o) is 'Hall-RightHall/LeftHall' (a T-junction) from the environment in Figure 2. The probability of getting that observation given a specific state and the prior probabilities ($p(o|s', b)$) is 0.0 for all of the states with the exception of B2-N and B2-S. The value of $p(o|s', b)$ for these two states is 1.0.

The probability of being in s' given the prior probabilities is the prior probability of being in that state (i.e., $p(s'|b) = 0.0277$). Finally, the probability of getting the observation given the prior belief ($p(o|b)$) is 0.055 or 2/36 (two states out of the 36 states give this same observation).

$$\begin{aligned} p(B2 - N|'Hall - RightHall/LeftHall') &= \frac{1.0 \times 0.2777}{0.555} && (2) \\ &= 0.5 && (3) \\ p(B2 - S|'Hall - RightHall/LeftHall') &= \frac{1.0 \times 0.2777}{0.555} \\ &= 0.5 \end{aligned}$$

We can extend this equation to give us the transition between beliefs after an action is generated. That is τ provides us with the likelihood of the resulting belief vector given our current belief vector and a specific action. (see fig. 4)

$$\tau(b, a, b') = \sum_{o \in O | p(s'|b,o,a)} p(o|a, b) \tag{4}$$

To illustrate the belief updating, we will use Equation (4) in addition to Table 3. For this example, let us assume that the observer generated the "Forward" action after its initial observation and that after this action the observer's observation is a "Wall". For 21 of the states the probability of seeing a wall is 1.0 (i.e., $p(o|s', b, o) = 1.0$). The probability of being in a particular state given the prior belief (B2-N=B2-S=0.5) and the action is 0.0 for all of the states with the exception of B3-S and B1-N. This can be shown by looking at Table 3 and looking at the intersection of the "Forward" column (a="Forward") and the B2-N and B2-S rows. This cell represents the resulting state given the action and initial state. Thus $p(B1-N|b, `Forward') = p(B1-N|b, `Forward') = 0.5$. The $p(o|b, a) = 1.0$. That is, given the current belief and a forward action, the probability of getting the observation of 'Wall' is 1.0. Equation 5 makes this updating function more explicit.

$$\begin{aligned} p(B3-N|b, `Wall', `Forward') &= \frac{1.0 \times 0.5}{1.0} \\ &= 0.5 \\ p(B1-S|b, `Wall', `Forward') &= \frac{1.0 \times 0.5}{1.0} \\ &= 0.5 \end{aligned} \tag{5}$$

Now let us imagine that the observer's next action is 'Rotate-Right' and following this observation it receives an observation of 'Hall-RightHall'. Given the prior belief and the action 'Rotate-Right', there are only two possible resulting states: B1-E and B3-W. The observation expected from B1-E is 'Hall-RightHall' while the observation expected from B3-W is 'Hall' (see Table 2). The probability of observing 'Hall-RightHall' from B1-E is 1.0, while the probability of getting this observation from B3-W is 0.0. Furthermore, the probability of making either one of these two observations ($p(o|b, a)$) is 0.5. As can be seen in Equation 6 the probability of being in B1-E is now 1.0 and the probability of being in B3-W is 0.0.

$$\begin{aligned} p(B1-E|b, `Hall-RightHall', `Rotate-Right') &= \frac{1.0 \times 0.5}{0.5} \\ &= 1.0 \\ p(B3-W|b, `Hall-RightHall', `Rotate-Right') &= \frac{0.0 \times 0.5}{0.5} \\ &= 0.0 \end{aligned} \tag{6}$$

With this observation we can see that all of the probabilities collapse onto the state B1-E – that is the observer has completely disambiguated itself within the environment.

Reward structure. Equation 1 formalizes how one would update their belief about their current state given their previous actions, observations and prior probabilities. However, it does not specify *which* action that the observer should take. In order to specify the optimal action one must specify a **Reward Structure** for the problem. A reward structure specifies the expected value (both positive and negative rewards) for generating an action in a particular state. For example, in Stankiewicz et al. (2006) the reward for generating a "Forward", "Roate-Left" and "Rotate-Right" action was the same (-1.0). This reward was the same for all of the states that the observer could be in.

However, for the experiments described here and the experiments in Stankiewicz et al. (2006) there was a goal state that the observers were attempting to reach. In the environment in Figure 2 the goal is C2-[N,S,E,W]. That is, the goal is to reach C2 facing in any particular direction. In order to model this we need to add a new action to the three shown in Figure 3. We will call this action 'Declare-Done'. We will assign a reward of 100 when this action is generated and the observer is in C2-[N,S,E,W] and -500 for all of the other states.

Given an explicit reward structure, one can begin to formalize the optimal action, given a specific belief about one's current state in the environment. In the following section, we will formalize how the optimal action can be selected given the current belief, transition matrix, and reward structure.

Choosing the optimal action

Using the Transition matrix, Reward Structure and Observation Matrix we can compute the optimal action for any given belief vector that the observer might have[5]. To compute the optimal action one must consider the *immediate reward* for generating a specific action in a specific state in addition to computing the future expected rewards. The following equation is the reward for generating a single action given that you have a particular belief vector (b). Figure 4 illustrates the set of actions that are available to the observer and the expected immediate rewards for generating each action (the equation for computing the immediate reward is in Equation 7).

$$\rho(b, a) = \sum_{s \in S} R(s, a) b(s) \tag{7}$$

ρ provides the *immediate* reward for making the action a and Figure 4 provides an illustration for the current example. As can be seen in Figure 4, there is no clear choice for the observer if their choice is myopic – that is, if they only consider the immediate rewards. If one were to choose the action that maximized the immediate rewards, at best, one would have to select from the three actions with an immediate reward of -1.

Given the structure of the problem, one can begin to consider the actions to help differentiate between which of the immediate actions is optimal. Computing the expected value when one is considering future actions is an iterative process – a process that has to consider future actions and beliefs. To compute the maximal expected value we want to compute the rewards that are acquired as the observer is moving through the problem. That is, the reward will be (Eq. 8):

[5] A belief vector is a particular probability distribution across all of the possible states in the environment.

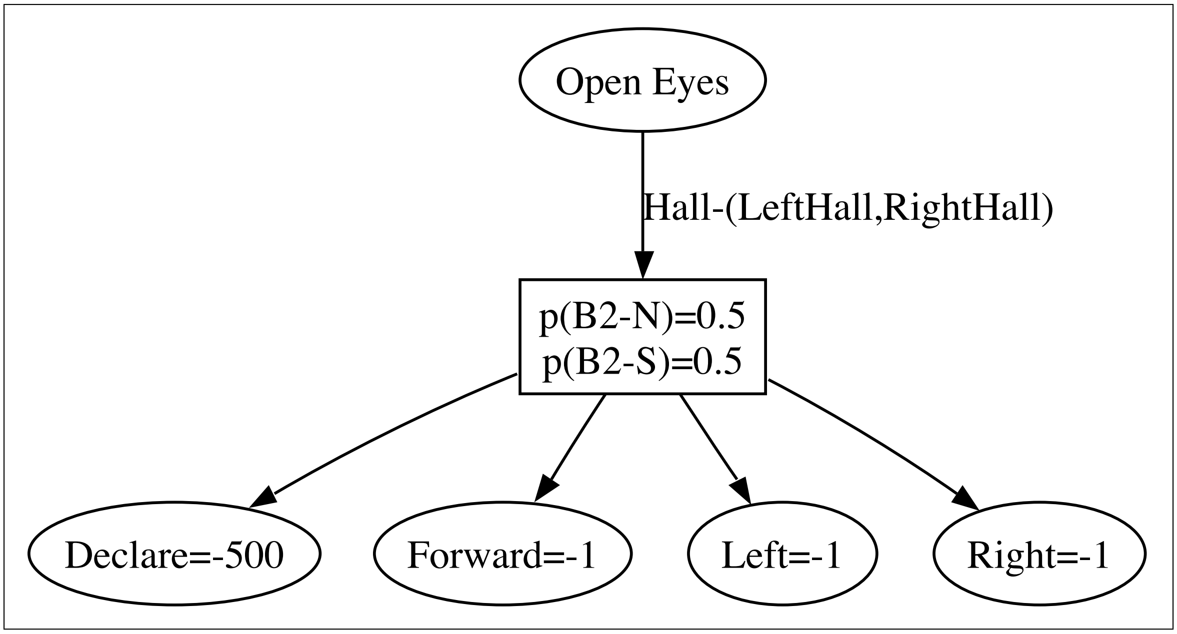

Figure 4. The **immediate rewards** (IR) for the four actions available to the observer given the current belief state.

$$R = \sum_{t=0}^{\infty} r(t) \tag{8}$$

To compute the optimal action for a given belief vector, one needs to consider all of the actions (A) that are available to the observer in the current state, the immediate reward for generating an action (ρ from Equation 7), and the transition function that specifies the likelihood that a new belief state will be generated (τ) and the expected states that one could be in given the observations that one could receive. γ specifies a *discount factor* that discounts future rewards relative to immediate rewards. By choosing the action that maximizes the following function, one can act optimally in the environment (see Eq. 9).

$$V(b) = \max_{a \in A} \left[\rho(b,a) + \gamma \sum_{b' \in B} \tau(b,a,b')V(b') \right] \tag{9}$$

Returning to our example, Figures 5 and 6 provide illustrations of the iterative manner of this problem. In Figure 5 the observer is faced with four different choices: Declare, Forward, Left (Rotate Left by 90°) and Right (Rotate Right by 90°). This is identical to Figure 4 with the exception that we have computed the **expected value** (V) instead of the immediate reward. The expected value considers the future rewards assuming that following the selection of the current action the observer will make the choice that maximizes the expected value. For the current example, given the a belief (p(B2-N=p(B2-S)=0.5) the expected value for declaring finished ("Declare") is -500 (see Section). However, if the observer moves "Forward" the expected value is 94, while the expected reward for rotating left or right is 92. If the observer chooses the action that maximizes its expected value, it should choose the 'Forward' action.

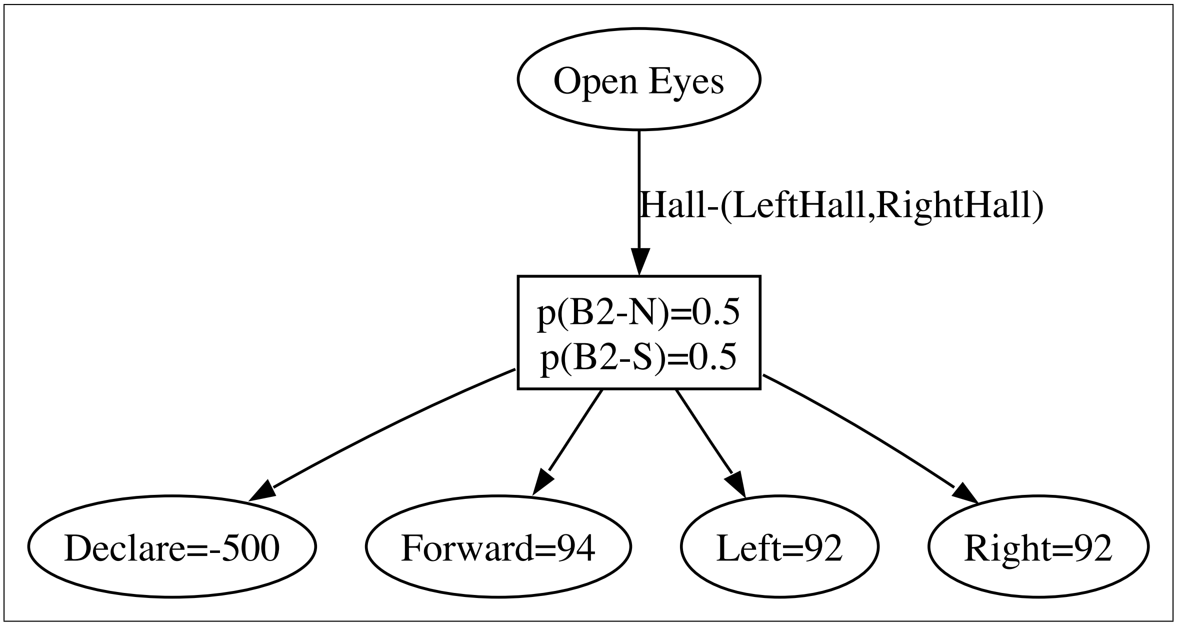

Figure 5. After the initial observation in our example of "Hall-LeftHall-RightHall" there are two different locations within the environment that the observer could be located with equal probability: B2-N and B2-S. The observer can choose between four different actions: Declare, Forward, Rotate Left by 90°("Left") or Rotate Right by 90°("Right"). Given the current belief state (p(B2-N)=p(B2-S)=0.5) we have computed the *Expected Value* (V) for generating each action. The expected value is the ultimate reward that the observer expects to receive if they generate the specific action and, following each action, they generate the action that has the greatest expected value. We have illustrated how the expected value is computed for the Forward action in Figure 6.

Figure 5 provides the end result of the expected value computation. We would like to provide an illustration of how this works with the current example. Figure 6 provides an illustration of how the value function might be used in Equation 9. In Figure 6 we show how the expected value of 94 is generated for the 'Forward' action. A similar process would be generated for all of the actions available (i.e., Left, Right, Declare).

There are a number of aspects that need to be pointed out in Figure 6 that illustrate how the POMDP works. First, note the main branching point in the illustration. This main branching point in the illustration is where the observer expects to receive one of two different observations given the current belief vector. When it is equally likely that it is at B1-N and B3-N the model expects one of two observations after rotating right by 90°– either a 'Hall' or a 'Hall-RightHall' (Right L-junction). The first thing to note is that given one observation or the other, the belief vector collapses on one state or the other (as shown by the new probabilities; p(B3-W)=1.0 if a "Hall" is observed and the probability of B1-E is 1.0 if a "Hall-RightHall" is observed). Second, note that selecting the action 'Right' is the weighted sum of the expected value of being in one state or the other.

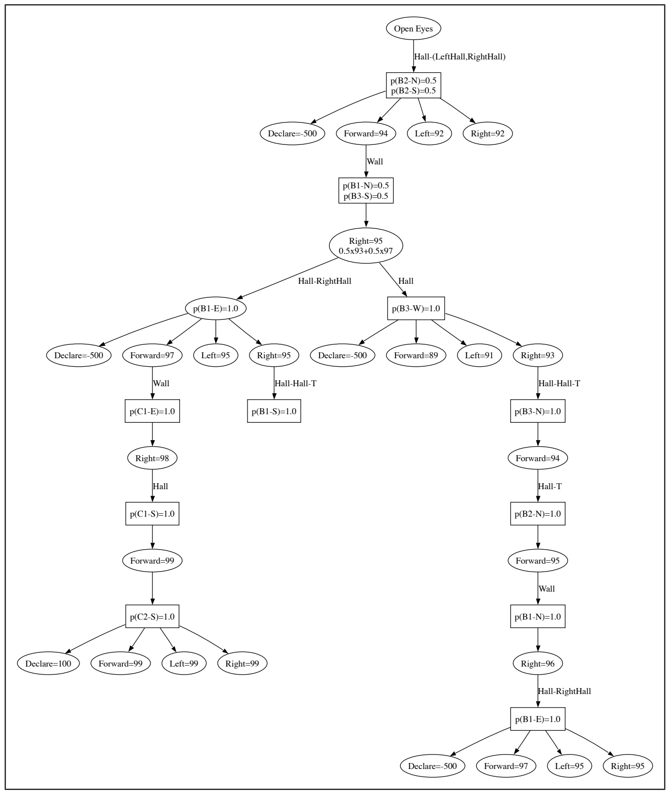

Figure 6. An illustration of how one might compute the expected value (EV) for a given action (Declare, Forward, Left, Right) given an initial belief vector (p(B2-N)=p(B2-S)=0.5). For each belief vector we compute the expected value if one were to select a particular action and do everything optimal following this action. However, if you start at the lowest node (Declare V=100) and work your way backward through the tree, one can see how these actions can be selected.

A second aspect to note from Figure 6 is that the observer can be in state B1-W in two different cases. The first is immediately following the 'Right' action at the divergent point mentioned above. The second is where the observer realizes that they started in B2-S, turns around and reaches that state a number of moves later. The important thing to notice is that the expected value at this state is the same regardless of how it was reached. One can take leverage of this property and take advantage of dynamic programming techniques to develop efficient algorithms for computing the expected values for a specific belief state.

Utility of using ideal observer analysis

Ideal observer analysis has been very useful for understanding specific perceptual issues that range from understanding the quantum limits of light detection (Hecht, Shlaer, & Pirenne, 1942) to many forms of visual pattern detection and discrimination (Geisler, 1989), to reading (Legge, Klitz, & Tjan, 1997; Legge, Hooven, Klitz, Mansfield, & Tjan, 2002) object recognition (Liu, Knill, & Kersten, 1995; Tjan & Legge, 1998; Tjan, Braje, Legge, & Kersten, 1995; Braje, Tjan, & Legge, 1995) and also in reaching tasks (Trommershäuser, Maloney, & Landy, 2003). The ideal observer provides a benchmark by which human observers should behave if they are behaving rationally. It should be noted that we are not arguing in this paper that humans *are* in fact optimal, but instead we use this technique to remove the effects of *task difficulty* from human behavior.

Human performance is going to be a function of both their internal cognitive functions **and** the inherent difficulty of the task (task difficulty). This task difficulty is simply the *objective* difficulty of the task. For example, imagine that in an experiment one is measuring the distance traveled to reach a known goal state. If one studies the human observer in different environments (e.g., ones of different sizes) one would expect that smaller environments are inherently "easier" than larger environments. That is, if one were to perform perfectly (i.e., no cognitive limitations), they would have to travel farther in the larger environments than the smaller environments. The ideal observer would be sensitive to these task demands. By comparing human performance to the ideal observer's performance we can remove aspects of these task demands to allow us to understand the cognitive limitations in completing the given task.

In the current studies, we are interested in using the ideal observer analysis because we are running participants in different environments. The environments are of the same 'size', but the degree of uncertainty that the observer may have to deal with may be significantly different in one environment over another. Furthermore, each individual trial within each environment has varying levels of difficulty. Therefore, we will normalize the human performance in these environments by using the ideal observer's performance.

Experiment 1

The main objective of these experiments was to determine whether the use of proprioceptive and vestibular information and the quantization of space has an effect on human navigation performance with uncertainty. In doing so, this experiment also evaluates whether the results of Stankiewicz et al. (2006) generalizes to more realistic navigation conditions.

To do this, these studies replicated aspects of the Stankiewicz et al. (2006) studies with participants in three navigation conditions: *Key Press*, *Joystick*, and *Immersive*. The Key Press condition replicated the original condition used by Stankiewicz et al. (2006) in which participants viewed the environment using desktop virtual reality using **discrete** actions. In the Joystick condition participants observed the environment using desktop reality, but made **continuous** movements through the space using a joystick. This condition included no extra modal information, but allowed the participant continuous movement, and required the participant to visually control his movements. The Immersive condition had participants move by walking in an immersive virtual reality system using a head-mounted display (HMD). This condition replicated the visual information of the original condition, but also gave the participant access to internal vestibular and proprioceptive information. The performance of the participant in the task under each of these conditions was normalized for task difficulty by an ideal observer model. In this way, the only variables changed between conditions will the two factors observed: action discretization and egothetic (vestibular and proprioceptive) information.

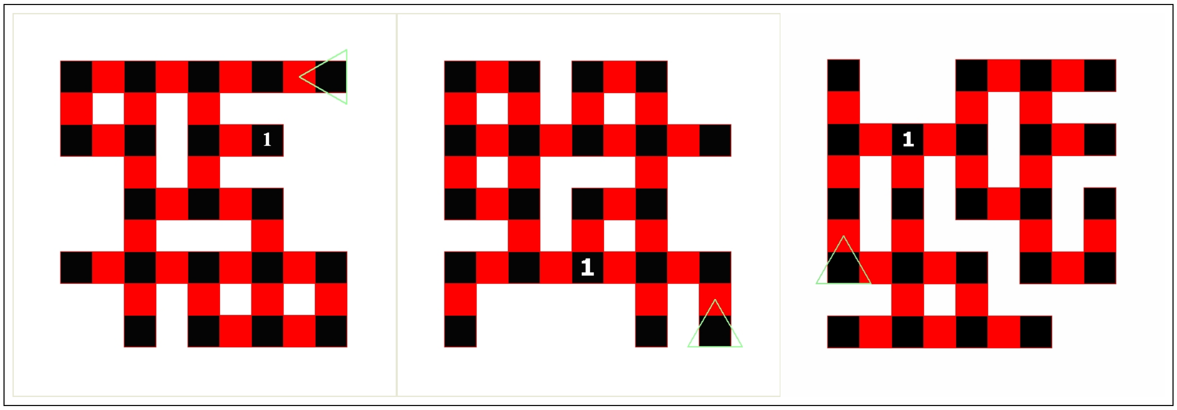

Figure 7. Maps of the environments used in the current studies. The triangle indicates the position in which participants started during the **Exploration Phase** of the study. The '1' indicates the goal state for each environment. During the *Exploration Phase* an auditory signal was given when participants walked over the goal state. During the *Experimental Phase*, participants were placed at a random position and orientation within the environment and instructed to move to the goal state in the shortest distance possible.

Methods

Apparatus. The experiment was run on an IBM Dell computer with a Virtual Research V8 head mounted display used for viewing the environments. The HMD has a 640x480 resolution display with a 60°diagonal viewing area. In the Immersive condition, participants were tracked using a WorldViz Precision Point Tracker that tracked the participants' position using video tracking and their orientation using a Intersense InertiaCube2. In the Key Press condition, the participant moved through the environment by making key presses that corresponded to a 90°clockwise rotation, a 90°counter-clockwise rotation or a forward translation of one hallway unit. After the participant made a key press, the computer would rotate or translate the virtual "camera" in the virtual space. In the Joystick condition participants moved through the environment using a ## joystick. In all conditions, participants viewed the environment using the V8 head mounted display.

Stimuli. Three different environments were used for the study that consisted of environments that consisted of twenty to twenty five identical hallway segments that intersected at 90°angles. The layouts from these environments can be found in Figure 7. All hallways were three meters high, 0.6 meters wide, and 0.6 meters long, making 1.2 meters between each intersection. Due to the limit of the room, the environment had to be encapsulated in a 5 m x 5 m square, making the longest hallway 4 segments or 4.8 meters. Red and white bars were placed at mid-height on all of the walls in the environment (See Figure 1). Red bars indicated an intersection while white bars indicated a hallway. This feature was intended to enhance the participants ability to discern the boundaries of the hallway.

Procedure. The current studies had three phases: *Exploration*, *Knowledge Test*, and *Experimental*. The purpose of the first two phases was to familiarize the participants with the environment and test their knowledge of the environment. The Experimental phase was the critical phase in which we tested participant's performances in these environments. Participants participated in all three phases in all three conditions (Key Press, Joystick and Immersive). The order of the conditions was counterbalanced across the six participants such that each condition occurred first, second and third equally often (twice). Each environment also appeared in each condition equally often.

Exploration Phase In the Exploration Phase participants explored a maze-like environment by navigating using either key presses, joystick or walking through the environment. The mode was dependent upon the current condition. Participants moved through the environment for a total of 100 meters. After which, their knowledge of the environment was tested (see Knowledge Test below). During the exploration phase, there was an auditory signal (in the form of a computer generated female voice announcing "Position One") when the participant reached the target location. Participants were informed that they needed to learn the environment well enough to draw the environment and the location of the target state. Participants were also informed about the Experimental phase in which they would start from a random location within the environment and would need to reach the target location in the minimum distance possible. If the participant needed to participate in the Exploration Phase more than once, they were always started at the same location within the environment.

Knowledge Test Phase A graphical user interface (GUI) was used for drawing the map. Although the Exploration Phase was conducted with the HMD, the Knowledge Test was completed using a 19-inch computer monitor. The GUI required the user to create the map from a top-down perspective. A small portion of the map was given consisting of the three or four hallway segments that the participant could see from their initial position in the Exploration Phase. Participants drew a map of the environment by selecting and de-selecting hallway segments until they felt that they had drawn the environment correctly.In addition to drawing the hallway structure of the building, the participants also specified the location of the target position. When they were finished, they selected a button indicating that they were finished. The computer evaluated whether the hallway structure that they generated matched the hallway structure of the environment that they explored. If participants did not draw the environment correctly, they returned to the Exploration Phase. They continued in the cycle of Exploration followed by Knowledge Testing until they could draw the environment correctly twice in a row.

Experimental Phase In the experimental phase participants started from each position of the map four times – once in each of the four primary directions. The order of the start position was randomized by the computer. Participants were instructed to move to the target location in the shortest distance possible. In the previous studies (Stankiewicz et al., 2006) participants were instantaneously "transported" to their start state (one advantage of using virtual reality). However, in the current studies we had people navigating in immersive spaces where they had to physically walk through the environment (the Immersive Condition). To get participants to move to their new location within the virtual environment, participants were "removed" from the maze-like environment and told to move to a large virtual arrow and face in the direction of the arrow. When the participants were standing on the arrow facing in the correct direction, they announced that they were "ready" and the experimenter pressed a button that made the observer re-appear inside the maze. To reduce the participant's ability to use their position within the physical space, the environment was randomly rotated in one of four different directions. This inter-trial movement allowed the participant to start from all states in the environment, while also providing little possibility that the participant was gaining information about their location at the start of the trial. Though it was possible to avoid this inter-trial movement in the joystick and key press conditions, it was included in all three conditions to keep the conditions as similar as possible.

Conditions

Keypress In the keypress condition, participants were confined to movement in between intersections of the environment. The participants pressed a button to move left and right by 90°and to move forward by one hallway segment. Participants moved at a constant speed of 1 meter per second and rotated at a constant speed of 90°per second.

Joystick Participants in the joystick condition were able to use the joystick to move continuously to every location in the environment. Preliminary measurements in the HMD environment showed that participants moved at a rate of 0.8 meters per second when walking at full stride. To match this, when participants pressed forward on the joystick, they moved forward at a rate between 0 and 0.8 meters per second. The participant moved at 0.8 meters/sec. when the joystick was fully engaged in the forward direction (slower speeds could be achieved by reducing the forward displacement on the joystick). When the participant pressed left or right, the participant would rotate up to a speed of 78°per second. Participants could not move through the walls of the environment. Unlike the Keypress, conditions, participants could also move backward in this condition.

Immersive In the Immersive condition, participants wore a V8 HMD. They walked in a 9-meter x 9-meter space. Though participants were told to avoid walking through the virtual walls of the environment, it could not be entirely avoided. Trials were removed in cases where participants moved through a wall to get from from one intersection to another. Participants were allowed to move backwards.

Participants

Six participants participated in the experiment, performing one session of 70-90 trials (depending on map size) for each of the conditions. The order of each condition was counterbalanced across participants. The participants included the two authors on this paper, two upperclass undergraduate students, and two university employees.

In order to keep the visual information in all of the conditions, participants viewed the environments using the HMD in all conditions. However, the HMD was only tracked in the Immersive Condition. Since the HMD was only being used as a monitor in the other two conditions, the display would not change when the participant moved his/her head. This, combined with the deliberate lack of peripheral vision, caused many of the participants

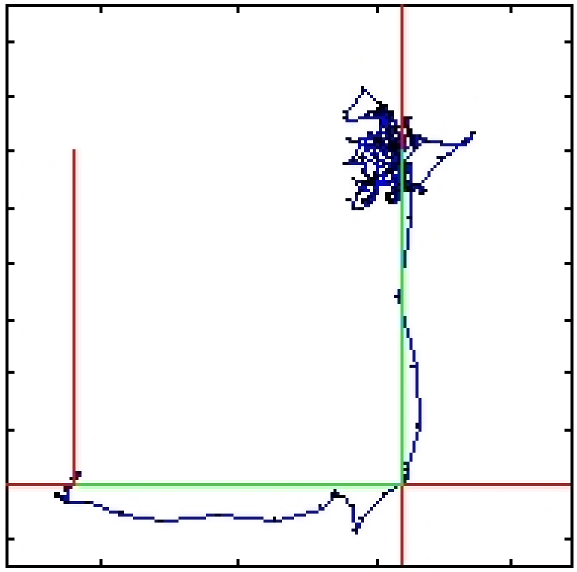

Figure 8. An example of small "micro-movements within a trial. The participant started at the upper right, repeatedly turned his head to look down the hallway, then proceeded to node at the lower left. By quantizing the movement, this trial gets recorded as 2.4 meters even though his actual accumulated trial distance is approximately triple this distance. allowed us to measure the distance traveled in all three condition on equal footing.

to become motion sick. Two participants, a graduate student from another lab and a paid volunteer, could not complete the joystick condition due to their level of uneasiness. All participants expressed some discomfort in all of the conditions. However, the joystick condition was found to be the most disturbing. All participants were encouraged to take 5-10 minutes breaks once every 10-15 minutes, according to their discomfort.

Results

Translation Efficiency

The main measurement in this study is the efficiency in which participants were able to reach the goal state. Efficiency was measured by the ratio of the distance traveled by the ideal navigator versus the human observer (see Eq. 10).

$$Efficiency = \frac{Distance(Ideal)}{Distance(Human)} \tag{10}$$

Measuring the distance in these three different conditions posed a series of challenges. One challenge was that, in the immersive condition, participants would make small, micro-movements (the head would move slightly forward and back and to the left and right) to look back and forth between hallways.

These micro-movements (see Figure 8) could not be made in the key press condition and were not typically made in the joystick condition. A second issue was that, in the

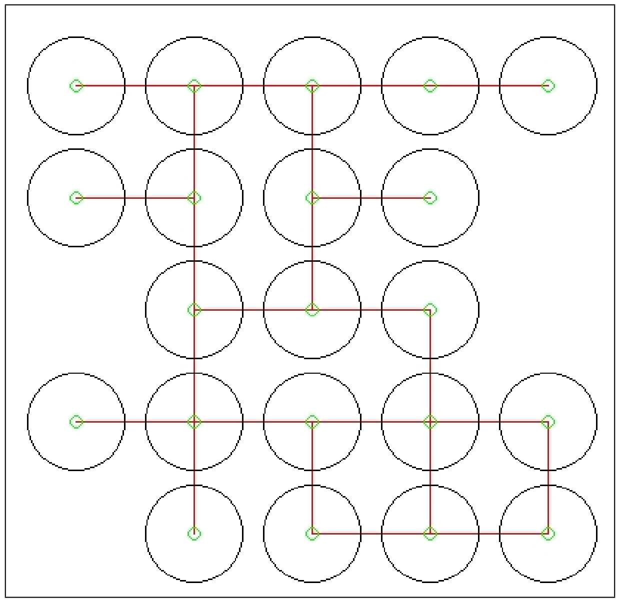

Figure 9. An illustration of the Nodal Quantization technique used to measure the participants distance traveled in the environments. A total of 1.2 meters were traveled when the participant left one node and entered another node within the environment. This allowed us to measure the distance traveled in all three condition on equal footing.

immersive and the joystick conditions, the participants could "cut the corners" to reduce the distance that they traveled. Because participants were restricted to traveling down the center of the hallways, this corner cutting strategy was not available in the key press condition[6].

To place all of the conditions in the same space, we used a *Nodal Quantization* technique. This form of quantizing created radii around each intersection that consisted of 40% of the distance between intersections or 0.48 meters (see Figure 9 for an illustration). A complete movement would be counted only if the participant left the radii of one node and moved into the radius of the next node or 60% of the distance of a hallway segment.

Using the Nodal Quantization technique we measured the distance that the participant traveled on each trial in each condition. We then computed the optimal distance traveled using the Nodal Quantization technique and computed the efficiency for each trial using Equation 10. Figure 10 shows the average efficiency for the six participants in this study.

To test the effect of egothetic cues on navigation performance we did a planned comparison two-tailed t-test between the performance in the Immersive and the Joystick condition and found no reliable difference ($t(5)=0.058$; $p=0.95$). To investigate if using quantized movements had a significant effect on the participant's navigation performance

[6] Although "cutting corners" and standing very still would allow the participants to be more efficient (as defined by the task), we are not interested in understanding these aspects of the navigation behavior. Instead, we are interested in understanding the role of egothetic cues in addition to quantization on macro navigation decisions.

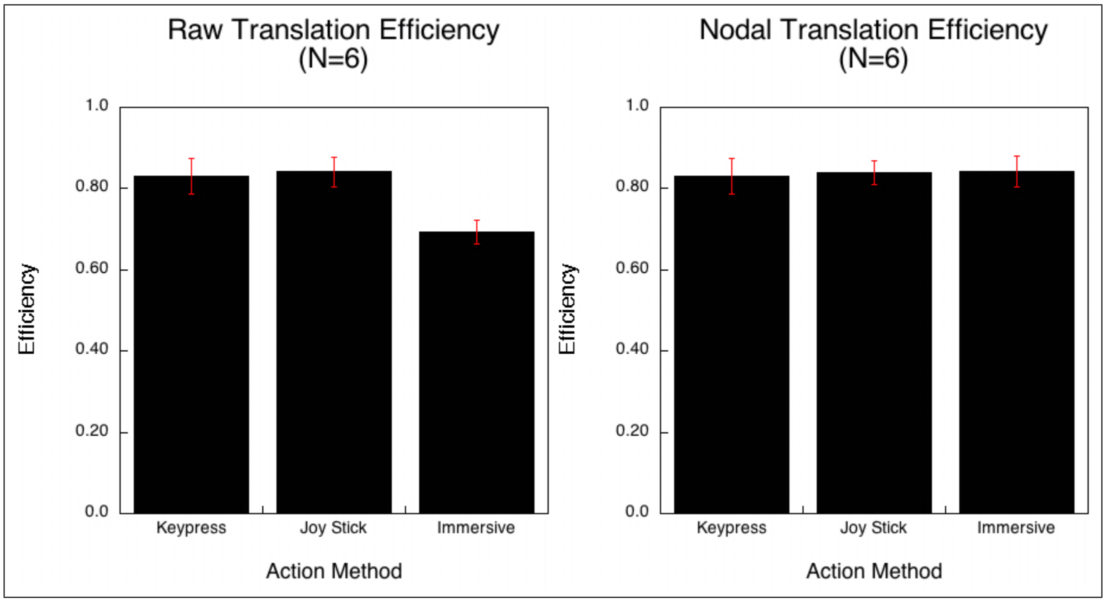

Figure 10. **Left.** Average navigation efficiencies for six participants who ran in all three navigation conditions. These efficiency performances are based upon the raw movements made by the participants. **Right.** Average navigation efficiencies for six participants who ran in all three navigation conditions. These efficiencies are computed by measuring the distances using the Nodal Quantization technique.

we did a planned comparison between the Key Press and Joystick Condition. There was no reliable difference between these two conditions either (t(**5**)=0.201; p=0.84). Finally, to see if adding both quantized actions and egothetic cues had any effect on navigation efficiencies we compared performance in the Key Press and Immersive conditions and again found no significant differences (t(**5**)=0.23; p=0.82). Power estimates were computed for all comparisons, and the number of participants with the same distribution of results that are needed to produce significance were greater than 1.3×10^4. We also computed performance on the raw distances which are also illustrated in Figure 10.

Learning

The method of interacting with the environment may have an effect on how quickly the subject is able to learn the environment. As a reminder, in the current study subjects explored the virtual environment for 100 meters. After the *Exploration Phase*, participants were instructed to draw the environment on a cartesian grid. Participants continued in the *Exploration phase* followed by the *Knowledge Test phase* until they drew the environment correctly twice in a row. We calculated the number of *Exploration phase* sessions the participant needed before drawing the first instantiation (of the two correct instantiations in a row) of the environment. The average number of *Exploration phase* sessions in *Keypress* condition was 2.667 (SEM=0.8327); *Joystick* condition was 2.667 (SEM=1.0066) and in the *Immersive* condition was 2.667 (SEM=0.5416). The number of *Exploration sessions* ranged from one to seven. There was no statistical difference between all three conditions.

Discussion

Recently there has been an influx of studies investigating how performance may differ in Virtual Reality versus Reality. For example, the important work by Thompson et al. (2004) has shown that humans consistently underestimate distances in VR environments that use photo-realistic renderings of real environments relative to their real environment counterparts. Work by Creem-Regehr, Willemsen, Gooch, and Thompson (2005) attempt to localize the cause of this underestimation by investigating whether the restricted field of view, monocular versus binocular viewing, or the inability to view one's feet might account for the compressed distance estimations in the immersive VR environments. However, none of these factors seem to explain the compression in the distance estimates found by Thompson et al. (2004)

Chance et al. (1998) showed that in a dead-reckoning task, human performance declined significantly when the vestibular cues were removed (i.e., rotation information). However, performance was not affected by removing the proprioceptive cues. These results suggest that using Desktop VR where participants navigate through an environment while sitting stationary might have a significant effect on performance.

In the current studies we compare performance in three conditions. The first condition, the Key Press condition, used Desktop VR in which participants could make quantized actions through the environment that were limited to rotate-left by 90°, rotate-right by 90°or move forward 1.2 meters (to the next node in the environment). In this condition, in addition to the Joystick condition which also used Desktop VR, there were no egothetic cues (vestibular or proprioceptive information). However, these cues were available in the Immersive condition. Results showed that there was no significant difference between the conditions that used the Desktop VR (Key Press and Joystick) and the Immersive condition. This result suggests that when navigating with uncertainty in structured indoor environments these cues may not be used, or do not play a significant role.

It should be made clear that these results may be limited to the types of environments that we are using. That is, participants in the current studies navigated in environments that were highly structured — they were indoor environments with clear hallway structure. We believe that the strong visual information that is available in these environments overrides any metrical error that might be generated in any *path integration* system. That is, participants may be estimating the distance travelled both from the optic flow information and from landmark based navigation (i.e., *piloting* to a landmark). There may be a great deal of uncertainty generated from the path integration system. However, the noise associated with the piloting system is most likely very small. Given this small variance, the navigation system might take a weighted estimate or it may simply ignore the path integration signal all together and base its estimates on the information from the landmark based system. In a related study by Foo, Warren, Duchona, and Tarr (2005) in which the authors pitted landmark knowledge against metric knowledge they found that participants relied more on landmark knowledge when landmark information was reliable. However, when the landmark knowledge turned out to be unreliable, the participants were able to use their internal metric estimations to complete the task.

The Use of Topology

The work by Foo et al. (2005) and the work here suggest that when navigating in large environments, humans may use a more "quantized" or topological representation of space. Recently, the work by Kuipers and his colleagues (B. Kuipers, Modayil, Beeson, MacMahon, & Savelli, 2004; B. J. Kuipers, Tecuci, & Stankiewicz, 2003; B. Kuipers, 2001, 2000; B. J. Kuipers, 1998) argues for a hierarchical representation of large-scale spaces with three levels of representation: Control, Causal and Topological. At the very top level of this theory is the *Topological* space in which *places* (e.g., rooms, buildings, etc.) are represented relative to *paths* (e.g., halls, roads, rivers, etc.). This space is an abstracted representation that does not make explicit the metric properties of a large-scale space, but instead makes explicit the categorical or relational structure of the space. The Causal Level provides an explicit representation that states if the observer is in a given state S and generates a specific action A, they will end up in a new state S'. Finally, at the lowest level is the Control Level that makes explicit the specific micro actions that convert the Actions from the Causal level into specific physical actions.

We hypothesize that when participants are navigating in an environment in which there is a clear topological structure, that the participants make use of this information to easily orient and navigate through the space. We suspect that if the space were less constrained, as in open-field navigation, that the results that we found in the current studies might change significantly. However, with that said, it seems that studying human wayfinding and orienting behavior using desktop VR works just as well as using fully immersive environments.

A Question of Task

There is no question that VR provides a rich system for testing and understanding human spatial navigation abilities. Using desktop VR one can study human performance in very large-scale spaces and manipulate aspects of the environment that would be difficult to do in real environments (e.g., starting participants from random locations in a large scale space). However, as can be seen by the different results, VR may be useful for understanding behaviors in some tasks but not others. It seems from the work of Thompson et al. (2004) that VR may not be a good medium for studying distance perception. Furthermore, the work by Chance et al. (1998) suggests that if a task relies on dead reckoning, one should be wary of using desktop VR. However, for structured large-scale environments, the results from the current studies suggest that desktop VR is just as good as Immersive VR.

Summary & Conclusions

The current studies set out to understand whether the reduction in efficiency in the Stankiewicz et al. (2006) generalize to conditions where participants could move continuously and where proprioceptive information was available. To do this, we tested participants in three VR environments where they moved through the environment by making keypresses, using a joystick or using an immersive environment. Using the ideal navigator algorithm we computed the efficiency of human navigation in these three conditions and found that there was no difference in performance in these three conditions. Furthermore, analysis of how many sessions it took participants to learn the environments to criterion also showed no reliable difference between the three conditions. These results suggest that in structured environments, desktop virtual environments can be used to better understand human wayfinding behavior.

References

Braje, W. L., Tjan, B. S., & Legge, G. E. (1995). Human efficiency for recognizing and detecting low-pass filtered objects. Vision Research, 35(21), 2955-2966.

Cassandra, A. R. (1998). Exact and approximate algorithms for partially observable markov decision processes. Unpublished doctoral dissertation, Department of Computer Science, Brown University.

Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting optimally in partially observable stochastic domains. In Proceedings of the twelfth national conference on artificial intelligence (AAAI-94) (Vol. 2, pp. 1023–1028). Seattle, Washington, USA: AAAI Press/MIT Press.

Chance, S. S., Gaunet, F., Beall, A. C., & Loomis, J. M. (1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. Presence: Teleoperators and Virtual Environments, 7, 168–178.

Creem-Regehr, S. H., Willemsen, P., Gooch, A. A., & Thompson, W. B. (2005, February). The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments. Perception, 34(2).

Darken, R. P., & Banker, W. P. (1998). Navigating in natural environments: A virtual environment training transfer study. In Proceedings of vrais 98 (pp. 12 – 19).

Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. Journal of Experimental Psychology: Human Perception and Performance, 29, 343–362.

Foo, P., Warren, W. H., Duchona, A., & Tarr, M. J. (2005). Do humans integrate routes into a cognitive map? map- versus landmark-based navigation of novel shortcuts. Journal of Experimental Psychology: Learning, Memory, and Cognition, 31(2), 195-215.

Franz, M., Schölkopf, B., Mallot, H., & Bülthoff, 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(4), 445-463.

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

Kaelbling, L. P., Cassandra, A. R., & Kurien, J. A. (1996). Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In Proceedings of ieee/rsj international conference on intelligent robots and systems.

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

Koh, G., Wiegand, T. von, Garnett, R., Durlach, N., & Shinn-Cunningham, B. (1999). Use of virtual environments for acquiring configurational knowledge about specific real-world spaces: I. preliminary experiment. Presence: Teleoperators and Virtual Environments, 8(6), 632-656.

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. 10.1–10.7). Cambridge, Massachusetts: University of Michigan.

Kuipers, B., Modayil, J., Beeson, P., MacMahon, M., & Savelli, F. (2004). Map-building with the hybrid spatial semantic hierarchy. In Ieee int. conf. on robotics & automation (icra-04) (p. xx-yy). (Submitted)

Kuipers, B. J. (1998). A hierarchy of qualitative representations for space. In C. Freksa, C. Habel, & K. F. Wender (Eds.), Spatial cognition: An interdisciplinary approach to representing and processing spatial knowledge (Vol. 1404, pp. 337–350). Berlin: Springer.

Kuipers, B. J., Tecuci, D., & Stankiewicz, B. J. (2003). The skeleton in the cognitive map: A computational and empirical exploration. Journal of Environment and Behavior, 35, 81–106.

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

Legge, G., Klitz, T., & Tjan, B. (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(4), 549-568.

Mallot, H. A., & Gillner, S. (2000). Route navigating without place recognition: What is recognised in recognition-triggered responses? Perception, 29(1), 43-55.

Ruddle, R., Payne, S., & Jones, D. (1997). Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. Journal of Experimental Psychology: Applied, 3(2), 143-159.

Ruddle, R. A., & Jones, D. M. (2001). Movement in cluttered virtual environments. Presence: Teleoperators and Virtual Environments, 10, 511-524.

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

Sondik, E. (1971). The optimal control of partially observable markov decision processes. Ph.d. thesis, Stanford University.

Stankiewicz, B. J., Legge, G. E., Mansfield, J. S., & Schlicht, E. J. (2006). Lost in virtual space: Studies in human and ideal spatial navigation. Journal of Experimental Psychology: Human Perception & Performance, 32(3), 688–704.

Steck, S. D., & Mallot, H. A. (2000). The role of global and local landmarks in virtual environment navigation. Presence. Teleoperators and Virtual Environments, 9, 69–83.

Thompson, W., Willemsen, P., Gooch, A., Creem-Regehr, S., Loomis, J., , et al. (2004). Does the quality of the computer graphics matter when judging distances in visually immersive environments? Presence: Teleoperators and Virtual Environments, 4(13).

Thorndyke, P., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. Cognitive Psychology, 14, 560-589.

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(21), 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. (1948). Cognitive maps in rats and men. The Psychological Review, 55(4), 189–208.

Trommershäuser, J., Maloney, L., & Landy, M. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. Journal of the Optical Society of America A, 20, 1419–1433.

Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models. In A. Frank & Campari (Eds.), Spatial information theory: A theoretical basis for gis, cosit '93 (Vol. 716, p. 14-24). Berlin: Springer-Verlag.

Waller, D., Hunt, E., & Knapp, D. (1998). transfer of spatial knowledge in virtual environment training. Presence: Teleoperators and Virtual Environments, 7(2), 129–143.

Witmer, B. G., Bailey, J. H., & Kerr, B. W. (1996). Virtual spaces and real world places: Transfer of route knowledge. International Journal of Human-Computer Studies(45), 413–428.