Development and Evaluation Of A Bayesian Low-Vision Navigation Aid

Brian J. Stankiewicz†, Anthony R. Cassandra,
Matthew R. McCabe and William Weathers
University of Texas at Austin
Department of Psychology
†Center for Perceptual Systems
†Consortium for Cognition and Computation

Abstract

Wayfinding in unfamiliar environments can pose a challenge to anyone, but can be especially challenging to someone who has some sort of visual loss. In this paper we describe an **indoor** navigation aid that uses Bayesian statistics to localize and guide an individual from an unspecified location within a building to a specific destination. We also present three studies investigating the efficacy of this system as a low-vision navigation aid. Two studies were conducted in virtual indoor buildings using desktop virtual reality and one study was conducted in a real building. All three studies investigated navigation performance with versus without the navigation aid. In all three studies subjects traveled a shorter distance with the navigation aid than without it. In the virtual reality studies, the navigation aid actually improved performance over navigating with normal vision.

Index Terms

Human Factors,Navigation,Handicapped aids,Belief maintenance

I. Introduction

Approximately eight-million Americans report having some form of visual disorder defined by difficulty reading print material in a typical newspaper with corrective lenses [1]. These visual disorders can range from being completely blind to having some sort of residual vision. We will refer to these visual disorders simply as *low vision*.

Common causes of low vision include, but are not limited to, macular degeneration, glaucoma, and diabetic retinopathy. Each of these forms of low vision affect an individual's vision in different ways. For example, someone who has age-related macular degeneration may loose vision in the central three-degrees of their visual field with the remaining portion of their visual field remaining intact. By contrast, individuals with glaucoma may have residual vision over their entire visual field, but the visual information is significantly reduced relative to someone with normal vision.

Although the specifics of the vision loss may vary significantly from one visual disorder to another (e.g., loss of contrast, acuity or loss of portions of the visual field) and the residual vision for a specific individual with the same **type** of visual disorder may vary significantly, it is commonly recognized that significant vision loss typically has a detrimental effect on an individual's ability to navigate independently [2].

Marston et al. [2] surveyed 30 individuals with low vision to get an idea of the types of activities that they did each week and whether there were activities that they wanted to participate in but did not engage in because of navigation difficulties associated with their visual disorder. They found that individual with low vision did not participate in 31% of the activities that they would have liked to engage in due to their visual disorder.

This research was supported by a grant from the National Institute of Health to BJS (EY016089), the Air Force Office of Scientific Research (FA9550-04-1) a University of Texas Undergraduate Research grant to WW and a University of Texas Undergraduate Research grant to MRM. Correspondence concerning this research should be sent to Brian J. Stankiewicz at bstankiemail.utexas.edu.

The current paper presents a description of a new low-vision navigation aid that is currently being developed. We also present three studies that test the concept of this low-vision navigation aid in both virtual and real environments. The findings suggest that the **NavAid** system can significantly reduce the distance traveled by a low-vision from an unspecified location within a large-scale building to a specific room or location within the building.

1) Developing a low-vision navigation aid: Developing an effective indoor low-vision navigation aid has the potential of having a significant impact on our society, especially the elderly, given the large population of individuals with low vision, and the fact that the population continues to grow due to age-related visual disorders. However, in developing a low-vision navigation aid, one needs to recognize that there are two fundamental consumers of the product. First, there is the end-user – the low-vision navigator who needs help with navigating through novel indoor environments. However, an indoor navigation aid will typically be worthless without the cooperation of the building owner. A successful system must be inexpensive both for the building owner and for the end-user. From the perspective of the building owner, there is the initial installation expense, but more importantly there is the maintenance expenses of the system. A system that requires constant maintenance may be undesirable to a building owner and if buildings do not maintain the system, then it eventually becomes useless to the end-user. An ideal system requires minimal or non-existent maintenance cost, inexpensive installation costs, is inexpensive for the end user, is easy to use and most importantly is reliable.

2) Current System: The current manuscript describes a low-vision navigation aid that uses Bayesian statistics to localize and guide an individual with low-vision in an unfamiliar indoor environment to a specific destination (e.g., a specific room, an elevator or a bathroom). The components of the system are quite simple and include a range-finder and a computer (e.g., a laptop, palm top or even a Braille computer all would suffice). A map of the environment and a set of pre-compiled files are used to determine the optimal action for the user given their degree of uncertainty. Later, we will provide a more elaborate description of the **NavAid** algorithm and concept. We will also present results from three studies investigating the efficacy of this system for low-vision navigation. Two of the studies (Experiments 1A and 1B) were carried out in virtual reality (VR) and one was carried out in a real building (Experiment 2).

A. Spatial Navigation

Fundamentally spatial navigation can be subdivided into two different processes: *Obstacle avoidance* and *wayfinding*. Obstacle avoidance is the process of maneuvering around objects within the navigator's immediate vicinity. Although obstacle avoidance is a challenge for someone with impaired vision, there are a number of effective obstacle avoidance systems. These systems include, but are not limited to, long canes, guide dogs and sonar devices. Although obstacle avoidance is an important aspect of navigation, Pelli [3] investigated the effect of reduced vision on obstacle avoidance behavior and found that even with significant reduction of visual information subjects were able to maneuver around objects with little difficulty.

The second component of spatial navigation, *wayfinding*, can pose a significant challenge for individuals with low vision. Wayfinding refers to the process of moving from one location in a large-scale space[1] to a specific destination (e.g., traveling from your office to the mail room). Although wayfinding can be a challenge for someone with impaired vision, under some conditions wayfinding does not pose as much of a challenge. Specifically, individuals with low vision are able to wayfind effectively through *familiar* environments. Familiar environments typically include, but are not limited to, their home, the apartment complex in which they live, the office building that they work in and even highly traveled city streets. This anecdotal evidence suggests that individuals are capable of building a useful *cognitive map* [4][2] of

[1]A large-scale space is an environment in which some positions are not observable from all of the other positions. For example a city or a typical office or university building are all large-scale spaces. A gymnasium is an example of a small-scale space.

[2]A Cognitive Map here is simply referring to an internal representation of an external environment.

the environments that they typically navigate through and that this mental representation allows them to navigate effectively. Therefore one challenge faced by individuals with low vision is navigating through *unfamiliar* spaces.

The challenge associated with navigating in unfamiliar environments with low vision is two-fold. First, much of the world is marked with signs that are typically read visually (e.g., street signs, door signs etc.). Normal-sighted individuals can use these signs not only to identify their goal, but to determine if they are heading in the correct direction by observing whether the room numbers are approaching the value of the desired room. Second, vision also allows one to easily and rapidly gather information about the distant structure of the environment. For example, is there an ‘outlet’ (e.g., intersecting hallway) at the end of this hallway; how long is the hallway, etc. Information about the structure of the environment may aid in the search strategy employed by the individuals when searching for a specific room in a building.

The current studies evaluate the effectiveness of a navigation aid for individuals with low vision who are wayfinding in unfamiliar, large-scale spaces.

B. Survey of Existing Low Vision Navigation Aids

Many research scientists and engineers have recognized the challenges faced by low-vision navigators. Recently, with the public use of global positioning system (G.P.S.) signals, researchers have developed navigation aids for outdoor navigation [5], [6], [7], [8]. These systems can be used to localize the user within a large-city, provide instructions to reach a specific destination in addition to providing information about local points of interests (e.g., “the Museum of Modern Art is to your left”). This system, which has been shown to be effective, has generated a great deal of excitement. However, the utility of these GPS-Based systems end at the front door of most buildings because GPS signals are typically unavailable indoors.

Many contemporary buildings are fitted with Braille signs that identify room numbers and other significant landmarks. However, the utility of these signs is marginal given that they must be localized (a difficult task without vision[3]) and studies have shown that many individuals with low vision either do not use Braille or cannot read Braille [9], [10]. This suggests that these Braille signs are not only awkward, but they may be ineffective for a large portion of the population that they are intended to help.

A low-vision, indoor-navigation aid has been developed by [11] which has been developed into the Talking Signs®system [12], [13]. This system uses beacons that transmit a modulated infrared light signal that generates a continuous signal. This beacon can be placed at an indoor destination (e.g., a door or an elevator) outdoor destination (e.g., a bus stop) or even on a moving bus [14]. The user carries a small hand-held receiver that when pointed in the general direction of the transmitter, translates the modulated infrared light signal into speech output that can be used for navigation. Bentzen and Mitchell [15] has shown that this system can be an effective system for navigating through large, unfamiliar environments.

Another low-vision navigation aid is the Verbal Landmarks®system. This system also uses a series of beacons similar to the Talking Signs®system. Verbal Landmarks®uses an inductive loop system that is activated when a portable receiver is within range (approximately 5-feet). When activated, a verbal message can be heard. Unlike the Talking Signs®, the Verbal Landmarks®are non-directional. Therefore, the auditory messages for the Verbal Landmarks®are usually different than those used for the Talking Signs®. The Talking Signs®typically announce the room number or landmark (“Bus Stop”) and allow the user to use this information to navigate. By contrast, the Verbal Landmarks®will give instructions to specific goal states [15].

Bentzen and Mitchell [15] investigated the efficacy of Talking Signs®versus Verbal Landmarks®at the Annual Convention for the American Council of the Blind held in a hotel conference center. In these studies, participants were given a collection of routes to follow. [15] found a significant reduction in both

[3]It should be pointed out that The American Disabilities Act Accessibility Guidelines (ADAAG) code 4.30.6 for Buildings and Facilities requires that all Braille signs be placed 60 inches from the finished floor on the side of the door that opens. This removes some of the localization problem.

distance traveled and the time to complete the routes for participants that used the Talking Signs®over those that used Verbal Landmarks®.

[15] showed empirically that Talking Signs®are a more effective low vision orientation aid than the Verbal Landmarks®system. Although the Talking Signs®system has shown its efficacy empirically, it has not been accepted broadly as a solution to the low-vision way finding problem. One reason for the lack of adoption may lie in that the system has a significant installation cost.

In order for a low-vision navigation aid to be accepted by the blind and low-vision community, it must be available in most buildings. That is, the technology needs to be ubiquitous. In order for a system to become ubiquitous, it must be inexpensive to install and inexpensive to maintain. The navigation aid described here requires little to no maintenance and the cost for installation is anticipated to be minimal since the system requires no physical apparatus within the buildings.

C. Current Research

In this paper we present a description of a low-vision navigation aid (**NavAid**) and two studies investigating the efficacy of the system. Study 1 consists of two experiments conducted in virtual environments. The experiments conducted in virtual environments provided us with a way to develop and test the navigation algorithm without having to build a physical system. Furthermore, the virtual environments allowed rapid modification of large-scale spaces (such as buildings), which would have been difficult to do in real environments.

Study 2 consists of a single experiment in a real environment. In this study, we examined human navigation with normally sighted subjects with degraded visual input. We used our first physical prototype of the NavAid system for navigating through a complex campus building

D. Description of the NavAid Navigation Aid

The *NavAid* system is a navigation device intended to assist people with low vision when navigating from an unspecified location within a building to a destination specified by the user. The current implementation is not designed to help with obstacle avoidance but, given that it contains range-finder technology, the device could be used for obstacle avoidance in the future. Although the system can be used under a variety of situations, we suspect that the NavAid system will be most useful under the following conditions:

1) Indoor environments with a complex hallway structure.
2) Environments that are unfamiliar to the user (i.e., the user does not have a cognitive map of the environment).
3) Room identification is not readily available to the user.

We believe that the primary utility of this navigation aid will be when an individual with low-vision has to navigate in a relatively complex (e.g., office building or medical building), unfamiliar, large-scale, indoor environment. The core contribution of the NavAid system lies in the algorithm used to localize and guide the user to another location in the building specified by the user.

NavAid uses a Bayesian algorithm that, given specific possible actions, costs for making those actions and a destination, can choose the optimal action despite uncertainty about the true state (position and orientation) of the user. This form of Bayesian statistics is referred to as Partially Observable Markov Decision Processes (POMDP) [16], [17], [18], [19], [20]. Previously, the POMDP algorithm has been used to guide robots in large-scale environments [16], [17], [18] and for measuring human efficiency when navigating with uncertainty [19]. The current system is not a robot, but uses some of the underlying mathematics to guide and instruct a low-vision user to a specific location within a large-scale space.

NavAid is comprised of four primary components: an *environment model*, a *range finder*, a *POMDP algorithm*, and a *computer*. The environment model makes explicit the set of states in the environment (positions and orientations), the expected observations for each state within the environment and the resulting states for each action from each state. The set of states in the current system are a series

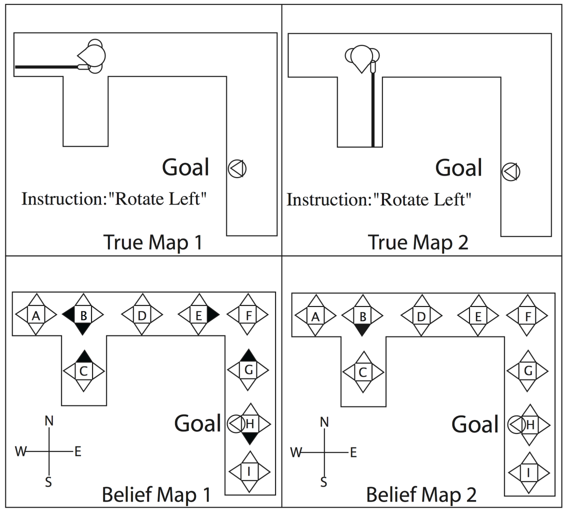

Fig. 1. A simple example of how the low-vision navigation concept works. The upper panel illustrates the *true* state of the world, and the lower panel illustrates where the algorithm believes that the user is located. There are 36 states that the observer could be in (9 positions [A-I] x 4 Orientations [N,S,E,W]. The circled triangles in the lower two panels indicate the goal state. The filled triangles indicate the model's current belief about the user's state. In the upper-left panel the user makes his first observation, given that observation the model generates a belief vector specifying the likelihood that the user is in each of the 36 different states (lower-left panel). Given the belief, the model computes that rotate left is the optimal action to do at this time. The user then carries out this action and then takes a second measurement (upper-right panel). Given this measurement in conjunction with the previous measurement and action the model generates a new belief vector (lower-right panel).

of locations within the environment crossed with four orientations at each location (see Figure 1). The expected observations from each state is simply a distance measurement from the user's current pose to the end of the hall (see Figure 1).

In addition to the environment model, there is also a range finder that is held by the user to take measurements from the user's current position and orientation within the environment to the end of the hall (see Figure 1). The distance measurements are the *observations* that are given to the POMDP algorithm to inform the system about the user's current state. Given a measurement, the algorithm references the environment model and its current **belief** about where it is in the environment. Thus, given an observation, the model can specify a likelihood that the observer is in any one of the states ($p(state|observation, belief)$); see Figure 1 upper-left and lower-left illustration). Given the likelihood of being in each state, or a *belief vector* (b), the model computes the action that will get the user to the goal state with the largest reward or minimal cost. The NavAid system instructs the user about the optimal action to take. Some actions may include, rotate-left by 90°, turn around, go to the end of the hall. After the user completes the instructed action, the user takes another observation. Given the new observation the model updates its belief about its state within the environment considering the initial observation, the action and the new observation (i.e., $p(s|a, b, o)$). The NavAid system continues through this cycle of observation, belief updating, and action selection until the optimal action is to declare "finished". The "finished" action is given when the model believes that the subject is at the goal state with enough confidence (i.e., likelihood) that outweighs the cost associated for declaring at the wrong state.

To further illustrate the basic concept proposed here, we will reference Figure 1. The upper-left and upper-right figures in Figure 1 illustrate the true-state of the user. In the lower-left and lower-right figures are the corresponding state spaces for the environment. That is, there are nine different positions that the environment model considers (A-I) and four different orientations (N,S,E,W)[4]. This gives us 36 different states that the model will consider. In the upper-left panel, the user is making his initial observation. The

[4]We use N,S,E,W as arbitrary directional units. We could have used up, down, left and right also, or any other notation so long as the orientations are consistent.

vector from the user to the wall represents the range finder measurement from the user to the wall. Given this initial measurement, the model computes the likelihood that the user is in each state. We represent the non-zero likelihoods as solid triangles and the zero likelihoods as open triangles. With the initial observation their are six different states that the user could be in: B-W, B-S, C-N, E-E, G-N, and H-S. The goal state for the system is H-W.

To simplify the illustration, we assume that there is no action or observation noise[5]. Because of this the likelihood for each non-zero states is 0.1666 (1.0/6.0). Given this belief vector, the model computes the optimal action to reach the goal state with the minimum cost or maximum reward. In this case, the optimal action is "Rotate-Left". This instruction is given to the user and the user then carries out the instruction. The resulting state is illustrated in the upper-right panel of Figure 1. The user then takes a second measurement from his new state. Following the second observation, the model updates its belief vector taking into account the previous belief vector (b), the instructed action (a) ('Rotate-Left') and the new observation (o) using Bayes' Theorem ($p(s|a, b, o)$ where s is a specific *state*) . The updated belief vector is illustrated in the lower-right panel of Figure 1. At this point, the model has no uncertainty about where the user is located. There is only one state in which this sequence of actions and observations could be made which is state B-S. The model continues with this cycle of observation, action selection, belief updating until the model was confident that the user had reached the goal state.

The previous illustration provides an intuitive understanding of how the algorithm and system works. However, it does not provide a formal description of the system. In the following section, we provide a formal mathematical description of the algorithm. We believe that the reader can skip the following section and still appreciate the fundamental concepts behind the NavAid system.

E. Formal Description of the POMDP Algorithm

The previous example illustrates an intuitive understanding of how the POMDP algorithm actually works. In the current section we will provide a more elaborate description.

The POMDP relies on defining the set of *states*, the expected *observations* from those states, the *action transition matrix* and the *reward structure*. Below, we describe each in some detail and show how these definitions are used to compute the optimal action given a specific belief about the user's current state.

1) Defining the state space: The ideal observer approach relies on converting a large continuous space into a quantized space. This is illustrated in the lower two panels of Figure 1. There are nine different positions that the user could be in crossed with four different orientations at each position to give a total of 36 states within the environment.

2) Defining the observation space: For each of these 36 states we can define the *observation* that is expected from that state. In the current implementation, the observation is a simple distance measurement from the user's state to the wall. However, in previous instantiations of the model [19] the observation was the actual view of the hallway structure from a given state in the environment.

Because the algorithm requires a quantized state space, the expected observations will also be quantized. However, the measurements from the real environment come in continuous values. In the current instantiation of the algorithm we simply round the actual measurement to the closest distance in the set of possible distance measurements within the environment.

3) Defining the action transition matrix: Spatial navigation is a dynamic process in which the actions can change the current state of the user. That is, translating or rotating will move the user from one state in the environment to another state within the environment. To properly define the environment model, we need to define the *action transition matrix*. The transition matrix makes explicit the probability of transitioning from one state (s) to another state (s') given a specific action (a). Or more formally, $p(s'|s, a)$.

This general form of the transition matrix allows us to model actions that are deterministic and probabilistic very easily. That is, for a given action and a given state the model could predict the resulting state with a probability of 1.0. However, often there might be error in the action. For example, in Study 2,

[5]When we actually model the environment we model it assuming that the actions are not deterministic.

we used probabilistic actions such that there was a probability that when the model instructed the user to rotate-left by 90°, the user might actually rotate-left, but, with a small probability, he might rotate right by 90°, stay in the current state or even turn around. Using probabilistic actions allows the system to be more robust to human error. We will discuss the role of probabilistic actions in more detail in the Introduction and Discussion of Study 2. Others have also used the probabilistic actions in POMDPs to model noisy actions [18], [17], [16], [21], [22], [20].

F. *Belief updating*

Once the state space, observation space and transition matrix are specified one can formalize the optimal method for updating the likelihoods for a given belief state – or *belief updating*. This belief updating will be dependent on the previous belief, the current observation and the action just generated. The equation for optimal belief updating is derived using Bayes' Theorem and is given in Equation 1.

$$\begin{aligned} SE_{s'}(b,a,o) &= p(s'|b,o,a) \\ &= \frac{p(o|s',b,a)p(s'|b,a))}{p(o|b,a)} \end{aligned}$$

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 (b), a specific action (a). The likelihood of transitioning into a new state is dependent upon the probability of a resulting observation (o)

$$\tau(b,a,b') = \sum_{o \in O|SE(b,o,a)=b'} p(o|a,b) \tag{1}$$

G. *Choosing the optimal action*

Using the Transition matrix, Reward Structure and Observation Matrix we can compute the optimal action for any given belief state that the observer might have[6]. To compute the optimal action one must consider the *immediate reward* for generating a specific action in a specific state. The following equation is the reward for generating a single action (a) given a particular belief vector (b).

$$\rho(b,a) = \sum_{s \in S} R(s,a)b(s) \tag{2}$$

ρ provides the *immediate* reward for making the action (a) and given the current belief (b). For choosing the optimal action **now** one must not only consider the immediate reward, but the resulting states and the rewards that one might receive when they are in those states. That is, the optimal observer needs to consider not only the current state, but future states too.

Computing the expected reward (ER) when one is considering future actions is an iterative process – a process that must 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:

$$ER = \rho(b) + \sum_{t=1}^{\infty} r(t) \tag{3}$$

where $r(t)$ is the reward received at time t and is based on the expected belief (b) at that time given the actions and observations.

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

To compute the optimal action for a given belief vector, one needs to consider all of the actions that are available to the observer in the current state (A) the immediate reward for generating an action (ρ from Equation 2) and the transition function that specifies the likelihood that a new belief state will be generated (τ). By choosing the action that maximizes the following function one acts optimally in the environment.

$$V(b) = \max_{a \in A} \left[\rho(b, a) + \sum_{b' \in B} \tau(b, a, b') V(b') \right] \tag{4}$$

H. Using the Model as a Navigation Aid

One needs a great deal of computation resources to calculate Equation 4. Because this function is so computationally expensive to compute one might think that it would be difficult to use this function in a real-time, low-vision navigation aid. If we were to compute this on-line, the user might remain standing in the middle of a hallway waiting for the computations to be completed. However, we can **pre-compute** this function for a specific layout to give us the optimal action for any belief vector that the user might find. Given that this function is pre-computed off-line, the remaining on-line functions are simple.

II. Study 1

Many of the ideas for the low-vision navigation system were born out of previous work investigating human navigation efficiency when navigating in large-indoor environments [19]. In working with an ideal navigator that used a similar navigation algorithm, it became evident that with limited visual information the navigator could quickly and accurately localize itself within a large-scale environment (Experiment 2 of [19]).

Experiment 1A investigates the utility of the NavAid system in a large-scale space starting from an unspecified location. The observations in Experiment 1 are the distances from the observer's position to the nearest wall in front of them. Study 1 investigated human navigation performance in two viewing conditions: *Normal Vision* and *Degraded Vision*. In the Degraded Vision condition participants navigated with the navigation Aid (Degraded Vision plus NavAid) and without the navigation aid (Degraded Vision). In Experiment 1A participants navigated through the environment by making quantized actions with key presses. In Experiment 1B, participants moved through the environment in a continuous manner using a joystick. In both experiments participants viewed the environment from a first-person perspective.

The primary goal of Study 1 was to evaluate and develop the underlying NavAid algorithm without first implementing a physical system. The results from Study 1 encouraged us to conduct Experiment 2 which tested the NavAid algorithm in a real, large-scale environment.

A. General Methods

B. Materials

Three virtual reality indoor environments were used. These environments were randomly generated on a Cartesian grid and consisted of forty hallway units. Within all of the buildings "signs" were placed on the wall that had numbers (see Figure 3). The numbers in the virtual building were placed in a typical structured format. That is, each sign in the same hallway had the same century value and were linearly ordered. The precise maps and room numbers of the buildings used can be found in Figure 2.

Participants viewed the environment from a first person perspective using a desktop display. Examples of these displays are shown in Figures 3 and 4. Participants navigated through the environment in three viewing/navigating conditions:

1) **Normal Vision** The participant navigated through the building with unimpeded vision. They could see all the way to the end of the hallway and they could see all of the signs on the current hallway.

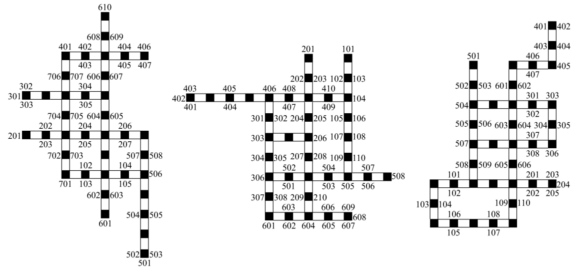

Fig. 2. Maps of the three environments used in Study 1. The numbers represent the sign numbers used in the virtual reality environments and their locations within the environment.

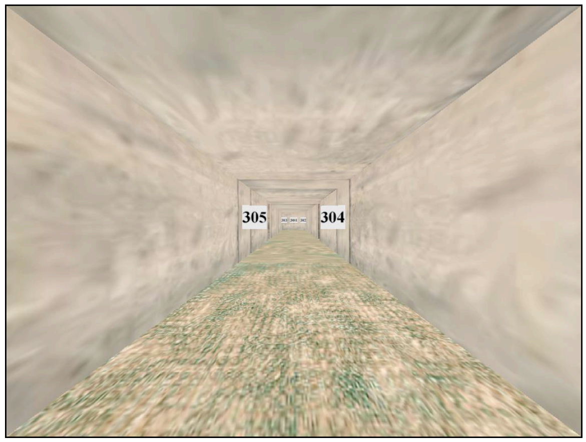

Fig. 3. An example of a view used in the “Normal Vision” condition. In this condition participants could see to the end of the hallway and they could see the signs from a distance.

2) **Degraded Vision** Fog was added to the virtual environment to simulate low vision navigation. The fog was presented such that the observer could not see the next sign, nor could they see the end of the hallway. However, they could perceive whether they were facing a wall, or an open hallway (see Figure 4).
3) **Degraded Vision + NavAid** This condition is identical to the *Degraded Vision* condition with the addition of a set of instructions given by the NavAid algorithm. In this condition the participant was given an auditory instruction to “Move Forward”, “Rotate-Left” or “Rotate-Right”.

1) Apparatus: A Windows computer with a 17-inch monitor, keyboard, mouse, and joystick were used in this study. The virtual environments were modeled using Virtual Reality Modeling Language (VRML) and rendered using Vizard 2.0 software. The three building layouts were created by a computer program that randomly generated floor plans. Forty to seventy room number signs were distributed throughout the buildings in a way to indicate the structure of the buildings. Participants viewed the environments from a first person perspective and navigated using either the keyboard or a joystick. When using the keyboard, the participant navigated using the forward, left, and right arrow keys. Pushing forward moved the view forward one unit, left rotated the view to the left 90 degrees, and right rotated the view 90 degrees to the right. When the participant used the joystick, she was allowed to move freely through the environment by pushing the joystick forward, left, right, and backward.

After the participant navigated to the desired room number, she would hit the space bar to indicate that she had arrived at the goal location. To repeat the announcement of the current goal location, the participant would hit the “G” key. The participant also used a few special functions unique to the NavAid condition. These included taking a reading of the hallway length by pressing the ‘+’ key, pressing the ‘0’ key to indicate that she had executed NavAid’s last instruction, repeating the last NavAid instruction by pressing the ‘M’ key, and resetting the NavAid device by pressing the ‘R’ key.

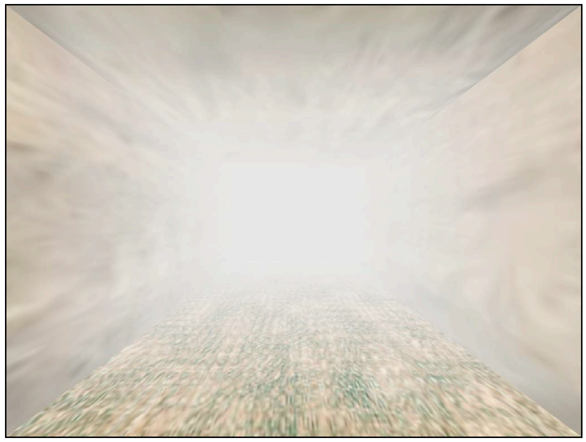

Fig. 4. An example of a view used in the Degraded Vision condition. Participants could not see to the end of the hallway and they could not see the signs from a distance. When participants were at a location where there was a sign they could rotate to see the sign.

C. General Procedure

Before beginning the testing phase of the study, participants ran in a demonstration. The demonstration familiarized the participant with the task, how to navigate, and how to follow the NavAid instructions. This was done in an environment that was smaller than the actual environments used in the study. The participants were given an opportunity to repeat the demonstration until they felt comfortable with the procedure used in the study.

The experiment was conducted using a within participants blocked design. Each participant ran in 50 trials in each of the three conditions: *Normal Vision*, *Degraded Vision*, and *Degraded Vision + NavAid*. The order of conditions was counterbalanced between participants. The start and goal locations for each trial were selected randomly for each layout. The layouts were counterbalanced across the different Viewing/Navigating condition such that each layout occurred in each Viewing/Navigating condition equally often.

At start of each trial, the computer placed the participant at the designated starting state within the environment. The computer then verbally announced the current goal room (“sign number”). The participant then moved to the goal using either key presses (Experiment 1A) or by manipulating a joystick (Experiment 1B). When the participant arrived at the goal state they indicated that they were there by pressing the space bar. After pressing the space bar, the participant was “teleported” to the next starting location. After this process was repeated 50 times (all of the start-goal state combinations), the participant started in the next condition.

In the *Degraded Vision + NavAid* condition, participants moved through the environment in much the same way as the other two conditions (Normal Vision and Degraded Vision). However, in this condition the participant received an auditory instruction from the computer about which action to take. The computer instructions consisted of the following instruction set: “Move Forward”, “Rotate-Left”, “Rotate-Right”[7], “Take a Measurement” and “Finished”.

Participants were informed that the “Move Forward” instruction corresponded to moving to the next intersection, or location where there would be an intersection. The “Rotate-Left” and “Rotate-Right” instructions corresponded to rotating counter-clockwise and clockwise respectively, by 90°. Furthermore, participants were instructed that when the computer instructed them to “Take a Measurement” that they should press the ‘+’ key on the keyboard. They were instructed that this would take a measurement from their current location to the end of the hallway, or the wall in front of them.

The first action that the NavAid system instructed the observer to do was “Take a Measurement”. The participant would then press the ‘+’ key. This would then be followed by NavAid instructing the

[7]Both rotation instructions had the implicit suffix of “by 90°”

participant to make a movement (forward or rotate action). The participant would then carry out this action and when they had completed the action they would press the "0" key. The computer would then announce "Take a Measurement", and followed again by the movement action after the measurement was taken. This process would continue until the algorithm would declare "Finished". This is when the algorithm believed that it was at the goal state and the observer would hit the space bar to indicate that they were finished.

An important aspect of the Degraded Vision + NavAid condition is that the instructed action and the action actually taken were not necessarily the same. That is, the algorithm may instruct the user to "Rotate-Left" but the user might "Rotate-Right" (or any other combination of errors). When these errors occurred, the algorithm would receive an observation that was impossible given the action it thought that the observer generated. So under these conditions the NavAid algorithm would "Reset" itself (start with a uniform belief vector over all of the states) and instruct the user to take another measurement.

The dependent measure in Experiments 1A and 1B was the average distance traveled between the start and goal states. As a result of the counterbalancing in the study, all conditions had the same average minimum distance (the minimum distance required to travel between the start and goal states.)

D. Experiment 1A

Experiment 1A was our initial attempt at using the NavAid system for guiding a user in an unfamiliar environment to a specific goal state. In Experiment 1A participants moved through the environment by making key presses that corresponded to moving forward 1 hallway unit, rotating left by 90°or rotating right by 90°.

1) Participants: Three participants affiliated with the University of Texas at Austin participated in this experiment. All three participants were male and had normal or corrected vision. The participants ranged in age from 22 to 37 years old. Two of the participants are authors on the current manuscript (MRM and BJS). The third participant was recruited from outside the lab. The participant who was not a member of the lab received $10 an hour as compensation for participating in the experiment.

2) Procedure: Experiment 1A: In the current study the participants moved by pressing buttons on the number pad of a typical computer keyboard. The movement actions were designated in the following way: "8" was assigned to forward 1 hallway unit. "6" was assigned to rotate clockwise by 90°and "4" was assigned to rotate-left by 90°. Thus, the state space was effectively discretized into a set of hallways and nodes; in moving through the environment, the participant was only able to move from node to node through the hallway segments. A single rotation (either left or right) or a translation took approximately 1 second.

E. Results: Experiment 1A

Figure 5 shows the mean distance traveled by all participants in the three conditions in Experiment 1A. The mean distance traveled in the Normal Vision condition was 43.63 m (SD = 6.09 m)[8], in the Low Vision condition was 81.38 m (SD = 14.13 m), and in the NavAid condition was 28 m (SD = 2.85 m). Planned paired comparisons were conducted between all three conditions. Because the variances of the conditions differed greatly, all t-tests assumed unequal variances. A two-tailed t-test between the Normal Vision and Degraded Vision conditions revealed a difference that approached significance (t(2) = 3.18, p = 0.09). This result suggests that the Degraded Vision manipulation had a negative effect on the participant's ability to navigate from the start state to the goal state.

A one-tailed t-test between the Degraded Vision condition and the Degraded Vision + NavAid conditions revealed a significant difference ($t(2) = 2.91, p < 0.05$). Finally, we conducted a two-tailed t-test between the Normal-Vision and Degraded Vision + NavAid condition which revealed no significant effect.

[8]The average shortest distance between all of the start and goal states was 22.96 m.

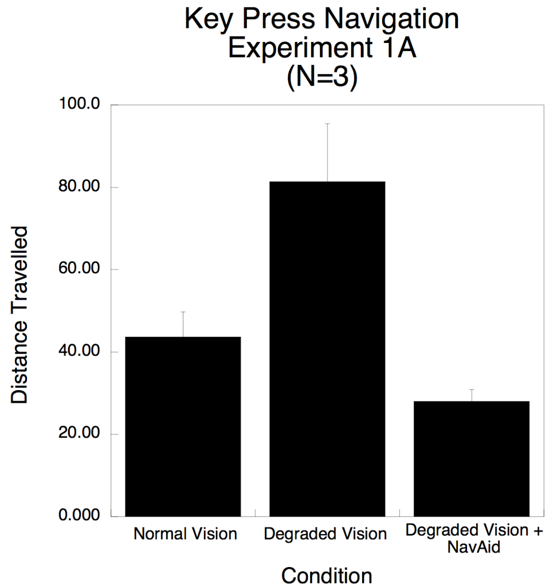

Fig. 5. Average distance traveled to reach the goal state in the three viewing conditions when participants made key presses to move through the environment. The *Degraded vision* and *Degraded Vision + NavAid* used the same viewing conditions, but the *NavAid* condition included the NavAid algorithm to help guide the participant to their goal state.

Although these results are encouraging, the small number of participants makes the results suspect. However, upon examining Figure 5, one can see the robust effect found between the Degraded Vision + NavAid and the Degraded Vision conditions.

1) Discussion: Experiment 1A: This initial study suggested to us that the NavAid algorithm may have some utility for low-vision navigation. While the findings were encouraging, we recognized that there were many limitations in conducting the study using quantized actions. Because participants moved through the environment in a quantized manner, many of the errors that someone with low vision might make would be difficult to generate. To this end we conducted a second study in the virtual reality environment in which participants moved through the environment using continuous actions and a joystick.

One error that will pose a difficulty for users with low-vision is being able to align with the hallway to get an accurate measurement from their current pose to the end of the hall. Because the participants in Experiment 1A could only rotate by 90°, alignment was never an issue. Second, the algorithm instructed the participant to translate forward one hallway unit at a time. This corresponded to a single "forward" key press in Experiment 1A. However, in real environments, users would not be able to simply press a button to move forward a certain distance. With continuous actions the precise distance that a user would travel may vary dramatically.

To address these issues, we ran the same study in Experiment 1B, but instead of moving through the environment with key presses, participants moved through the environment using a joystick. Using the joystick eliminated the quantized set of actions that were available to the participant. Thus, many of the errors mentioned above may rear their ugly head.

F. Experiment 1b

We designed this experiment to use continuous space to test NavAid under more realistic conditions. To take a reading in this experiment, the participant had to align herself with the view of the hallway in order to obtain an accurate measurement with the simulated range finder. Thus, there was a chance that she would make an observation error if she did not aim the NavAid device correctly and took a distance measurement from a nearby wall instead of the wall at the end of the corridor. When this happened, NavAid would not have an accurate measure of the distance of the hallway in front of the participant.

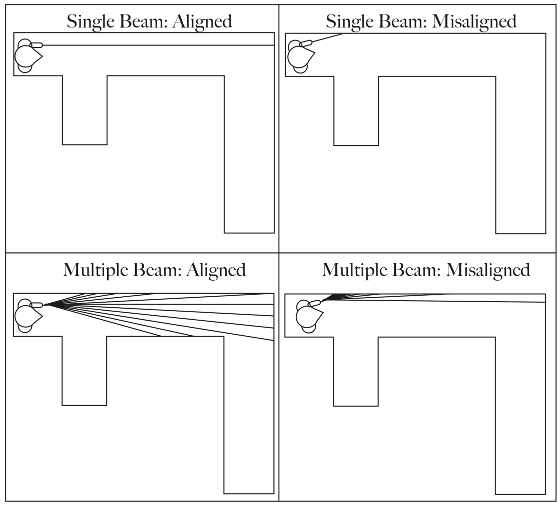

Fig. 6. An illustration showing the utility of the multi-beam system. The upper two panels show how a single-beam may produce the wrong distance measurement if the user is not aligned with the hallway. The lower two panels shows how the multi-beam system would compensate for a certain amount of misalignment.

1) Participants: Six participants affiliated with the University of Texas at Austin participated in this experiment. Four were male and two were female. All of the participants had normal or corrected vision and ranged in age from 21 to 26 years old. Five of the participants were members of the researchers' lab and one was recruited from outside the lab. The participant who was not a member of the lab received $10 an hour as compensation for participating in the experiment. Note that these participants did not participate in Experiment 1A.

2) Apparatus: Experiment 1B: The apparatus was very similar to what was used in Experiment 1A. Two important modifications were made. First, participants navigated using a Logitech Attack 3 joystick. This allowed free movement and rotation of the participant throughout the study. Second, to help compensate for not aligning directly with the hallway in front of them, we used a simulated multi-beam distance finder. The multi-beam range finder returned multiple distances from forty simulated beams (see Figure 6). The beams were distributed across 20 degrees in 0.5 degree increments. The algorithm returned the maximum length distance which was used to update the algorithm's belief vector.

The underlying algorithm for estimating the participant's state within the environment was the same in Experiment 1B as it was in Experiment 1A. That is, as in Experiment 1A, the algorithm quantized the space into four orientations at each node in the environment. Because the participants were freely moving through the environment, it was unlikely that they would stop exactly at the center of a particular node. Therefore, the NavAid algorithm was not given the precise measurement from the multi-beam range finder, but instead was given the measurement as if they were standing at the nearest node. More specifically, we rounded the value to the nearest 4 meter measurement (the distance from one node to another in the environment).

G. Results: Experiment 1B

The mean distance traveled by all participants in each condition in Experiment 1B is shown in Figure 7. The mean distance traveled in the Normal Vision condition was 38.61 m (SD = 8.48 m), in the Degraded Vision condition was 51.07 m (SD = 13.47 m), and in the Degraded Vision + NavAid condition was 28.26 m (SD = 1.49 m). We conducted a t-test assuming unequal variances to compare the three means. There was a significant difference between the Normal Vision and Degraded Vision conditions ($t(5) = -1.92, p \leq 0.05$). A one-tailed t-test to comparing the Degraded Vision condition and the Degraded Vision $+$ NavAid conditions also found a reliable difference($t(5) = 4.12, p \leq 0.05$). We also found that

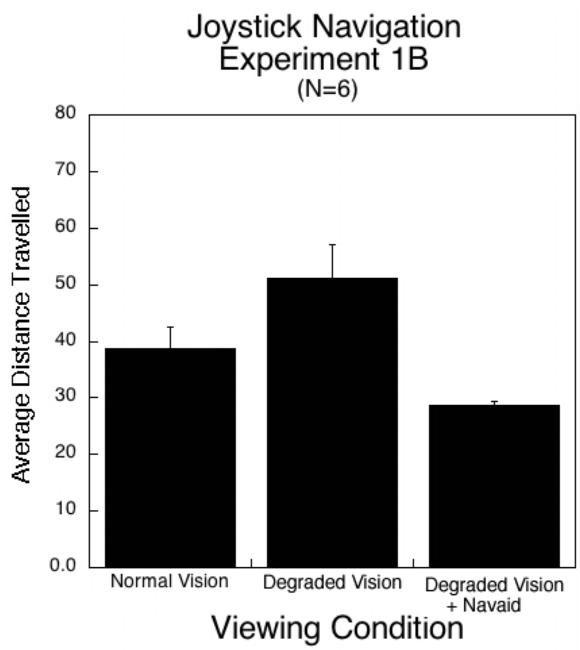

Fig. 7. Average distance traveled (error bars represent one standard-error of the mean) to reach the goal state in the three viewing conditions. The *Degraded Vision* and *Degraded Vision + NavAid* used the same viewing conditions, but the *Degraded Vision + NavAid* condition included the NavAid algorithm to help guide the participant to their goal state.

participants traveled a shorter distance in the Degraded Vision $+$ NavAid condition than they did in the Normal Vision condition ($t(5) = 2.94, p \leq 0.05$).

H. Discussion

In Experiment 1, we hoped to show that our simulated NavAid device would effectively guide participants in a virtual building. We first hypothesized that participants would navigate more efficiently when they were in the Degraded Vision + NavAid condition versus the Degraded Vision condition. We found support for this hypothesis in both experiments; it appears that NavAid helps to guide participants in both discrete and continuous environment spaces.

We also hypothesized that participants would navigate more efficiently in the Normal Vision condition compared to the Degraded Vision condition. In Experiment 1A, we did not find support for this hypothesis. However, the result did approach significance. We believe that the lack of significance was the result of the small number of participants in the first experiment. In Experiment 1B, we did find a significant difference between the Normal Vision and Degraded Vision conditions. Specifically, the participants navigated more efficiently in the Normal Vision condition than in the Degraded Vision condition. This significant difference was most likely the result of the larger number of participants in the second experiment compared to the first experiment.

Given that we found a significant difference between these two conditions in Experiment 1B, it is clear that adding fog to a virtual environment makes the navigation task more difficult. Thus, we believe that the Degraded Vision condition effectively simulates the task of navigating with a degree of reduced vision.

Another interesting finding was the fact that participants navigated more efficiently using the NavAid device compared to the Normal Vision condition in Experiment 1B. Thus, when a fog was added to the environment and the participant used the NavAid device, she performed better than when there was no fog in the environment and she was not able to use NavAid. This result suggests that people with normal vision may benefit from using the NavAid device when navigating in unfamiliar indoor environments.

III. Study 2: Evaluation of NavAid in a Real Environment

Study 1 demonstrated that the NavAid system may have some utility for someone with low vision when navigating in an unfamiliar indoor environment. However, being able to use the system in a simulated indoor environment is much different than using the system in a real environment. There are aspects of the virtual environment that are not available in real environments. Some of these include, but are not limited to:

1) While in the virtual environments there were clear node markers, these markers would not be available in the real environments.
2) Although buildings may have a basic Cartesian structure, this may not always be the case.
3) When taking measurements in the virtual environments there was no vertical movement of the laser-range finder. In real environments, in addition to possibly pointing at the wall beside them, a person may also point at the floor or ceiling.

After conducting Experiment 1B and during initial development of the NavAid system for real environments, we recognized the need to modify the algorithm somewhat. We made three significant changes to the algorithm. The first was that we increased the number of actions that the NavAid system could provide as instructions. These instructions would reduce the number of instructions that would be required for the participant to reach the goal. These instructions included “End of Hall”, “Turn Around”, “Forward 2”, “Forward 3” and “Forward 4”.

The second significant change was the use of noisy actions. In Study 1, the algorithm always assumed that the user correctly carried out the given instruction. Often the participant would make a mistake and the algorithm would need to reset. In developing and testing the algorithm we recognized that, for the algorithm to successfully guide the user, it needed to consider not only the action that was instructed, but a collection of other possible actions that the participant might make. For example, if the NavAid system gave the instruction “Forward 2” the algorithm assumed that the participant most likely moved forward 2 units, but that there was a small possibility that the participant actually moved forward 1 unit or 3 units.

The third change that was made to the algorithm was limiting the longest possible measurement that the algorithm would trust. Again, in developing and testing initial designs we realized that it would be very difficult to accurately point to the end of the hallway when the hallway was greater than 15 meters. This was especially true when the participant was wearing the low-vision simulating goggles. Often when the hallway was greater than 15 meters the participant would take a measurement off of the floor or ceiling far down the hallway. When this happened, the laser range-finder would either return an error (which we entered as greater than 15 meters) or it returned a value greater than 15 meters. If the laser-range finder returned a value, we entered that value into the algorithm.

The goal of Experiment 2 was to evaluate the effectiveness of the low-vision navigation aid in a real environment. In the study we used a building on the University of Texas at Austin campus. We ran the participants in the evening or on the weekends to reduce the traffic within the hallways during the studies.

A. Methods

1) Participants: Six participants were recruited from the University of Texas at Austin. The participants were paid $10 / hour for participating in the study. There were two males and four females ranging in age from 18-23 years of age.

2) Materials: Participant used a hand held, Disto laser-range finder (Pacific Laser Systems PLS1) to obtain distance measurements in the environment. Participants wore goggles that had lenses that were scratched with an abrasive surface. The scratched surface severely limited the participant’s ability to identify signs and see details of the environment. However, it did not prevent them from successfully avoiding obstacles [3].

A Macintosh Titanium laptop was used in the study. The laptop served multiple purposes. First, the NavAid program was loaded on the laptop. The computer ran the NavAid algorithm and generated the auditory instructions for the participants. The computer also randomized the conditions that the participants

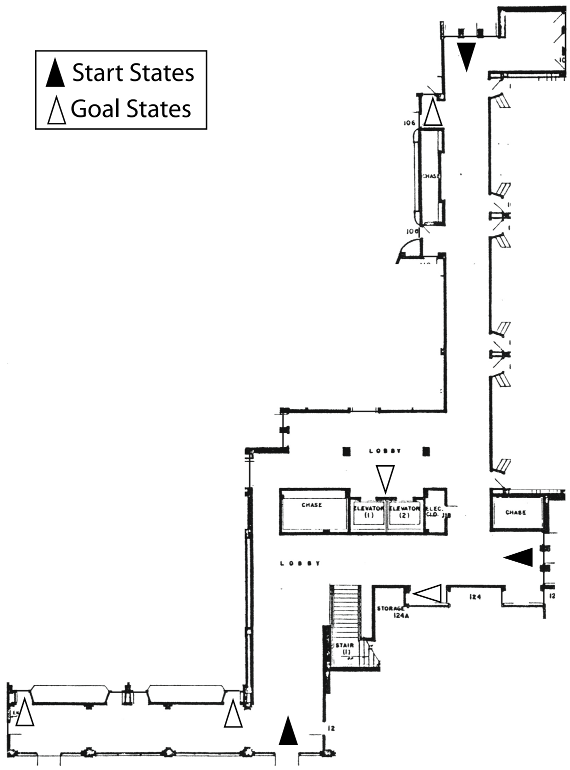

Fig. 8. Map of Burdine Hall which was used in Experiment 2. The filled triangles represent the starting positions and orientations. The open triangles represent the goal states and orientations

ran in during the study (Start State, Goal State, NavAid vs. No NavAid). To prevent learning of the environment and to simulate navigating in an unfamiliar environment, we also randomized the room numbers (in a structured manner, see below). The computer was used to randomize the room numbers in the No-NavAid condition. Finally, the computer was used to collect the participants' performance in the task (time to goal, distance traveled and distance to goal).

To measure the distance traveled to the goal a Trumeter 5605 Mini Measure Maxx rolling measuring device was used. The measuring device was reset at the beginning of each trial and held by the participant to measure how far they traveled during each trial.

The first floor of the Burdine Hall at the University of Texas at Austin was used. A map of the building is shown in Figure 8. Three starting positions and five goal states (four rooms and an elevator) were used. These states are indicated on the map by filled triangles (starting states) and open triangles (goal states). Each participant ran from all three starting positions to each of the five goal states, for 15 trials in each condition (NavAid and No NavAid) for a total of 30 trials.

3) Procedure: Participants engaged in three phases during the study: *Calibration*, *Introduction* and *Testing*.

Calibration. The NavAid system quantizes the environment space into three-meter regions. The translation instructions given by the NavAid system specify how many regions forward that the participant should walk. These distances were converted into number of steps. However, because there can be variability in stride length we first determined how many steps corresponded to a three-meter translation.

Using the range finder, participants took a measurement to the end of the hallway. The participants were then asked to take a specific number of steps (5 steps, 10 steps, 15 steps, and 20 steps), stop, and take another measurement in the direction they were facing to the end of the hallway. The distance the participant traveled was recorded and used to calculate their average stride length. This procedure was done twice for each of the specific number of steps (5, 10, 15, and 20). The average stride length was then used to determine how many steps the participant needed to take to move one unit, or 3 meters.

Because the participants stride length may have changed when they had reduced vision, the calibration was conducted with the low-vision goggles.

Introduction We also familiarized the participant with using the laser-range finder. Pointing the laser-range finder to the end of a hallway can be difficult. Slight variations in the angle can produce large movements of the laser range-finder's position at the end of the hall. This was especially difficult with the goggles because the range finder's laser point was unobservable with the goggles. To familiarize the participant with this process we allowed the participant to try and point the laser range finder to the end of the hall. They then took off their goggles to see where it was pointing. They also walked forward and/or rotated and then tried pointing the laser-range finder.

Participants were also familiarized with the set of actions that the NavAid system would use. Most importantly, they were familiarized with the translation instruction. We told the participant how many steps corresponded to one three-meter movement (typically this was about 4). We then instructed them that they should take that many steps for each forward action that the NavAid system instructed (e.g., "Forward 3" would correspond to twelve steps).

Testing. The participants ran in two conditions: NavAid and No-NavAid. The participant began from one of the three start positions. In the NavAid condition, the participant began by taking a measurement to the end of the hallway or the wall directly in front of them using the range finder. The experimenter read the measurement from the range finder and entered it into the laptop equipped with the NavAid program. Based on this measurement, a verbal command was issued to the participant by the laptop. The verbal commands were: "Rotate-Right", "Rotate-Left", "Forward 1", "Forward 2", "Forward 3", "Forward 4" "End of Hall", "Turn Around" and "Finished". "Rotate" indicated that the participant should rotate 90°in the direction of the command. Each number associated with the "forward" command indicated how many units of movement the participant was to move forward (recall each unit corresponds to three meters of movement). "End of Hall" directed the participant to move all the way to the end of the hallway in the direction they were facing. "Turn Around" directed the participant to rotate 180°. After completing the designated action, the participant took another measurement which was entered into the computer by the experimenter followed by another instruction. This process (Observation, Instruction, Action) continued until the algorithm determined that the participant was at the goal. At this time the algorithm would announce "Finished".

When the computer announced "Finished" the participant was instructed to point to the door that they thought was the goal state. The experimenter entered whether the participant pointed to the correct door or not. The distance from the goal was also measured using the laser range finder and the distance traveled to the goal was also recorded.

In the *No-NavAid* condition, the participants were given a goal room in the form of a room number. The task was to find this room within the building by searching through the building. Because the participants were wearing the optically limiting goggles they were unable to see the signs for each room. To allow for the participant to find the goal room we had participants touch doors within the building and after touching the door the experimenter would verbally tell the participant the room number associated with that door.

Because we ran participants to the same room multiple times we did not want the participant to build a cognitive map of the building that contained the positions of each room. Remember, the intent was to test human way finding when participants were unfamiliar with the environment. Although we could not prevent the participants from learning the topology of the building, we could prevent them from knowing exactly where the goal room was located. To achieve this, each room was assigned a new room number by the computer. These numbers were generated to have the same underlying structure as the actual room numbers. That is, the rooms in each hallway had the same century value and the room numbers also progressed in a linear order as a function of their position in the building (similar to how the actual room numbers were structured in the building). This allowed the participant to make use of any information that might be available from knowing the room numbers of the building (such as linear order) but prevented participants from being able to recognize the room number and move directly to the goal state.

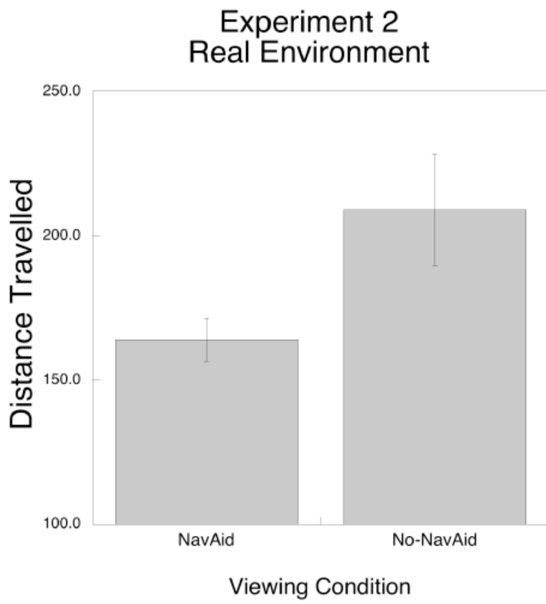

Fig. 9. Distance traveled to reach the goal state when participants used the NavAid system versus when they did not use the NavAid system. Error bars represent one standard-error of the mean.

The room number re-ordering was done by the computer on each trial. Each hallway was assigned a random century value and the room numbers within a given hallway were assigned numbers that were linearly ordered in the same manner as the actual hallway. The only modification that we made was whether the number increased in the same way as the actual hallway, or if the linear ordering was inverted.

In addition to preventing participants from knowing the position of the goal state given only the room number of the goal, this method of finding a door and touching it is similar to someone reading a Braille sign. We recognize that touching a door is much simpler than trying to find a Braille sign and reading it when one has limited vision. However, simply touching a door and receiving an auditory signal of the room number is probably as fast or faster than using Braille signs.

Similar to the *NavAid* condition, the participants held the rolling distance measuring device. Furthermore, the participants indicated when they had reached the goal state. As before, we recorded the distance traveled, whether the participant reached the correct goal or not and the time to reach the goal state.

Four variables were measured: travel time, distance traveled, distance to target, and correct identification of target door. The travel time was measured by the computer. Distance traveled was measured by the rolling distance measuring device and was recorded after each trial. The distance to the target was measured using the laser range finder. When the participant (in the No NavAid condition) or the NavAid program (in the NavAid condition) deemed the trial was over and the participant was at the goal state the laser range finder was used to measure the distance from where they finished to the goal state. This was done by measuring the distance from the participant to the door of the goal state. The correct door determination was assessed by asking the participant to point to the door he thought was the door of the goal state. A ”1” was given for a correct answer and a “0” was given for an incorrect answer.

B. Results

One of the first measurements that we were interested in evaluating was how close the participant got to the goal state when using the NavAid system. The average distance from the goal was 0.16 meters. We also instructed the user to inform us of which door they thought was the goal state. All participants, with the exception of one, identified the correct goal state (door) in all trials in both conditions. One participant misidentified the goal state twice in the NavAid condition.

A significant difference ($t(5) = -2.39; p \leq 0.05$) was found between the two conditions for distance traveled using a repeated measures t-test. The mean for distance traveled in each trial in the NavAid condition was 164.38 ft (SEM=7.48 ft) and in the NoNavAid condition the mean was 209.85 ft (SEM=19.397 ft)[9].

For travel time, a significant difference was found ($t(5) = 7.02; p \leq 0.05$) indicating that in the NavAid condition the participant took longer to get to the goal state than in the No NavAid condition. The mean time for the NavAid condition was 198.8s (SEM=11.37s) and the mean time for the No NavAid condition was 114.05s (SEM=6.74s).

C. Discussion

The current study was interested in evaluating whether a Bayesian low-vision navigation aid might benefit someone navigating through an unfamiliar indoor environment. The results showed a 22% reduction in the distance traveled from the start state to a goal state when the participant was using the NavAid system. This result suggests that the system might have some benefit navigating benefits for individuals with low vision.

However, we also showed that it took longer for the participant to reach the goal state with the navigation aid than without it. We suspect that much of this was due to the fact that the experimenter had to read the distance value off of the laser range-finder and enter the value into the computer. With a unified system in which the range finder is directly connected to a computer, we believe that this time will be reduced.

IV. General Discussion

Navigating in an unfamiliar, complex indoor environment can be a very daunting task. To help a naive navigator in a building, a number of aids have been installed in most buildings. Often times a building will be equipped with a map with a “You Are Here” symbol. This map may contain information about where all of the rooms are within the building in addition to where the elevators, stairwells, and exits are located. Furthermore, in most buildings each room is assigned a room number. These room numbers are almost always posted on visual signs and there is a particular structure to the assignment of room numbers. For example, the millennium value may specify which floor the room is on, while the century mark specifies the hallway and the remaining values specify the specific room. In addition to that, the rooms within a particular hallway typically increase linearly as one progresses down the hallway in one direction.

Even for someone with normal vision, it can be difficult to navigate a complex indoor environment using these aids. The problem becomes even more complex for someone who has low vision. Most of these navigation aids are presented in a visual format and are therefore unaccessible to someone who has low vision or who is blind. Although most buildings are equipped with Braille signs, many people with low vision do not read Braille, and often times it can be difficult to locate these signs within a building.

Because of this many researchers have proposed low-vision navigation aids for indoor navigation. In this manuscript we described a low-vision navigation aid that uses Bayesian statistics to localize and guide the individual to their destination. We tested the system in both virtual reality environments and in real environments using simulated low vision. In all cases the system decreased the distance that it took participants to reach the goal state.

A. Some Remaining Issues

We found the current studies encouraging and they suggest that the NavAid system may be useful for low-vision navigation. However, there are a number of issues that remain to be addressed. First, the current studies were conducted with normally sighted subjects with simulated low vision. Given some of

[9]The average distance between all of the start and goal states was 107 ft.

the properties of the system, namely the need to point the laser-range finder down a hallway, it remains to be seen whether the NavAid system is useful for people with **any** visual disorder, or if it will only be useful for individuals within a specific visual loss range. For example, on the upper end, the system might not benefit someone who has normal vision at all[10]. On the lower end, someone who is completely blind may find it challenging to point the range finder down a hallway. For the low end users, current research and development in the lab is working to overcome this challenge from three perspectives. The first is actually evaluating how well a low-vision user can point a laser range finder down a hallway. If blind users (or severely visually impaired subjects) can accomplish this then there is no challenge. If blind and severely visually disabled subjects are unable to successfully point the range finder down the hallway then we have specific engineering designs and techniques that we believe can overcome this challenge. We are also working to modify the software to take into consideration the types of observation errors that are made by individuals who are blind (or who have very little residual vision). We would like to re-iterate that the POMDP approach is specifically designed to handle noisy observations.

We also have reason to believe that for certain environments the system may not have that much of a benefit. Imagine for a moment a building that consists of a single hallway. Although the system would work, the benefit of using the system over not using the system may be small. A second issue related to the environment layout has to do with symmetry. If the environment is two-fold or four-fold symmetric, the system will never be able to reduce its uncertainty. Future implementations of the system may utilize a compass to give the general heading of the user while translating down a hallway. These compass readings will allow the system to 'break' this symmetry.

In addition to issues with the environments, there are also some remaining issues with the underlying approach. The current implementation of the NavAid system utilizes a laser range-finder to measure the user's position relative to the end of the hall (or nearest wall). Given this, there are two potential issues with this type of system. The first is that there might be *dynamic obstacles* and/or *obstacle clutter*. Dynamic obstacles refer to objects within the environment that might move from one time to another (e.g., pedestrians) and obstacle clutter are objects within the environment that might not be part of the building structure (e.g., a plant or cabinet in a hallway). Each of these obstacles can generate an inaccurate observation (i.e., an observation that is shorter than the expected observation from that state to the end of the hall). To handle this issue, we have been modeling the environments with probabilistic observations in addition to probabilistic actions. That is, for each state the system expects a collection of observations, each with varying probabilities. For example, a state that is twelve meters from the end of the hall might expect the distance of 12 with a probability of 0.8. A measurement of 9 with a probability of 0.1, a measurement of 6 with a probability of 0.05, etc.

A second remaining issue has to do with getting someone who has low vision to be able to accurately point a measuring device down a hallway. Pointing is a difficult task in the absence of vision. We have multiple solutions to this problem. Some of which are engineering solutions that are beyond the scope of this manuscript. However, in addition to the engineering solutions, we have found that the probabilistic observations described above allow for inaccurate pointing. Inaccurate pointing will generate an observation that is shorter than the true observation. This is generated by pointing the range finder off of the floor, ceiling or one of the side walls. Again, by expecting a collection of observations, the system should be robust to these types of errors.

B. Summary and Conclusions

For individuals with low-vision, navigating through an unfamiliar, large-scale environment can be a daunting task. The challenge lies in difficulties associated with both *obstacle avoidance* and *wayfinding*. The current studies investigated the efficacy of a low-vision navigation aid that uses a Bayesian model (Partially Observable Markov Decision Process) to guide a user from one unspecified location to the

[10] Although we suspect that in very complex environments one might find that the system has some benefits for even individuals with normal vision.

user's goal state. In all three experiments, participants with the NavAid system outperformed individuals without the NavAid system. In fact, in Experiment 1A subjects with the NavAid and degraded vision performed as well as someone with un-degraded vision and in Experiment 1B subjects performed better than subjects with un-degraded vision. In Study 2, we tested the NavAid system in a real environment. Again, we found that the performance with the NavAid system was better than performance without the NavAid system.

The results from these initial studies are encouraging and we believe that the NavAid has the potential to serve as an effective low-vision navigation tool in the near future. Future research will evaluate the NavAid system with individuals who have low-vision in real buildings. We are very interested in how well the system works for the intended population. There are a number of user interface challenges that remain to be resolved for the low-vision user. We are also interested in investigating how effective the NavAid system is within buildings that have dynamic obstacles (e.g., other pedestrians).

Although there are a number of issues that remain to be resolved, we are encouraged by our initial findings with such a rudimentary system. The issues discussed above and other design issues remain to be addressed, but we are confident that most, if not all of these, can be addressed, tested and resolved in the near future.

REFERENCES

[1] J. McNeil, "Americans with disabilities: 1997," U.S. Government Printing Office, Washington, DC, Current Population Reports P70-61, 2001.

[2] J. R. Marston and R. G. Golledge, "The hidden demand for participation in activities and travel by persons who are visually impaired," *Journal of Visual Impairment & Blindness*, vol. 97, no. 8, pp. 475–488, August 2003.

[3] D. G. Pelli, "The visual requirements of mobility." in *Low Vision: Principles and Application*, G. C. Woo, Ed. New York: Springer-Verlag, 1987, pp. 134–146.

[4] E. Tolman, "Cognitive maps in rats and men," *The Psychological Review*, vol. 55(4), pp. 189–208, 1948.

[5] R. G. Golledge, J. M. Loomis, R. L. Klatzky, A. Flury, and X. Yang, "Designing a personal guidance system to aid navigation without sight: progress on the gis component," *International Journal of Geographic Information Systems*, vol. 5, no. 373-395, 1991.

[6] R. G. Golledge, R. L. Klatzky, J. M. Loomis, J. Speigle, and J. Tietz, "A geographical information system for a gps based personal guidance system," *International Journal of Geographical Information Science*, vol. 12, pp. 727–749, 1998.

[7] J. M. Loomis, R. G. Golledge, and R. L. Klatzky, "Gps-based navigation systems for the visually impaired," in *Fundamentals of wearable computers and augmented reality*, W. Barfield and T. Caudell, Eds. Mahwah, NJ: Lawrence Erlbaum Associates, 2001, pp. 429–446.

[8] ——, "Navigation system for the blind: Auditory display modes and guidance," *Presence: Teleoperators and Virtual Environments*, vol. 7, pp. 193–203, 1999.

[9] A. P. for the Blind, "Annual report october 1, 1998-september 30, 1999," Louisville, KY, 1999.

[10] L. Goldish, *Braille in the United States: its production, distribution, and use*. American Foundation for the Blind, 1967.

[11] W. Loughborough, "Talking lights," *Journal of Visual Impairment and Blindness*, vol. 73, p. 243, 1979.

[12] J. A. Brabyn and L. A. Brabyn, "Speech intelligibility of the talking signs," *Journal of Visual Impairment and Blindness*, vol. 76, pp. 77–78, 1982.

[13] L. A. Brabyn and J. A. Brabyn, "An evaluation of talking signs for the blind," *Human Factors*, vol. 25, pp. 49–53, 1983.

[14] J. R. Marston and R. G. Golledge, "Improving transit access for the blind and vision impaired," *Intellimotion, Research Updates in Intelligent Transportation Systems, Transit Research Issue*, vol. 7, no. 2, pp. 4–11, 1988.

[15] B. Bentzen and P. Mitchell, "Audible signage as a wayfinding aid: Verbal landmark versus talking signs," *Journal of Visual Impairment & Blindness*, vol. 89, pp. 494–505, 1995.

[16] L. P. Kaelbling, A. R. Cassandra, and J. A. Kurien, "Acting under uncertainty: Discrete bayesian models for mobile-robot navigation," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1996.

[17] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, pp. 99–134, 1998.

[18] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman, "Acting optimally in partially observable stochastic domains," in *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, vol. 2. Seattle, Washington, USA: AAAI Press/MIT Press, 1994, pp. 1023–1028. [Online]. Available: citeseer.nj.nec.com/cassandra94acting.html

[19] B. J. Stankiewicz, G. E. Legge, J. S. Mansfield, and E. J. Schlicht, "Lost in virtual space: Studies in human and ideal spatial navigation," *Journal of Experimental Psychology: Human Perception and Performance*, In Press.

[20] E. Sondik, "The optimal control of partially observable markov decision processes," Ph.D. Thesis, Stanford University, 1971.

[21] S. Thrun, S. Gutmann, D. Fox, W. Burgard, and B. J. Kuipers, "Integrating topological and metric maps for mobile robot navigation: A statistical approach," in *Proc. 15th National Conf. on Artificial Intelligence (AAAI-98)*. AAAI/MIT Press, 1998, pp. 989–995.

[22] W. Burgard, A. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "The interactive museum tour-guide robot," *Proceedings of the AAAI Fifteenth National Conference on Artificial Intelligence*, 1998.