Customizing Directions in an Automated Wayfinding System for Individuals with Cognitive Impairment

Alan L. Liu*, Harlan Hile*, Gaetano Borriello*, Pat A. Brown#, Mark Harniss#, Henry Kautz@, Kurt Johnson#

*Computer Science & Engineering
University of Washington
Seattle, WA 98195-2350 USA
{aliu,harlan,gaetano}
@cs.washington.edu

#Department of Rehabilitation Medicine
University of Washington
Seattle, WA 98195-7920 USA
{pabrown,mharniss,kjohnson}
@u.washington.edu

@Computer Science
University of Rochester
Rochester, NY 14627-0226 USA
kautz@cs.rochester.edu

ABSTRACT

Individuals with cognitive impairments would prefer to live independently, however issues in wayfinding prevent many from fully living, working, and participating in their community. Our research has focused on designing, prototyping, and evaluating a mobile wayfinding system to aid such individuals. Building on the feedback gathered from potential users, we have implemented the system's automated direction selection functionality. Using a decision-theoretic approach, we believe we can create better wayfinding experience that assists users to reach their destination more intuitively than traditional navigation systems. This paper describes the system and results from a study using system-generated directions that inform us of key customization factors that would provide improved wayfinding assistance for individual users.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology, User-centered design, Prototyping*; K.4.2 [**Computers and Society**]: Social Issues—*Assistive technologies for persons with disabilities*

General Terms

Design, Human Factors

Keywords

Wayfinding, user interface, cognitive impairments, Markov decision process

1. INTRODUCTION

Wayfinding is a concern for individuals with cognitive impairments. Being unable to find their way safely and independently limits their ability to fully participate in their community, while also placing a burden on their caregivers and community services. This paper presents the design of an automated system for producing real-time wayfinding direction delivered on a mobile device (see Figure 1), and a user study of the system with potential users that investigates individual preferences and abilities to follow the different types of supported directions.

Our system generates directions of two core types:

1. **Landmark-based directions** that direct a user in relation to a landmark. Landmarks are useful for giving users a better sense of where they're going and whether they're going the right way. Landmarks may also require higher cognitive effort due to the need to recognize visual features from a photo.
2. **Turn-based directions** that direct a user more precisely using short messages with an iconic presentation (see Figure 2). They usually require less cognitive effort but leave out possibly useful wayfinding detail, which can lead to users paying closer attention to the system. In the worst case, a user might focus too much on these directions and pay too little attention to traffic. They also require highly accurate location to minimize problems with message timing.

Figure 1: Wayfinding system front-end running on a Nokia N95 mobile phone. A landmark-based direction is shown with a selected photo, overlaid arrow, and text. Audio equivalent to the text is also produced using the built-in text-to-speech function. The phone is used in landscape mode to utilize a larger portion of the screen for images. The screen remains slid out to expose the Global Positioning System (GPS) antenna and keypad, which the top row of keys used to repeat directions. The center face button is used as a Help button that requests a different direction.

ASSETS'09, October 25-28, 2009, Pittsburgh, Pennsylvania, USA.

(a) Straight and standard left/right turns use solid arrows

(b) Stop and U-turns use standard street signs

(c) Modified street sign for non-standard intersections

Figure 2: Examples of turn-based directions.

Our system consists of two parts: a landmark-selection system that produces landmark-based directions using a photo database, and a direction-selection system that uses a decision-theoretic approach to control routing and message delivery.

Studies have shown that landmarks are the predominant way of producing wayfinding directions for people without disabilities [10] and that they can help older people by reducing the cognitive load required to wayfind [6], however until recently integrating landmarks into wayfinding directions has not been practical. Because of technological advances that promise to provide scalable and ubiquitous access to landmark information, those type of directions are now possible. Our landmark-selection system uses a database of geo-tagged digital photos (photos containing landmark and GPS location metadata) to determine landmark visibility. It selects the most appropriate photo of a desired landmark and augments the image with an overlaid arrow. See Figure 3 for example images produced by the system. Section 3 describes the landmark system.

Our prior work showed that individuals with cognitive impairments would benefit from a wayfinding system that is capable of supporting customizable and adaptable direction selection. This can be achieved by using a decision-theoretic approach where system actions are chosen based on knowledge of a user model. In many ways this is similar to the path planning problem in the robotics community. Various techniques for path planning under uncertainty have been developed by that community [15], which we can apply toward creating an automated wayfinding system. We model the problem of direction selection using the Markov decision process (MDP) framework. An individual user is represented as an agent in the MDP, whose state contains location and orientation information (among other data). A user's wayfinding preferences and abilities modeled as action costs and state transition probabilities. Solving the MDP produces a policy (an optimal set of rules) that determines the system action to take at each decision point. Section 4 describes in more detail the MDP we use.

To better understand key customization factors to produce personalized directions, we conducted a study described in Section 5. The study used two different user models to explore potential users' ability to follow different directions, and their reaction to those directions. In Section 6, we discuss how our system can be changed to support those customization, through both enhancements to the landmark-selection system, the user model, and the user interface. Finally, we end the paper with conclusions in Section 7.

2. RELATED WORK

The design of our system is a result of several previous iterations that focused on gathering feedback from potential users and observing how they respond to different types of wayfinding directions. Our first user study involved a prototype designed to support indoor navigation using combinations of photos, symbols, text, and audio [8]. The results from that study were positive, showing that potential users were largely able to use the directions to reach a destination in an unfamiliar building. However, we also found a wide range of preferences for the types of directions, with some participants preferring visual directions while others found them more cognitively challenging. Therefore, the ability to modify system behavior to reflect individual preference was a necessary feature.

In a second round of user studies, we found that the directions used in indoor navigation had several issues when used outdoors [9]. Photos without clear landmarks were difficult to match with the real world, and the step-by-step process involved constant user attention to the detriment of user awareness for safety. We therefore incorporated landmark-based directions and tested their feasibility in a controlled study. The findings from the controlled study found that heading toward a landmark was significantly easier and less error-prone than any other direction type using landmarks. However, we did not use the landmark-based directions as part of a route for users to go through, nor were we able to compare the landmark-based directions to our previous directions. The system in this paper can produce all the types of directions we have studied, and our user study involves examining how individuals reacted to the direction types.

Like our initial prototype, the system tested by [4] used augmented photos to direct potential users with cognitive impairments. However, because our subsequent studies found drawbacks to relying on those type of directions, the system we designed and study in this paper can choose between such "low-level" or precise directions and "higher-level" landmark-based directions capable of more than one type of direction.

Related work by [3] focused mainly on the infrastructure to deliver handcrafted directions to potential users with cognitive impairments, as opposed to automatically generated ones like the system described in this paper. Other work has focused on predicting a user's desired destination [11]

An obvious related technology for wayfinding are the existing commercial navigation systems. Unlike our system, current systems are limited in their ability to support customization. For example, Global Positioning System (GPS) navigation devices give users the ability to choose between

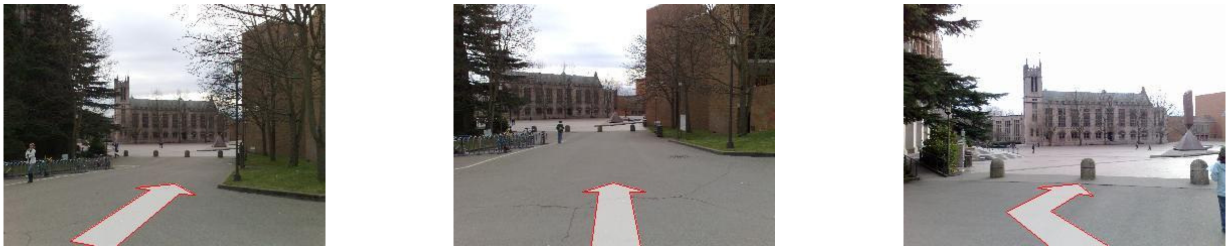

(a) An example landmark and overlaid path arrow

(b) The same direction as in 3a but from a different location and orientation

(c) A example showing a jointed arrow to denote an upcoming turn

Figure 3: Examples of augmented landmark-based photos.

quickest and shortest routes, but every user that chooses the same route will receive the same type of direction, without regard to user preference. The GPS unit only modifies its calculated route if a user chooses to deviate from it, however we have found that a common behavior when a user becomes confused is to remain in one place. Because GPS units do not have alternative methods for delivering directions, they cannot produce different levels of help that a user may need in those instances. Finally, current devices do not support incorporating landmarks into directions, despite the utility of landmarks in pedestrian wayfinding. This leads to the need to constantly pay attention to the directions given by the system, leading to safety issues such as not watching out for traffic [12].

Recently, systems to choose appropriate landmarks for directions [1] and generating user-specific directions based on familiar landmarks [5] have been studied. In [1], a mobile phone game is used to create a set of photos where landmarks are more easily found, compared to randomly chosen, nearby geo-tagged photos from online photo sites. Like their system, our current system requires an initial seeded set of images with landmark name and location tagged. Unlike their system, our system can immediately use those photos without a second phase of user input, because it uses heuristics such as landmark popularity to choose appropriate landmarks. While [5] shows that using certain landmarks familiar to users may be beneficial, the study used directions generated statically (beforehand) as text, rather than delivered visually in a just-in-time manner to users who may not be familiar at all with their route.

3. LANDMARK SELECTION SYSTEM

Automatically generating landmark based directions requires selecting an appropriate landmark and an image of that landmark. The landmark selection system leverages existing collections of geo-tagged images (images with location and landmark metadata) to retrieve suitable images of landmarks [7]. This makes it possible to select an image from the database that relates to the user's current location and intended direction, for example, to select an image of the building they should walk toward in a perspective close to their current position. Additional aspects of the image database make it possible to choose landmarks by popularity (heuristically, based on the number of images in the database) and to choose a quality representative view from the possible choices. These images can also be augmented with arrows See Figure 3 for some example images that were automatically constructed by the system and `http://vcui1.nokiapaloalto.com/marwebapi/marwebapi/apiindex` for a Web-based front-end.

4. MARKOV DECISION PROCESS

The robotics community has developed various techniques for path planning under uncertainty [15]. A key concept in this context is the Markov decision process (MDP), which provides techniques for generating navigation plans even when observations and the outcome of navigation actions are uncertain [14]. For example, partially-observable MDPs have been used to assist persons with dementia through tasks such as hand-washing [2]. The framework we chose builds upon these techniques to model uncertainty in whether or not a person will follow the guidance provided by our system, but to reduce the state space necessary to solve our MDPs, we rely only on observable action results.

MDPs are defined by **state** and **action** sets and one-step **transitions**. States have associated **rewards**, and a solution to an MDP is a **policy** that maps states to actions in order to maximize expected reward. A key aspect of MDPs is that they can be used as a framework for learning and adaptation – transition probabilities may be approximated at first and then updated given observed behavior. Techniques for solving MDPs have been shown to enable the generation of navigation plans in robotics, and we believe they map well to our problem of choosing directions along a route for an individual with cognitive impairment that maximizes the chances of success.

There are some differences between previous applications of MDPs to robot navigation tasks versus producing wayfinding assistance to a person. One noteworthy alteration to traditional MDPs that we are using are **options**, which we use to represent each direction [13]. Traditional MDPs rely on fixed time slices and actions take only a single time step, but obviously users may take a variable amount of time when following directions. Options can be considered to be sub-policies that dictate when certain directions can be given, what expected transitions will occur within the option, and when they should terminate so that a new option may be selected. Incorporating options allows us to reason over temporally-extended actions with a lower computational cost with a trade-off in terms of the optimality of the eventual solution.

4.1 State

The state consists of variables that determine available options and affect transition probabilities. The most accurate model would incorporate every valid, observable variable into the state, but at the expense of a huge state space. The initial design of our model contains only a very small set of state variables, with the intention of investigating whether it is sufficient, and if not, what other variables would be the most beneficial to incorporate into the model.

The state consists of user position (location and orientation), the current option selected by the system, and a options blacklist. User position is a critical variable. Location is represented as a node in a graph network, while orientation is represented as either a heading toward an adjacent node or from such a node. We consider the current option as part of the state for the sake of simplicity, because the choice of options is determined by a combination of the current option and the (other) state variables. The options blacklist provides a way to temporarily avoid using directions that have been unsuccessful in the recent past. It is represented as an array of decrementing counters, where an option is given a positive value when a user incorrectly interprets that option, and that option is not available for that many subsequent future states. A new model is computed when the blacklist is changed, generating an alternative set of directions.

4.2 Option

An option is a sub-policy with its own set of termination states. To simplify our model we define the set of termination states of an option to be any state where the user has changed to an adjacent location, changed orientation, or not moved at all after an abnormally long time. The sub-policy determines when a state change has occurred. This determination is not yet implemented in our system, so we currently simulate that functionality manually.

The options available in our system are:

- **Straight:** This option is given to make a user move forward in the direction that they are facing. It might be used to provide further assurance to the user that they are moving in the correct direction, or to prompt a user to continue moving in case the user expects a new prompt from the system.
- **Turns:** This category of options covers making a turn. There are different turn options depending upon the type of intersection. A regular left/right turn is available if there is a standard intersection, while slight turns are usable when the angle of the turn is closer to straight. First and second turns are supported in the case where there are more than one possible path on that side. We choose not to support hard turns to keep options as simple as possible for users.
- **Turn around:** This is the U-turn option. It may not be as effective, as we have found that turning around can be disorienting for some, but it can be used to correct a user moving in an undesirable direction.
- **Landmark:** Based off our landmark analysis study [9], we found that a number of users are capable of identifying and moving toward a landmark provided a photo. Because users did significantly better moving toward a landmark in front of them, we choose to make this option available only when landmarks are visible and in front of a user.
- **Stop:** This option is given to make a user stop. Reasons for stopping include notifying the participant of reaching a destination, or introducing a pause in wayfinding when close to traffic for extra caution, or to prevent a user from moving too quickly away from a "better" (from a wayfinding standpoint) location.

4.3 Transitions

Once states are defined, transition probabilities are assigned to map initial state-option pairs to subsequent states. In other words, if the system were to choose an option while the user is in a given state, how will the user react? This is where user customization plays a major role. Transition probabilities can be seeded based on general intuition of wayfinding (*e.g.*, the findings in our study that directing a user toward a nearby landmark has a higher likelihood of success) can be defined manually if enough information is known about the user's preferences and capabilities. Otherwise, the user might go through some evaluation routes to collect initial data on their tendencies.

Health conditions with impact on wayfinding can dictate an initial distribution over the probabilities. For individuals with visual impairments, directions that use audio and text may be preferred over those that involve detailed landmark photos, so the chances that the former type of directions are successfully followed should be greater than the latter's.

4.4 Costs and rewards

Ultimately reaching the goal location is a highly desirable state, while giving a new direction has a cost both in terms of cognitive load, distraction, and time. Some of our study participants mentioned that reaching a destination in the quickest amount of time or taking the shortest path is less important to them than being more aware of their route or using less physical effort. Some users benefit from "high-level" directions that apply beyond a single turn, because they can plan their own routes farther, while others interpret directions more literally, so directions for them may need to be step-by-step or include additional wording to avoid ambiguity.

4.5 Producing a Policy

A policy is a mapping from every state to an option that produces the highest-valued expected reward. It is calculated by iteratively backing up state values. Each state's value is the sum of its reward and a discounted expected value of all next states of the best option. The discount factor puts a preference on nearer-term rewards. We use value iteration to determine the policy, with a termination threshold that stops iterating when no state's value changes between iterations by more than a small delta.

5. USER STUDY

To be useful, generated directions for an individual requires an accurate user model. We therefore conducted a study to learn key individual preferences and abilities that would suggest the types of customizations needed to be supported by the system. Using these findings, we would be able to support system customization for future users, who might initially either go through similar trial usage sessions and/or answer a set of pre-defined questions to "bootstrap" their system's model.

The study involves wayfinding through two routes on our university campus. Two different static user models are used to create policies for these routes. Similar to our methodology in prior user studies, user position and message timing are still controlled by a *wizard*, a person who shadows

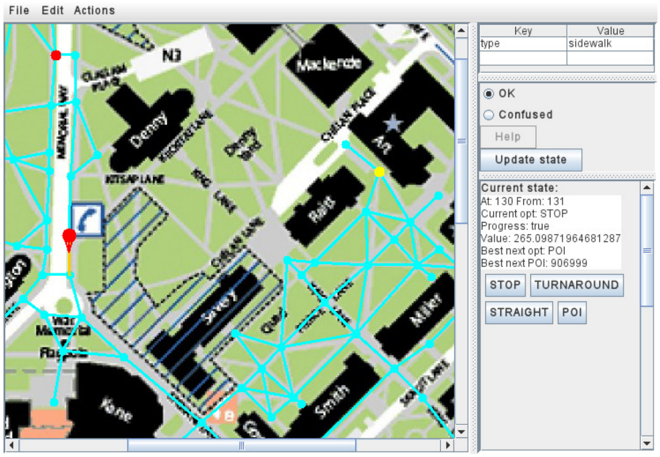

Figure 4: The Tablet PC Wizard-of-Oz interface used in our latest user study. The left pane shows the graph network with Route 2 in our study from the start (top-left node in red) to the destination (south of the Art building on the right, in yellow). Metadata for the route segment ahead of the user's position (the dot and arrow) is shown in the upper right table and can be used by the model or by the prompting system. The lower right pane shows the current state variables and available options.

the participant with a Tablet PC (see Figure 4 for the wizard/system interface). Unlike our prior studies, all system prompts are automatically picked by the system calculating a policy using the active model. In addition to the wizard, two researchers shadow participants to note how they react to the system, and to intervene if there are safety concerns. Success and difficulties wayfinding due to choice of prompt, wording, timing, etc. and wayfinding behaviors are noted.

5.1 Initial Models for Study

We defined two models to study that use the same state variables, but have different transition probabilities and costs.

Model 1: The Naïve Model: The first model incorporates few assumptions about wayfinding. It assumes that every option will be correctly followed 100% of the time, and that all options have the same base cost. An additional cost for distance traveled is added so that the distance of a route is also factored into the policy. A high reward is given for reaching the goal location, and no discount factor is applied to future state values. When more than one option has the same expected value, one is picked pseudorandomly based on a distribution seeded by the state. This ensures that there is no variation in the policy across different users.

Model 2: Landmarks for Longer-range Movement: The second model adds some additional assumptions about wayfinding to move closer to a model suggested by our previous study results. First, it penalizes the Stop and Turn Around options by assigning higher costs to those options, since users tend to prefer that the system adapt to their chosen route rather than be corrected after every minor route deviation. Second, it assumes that lower-level options are more likely to be followed correctly, but gives an additional bonus to use a Landmark option when the same landmark was just followed correctly. This models the assumption that a landmark option potentially requires more cognitive processing by a user at first, but that there is an advantage to using a landmark for longer, continuous stretches of a route. All other positions are given a small, non-zero probability, and the total weights are normalized to represent the transition probability of the options. A discount factor of 0.98 is applied to future state values to reflect the uncertainty of future options and slightly prefer nearer-term rewards.

Table 1: Participant demographics. *: *Uses powered chair*

#	Age	Sex	Health Condition
1	48	Female	Traumatic brain injury
2	27	Male	Asperger's
3	49	Male	Traumatic brain injury
4	22	Male	Cerebral palsy*
5	21	Male	Traumatic brain injury
6	25	Male	Cerebral palsy*
7	45	Male	Multiple sclerosis

5.2 Participants

We recruited 7 participants with cognitive impairments through the University of Washington Center on Outcomes Research in Rehabilitation and an outpatient rehabilitation clinic. Participants ranged in age from 21-49 (mean 34) with 1 woman and 6 men (see Table 1 for the demographics). We also conducted the same study with 6 participants without impairments, ranging in age from 28-46 (mean 36.7) with 2 women and 4 men. As the results from the participants without impairments were similar to the results found from the participants with impairments, this paper will focus on the findings from studying the participants with impairments.

5.3 Results

As we were only able to recruit a modest number of participants for the study, we focus more on the participants' notable actions and comments during the study and debriefing session. Of the 7 participants with cognitive impairments, 6 were able to follow the directions with only minor errors such as wrong turns or other issues due to our prototype system's limitations. Participant 4 had some minor problems interpreting landmark photos along his first route when using Model 1 because he did not notice the overlaid arrows on the photos, while on the second route using Model 2, he needed a significant number of rerouting directions from the system to recover from several wrong turns.

5.3.1 Landmark-based directions

Landmark-based directions were well-received by some but found less useful than the turn-based directions by others. Participants 1, 2, 3, and 5 all expressed the usefulness of using the photos to wayfind. Participant 3, in particular, would often press the Help button when given other types of directions in order to receive a photo-based direction. Suggesting the usefulness of Model 2, all of these participants were able to wayfind flawlessly when a follow-up direction used the same landmark (*e.g., "Continue along the path toward <landmark>"*).

P1: *"When you have an actual picture of what's in front of the person, that's excellent."*

P2: *"I can coordinate [matching the picture to what he can see] pretty well."*

Participants 4, 6, and 7 preferred the lower-level, turn-based directions. Participant 4 and 7 had difficulty seeing the images on the phone, while Participant 6 simply found the turn-based directions to be easier. During the debriefing session, Participant 7 confirmed that the lighting conditions outside made it difficult for him to see the images, but under better conditions he was able to see the images more clearly, suggesting that a better (larger, brighter, and/or higher-contrast) screen might have made a difference in his wayfinding experience.

The other notable usability issue with the landmark-based directions occurred when the provided photo was taken from a different vantage point than the participant's actual location. In some cases, a photo would be used from the other side of the street. In these cases, participants would either trust their own judgment and choose a direction to go in, or request a different type of direction.

> ***P1:*** *"The towers [of the Campanile] were behind trees, but I trusted myself to go forward."*
> ***P2:*** *"I think [the directions were] pretty clear... if I made a mistake it could tell me to turn around."*
> ***P7:*** *"I couldn't see the photos, there was too much glare, but the façade and context helped in certain cases."*

5.3.2 Using landmark names

In our study, we chose to include the name of the landmark in every landmark-based direction. Participants expressed different opinions on using the names. Some participants, though unfamiliar with the campus, stated that they would try to match the names to signs as a way of confirming their direction, and that it could be helpful in remembering the landmarks for future use. Other participants stated that they ignored the names. In one case where using the landmark name was detrimental, Participant 6 made an incorrect turn based on an incomplete knowledge of where a certain landmark was on the campus.

> ***P2:*** *"Some of the signs on the buildings are tricky, you just have to figure out where the signs are."*

Participants suggested that only referring to a name in limited cases, such as to describe their destination, would allow them to focus on finding the landmark carefully only when it was necessary.

5.3.3 Using compound directions

In our previous studies we found that sequences of turns could be problematic when they occurred close together. Some of our directions were therefore compound directions such as *"Go along the sidewalk toward <landmark> and take the next left."*, followed by a reminder to turn once the participant came closer to the intersection. This caused some confusion for participants, who interpreted the second turn direction as a separate instruction to be completed after the first, compound direction. Participants noted that they felt that they were falling behind while the directions were delivered too quickly.

At other times, the compound directions were misinterpreted. For instance, Participant 6 would sometimes perform the straight and turn segments in reverse order.

5.3.4 Other direction issues

Some participants had a few issues with non-standard turns (slight turns, or "second left"-style turns such as in Figure 2c). Participants 4 and 6 missed a segment of the second route that the system attempted to direct them towards when they did not spot the accessible ramp to the side of a small flight of stairs.

5.3.5 Using help

Participants 3 and 4 used the Help function often to get different directions, for opposite reasons. Participant 3 preferred the landmark-based directions and would use the button to request them, while Participant 4 did not find them helpful. The other participants used the help function sparingly, or not at all, which matches our previous observation that potential users may be better served with some form of automatic assistance.

5.3.6 Attention between device and environment

We observed that Participants 2, 3, 5, and 6 were able to balance their attention between the device and the environment. Participant 4 heavily relied on the text-to-speech audio and did not pay attention to the environment at most times, though unlike in our previous studies, no participant crossed a street at an unsafe location.

Participant 1 explained that some combination of being nervous and unfamiliar with the device might have contributed to her focusing too heavily on the device, while Participant 7 also mentioned that increased familiarity with the device would have changed how he used it.

> ***P1:*** *"I would have enjoyed the walk a lot more if I just would've relaxed. [I] know it's going to buzz in [my] hand so then [I] can look at it, but I looked at it the whole time."*
> ***P7:*** *"I was thinking for a while, maybe I should stick this in my pocket and react when the thing goes off. I didn't do it because [I thought], 'Well, by the time I take it out, will I miss the instruction?' but obviously not, the instruction stayed on which was good, so I should have."*

5.3.7 Hardware usability

We observed some drawbacks to using the N95 as the sole device for this study. As previously noted, some participants found the screen on the N95 to be insufficiently bright for displaying images, though others were able to see them. Participants 4 and 6, who have cerebral palsy and use motorized wheelchairs, found simultaneously holding the device and pressing buttons to be uncomfortable. Participant 4 suggested that an arm attachment would allow him to detect the vibratory new direction alert, while Participant 6 suggested a swivel mount that could be attached to one armrest of his chair.

Another issue during the study were accidental keypresses that switched phone applications. As we were unable to disable many of the keys on the N95, this caused issues where miscellaneous applications would take the foreground and had to be exited manually by a researcher.

Aside from these issues, the size and form factor of the N95 was deemed acceptable. Participants felt comfortable with using it out in the public without stigma.

> ***P3:*** *"I don't find [the device] obtrusive. Nowadays everybody's carrying something."*

5.3.8 Overall participant impressions

The overall impressions that our participants had were highly positive. Participants expressed the desire to have the system as soon as it becomes available, with most volunteering to participate in further studies as the system is enhanced.

P1: *"I'd take it to a new city. That would be fantastic!"*
P5: *"I use my iPhone for directions, but if I had this, I would definitely prefer to use it for the images."*

6. DISCUSSION AND FUTURE WORK

The results showed that the system, with a few adjustments, could be suitable for potential users to independently wayfind. It also confirmed that there would be individual, disparate preferences and interpretations of the directions. Below are some of the customizations that could be needed to produce high quality directions suitable for the user.

Individual preference for types of direction: Solving an MDP produces a policy that chooses the direction that the model believes most likely produces the highest reward. The transition probabilities that the model relies on to calculate this policy can be initially populated with approximate values based on our user studies. For example, the probability that a user will successfully follow a direction that involves moving toward a landmark ahead of her would be higher than that of a landmark behind her. A direction that uses a photo with an overlaid arrow might have an even higher chance of success, but only lead the user to a nearby location, while the landmark-based direction could allow the user to make more progress towards the destination. In addition, some study participants have stated that they preferred to have a better idea of where they are going, so landmark-based directions might again be preferred and states involving their use would be assigned higher rewards.

Leveraging user familiarity: Landmarks familiar to a user should be used more often than those less familiar. Similarly, users comfortable using cardinal directions might be given them more often. The model can reflect this by boosting the probabilities of success for options that use familiar places or terms.

Level of detail in a single direction: Our studies showed that while minimizing the cognitive load necessary to interpret a direction is beneficial, some users may benefit from more "high-level" directions that cover more than one route segment (see Figure 3c), because they can plan their own routes farther. Some individuals interpret directions more literally, so directions for them may need to be more in the step-by-step fashion or include additional wording to avoid ambiguity.

Health conditions that impact route planning: One critical aspect to customization is that people with cognitive impairments often have other health conditions that impact their device use and the appropriateness of routes. Visual or hearing impairments would affect the efficacy of landmark or audio directions, respectively. For individuals who may be concerned about fatigue, such as Participant 7, the MDP can increase the weight of distance to cost so as to minimize the total effort involved. For those who require accessible pathways and entrances, options involving stairs or other inaccessible paths can be heavily penalized.

Detecting errors and intervening: Different users make different errors. Some find that they cannot follow a direction, stop to look around, and become frustrated. Others will pick a direction even when unsure. Some want to be corrected right away if they are taking a path that makes it difficult to recover later (for example, people with multiple sclerosis, who get fatigued). Others are fine with their decision and just want to proceed rather than backtrack if it's unnecessary. Encoding these errors as states involves identifying the behaviors based on the individual's tendencies and assigning large penalties to them, while also terminating the option that caused the error so a new option can be selected.

Customizing for safety: We observed that some individuals tend to be more careful and maintain an awareness of their surroundings, whereas others do not always remember to look out for vehicles and other hazards. While teaching individuals about traffic safety is outside the scope of our design, our system must be able to be customized for the range of user behavior in regards to safety. For users who are not in the habit of crossing only at crosswalks or do not always look out for traffic when crossing streets, the system might explicitly direct the user to a crosswalk and instruct them to wait until the street is clear of traffic before having them cross. For users who tend to be more careful, an occasional safety reminder might suffice.

Better landmark images: Aside from model enhancements, we have begun working to produce better images when there is a discrepancy between between a user's location and where a photo was originally taken. Figure 5 shows alternative images synthesized by the landmark-selection system. We have preliminarily evaluated these images with the group of participants without cognitive impairments, and the results suggest the potential for improved understandability of the directions, especially Figure 5d.

Alternative form factors: We are considering other hardware form factors to accommodate usability issues. For example, a touchscreen phone would allow us to reduce the number of physical buttons and also present custom, labeled buttons for various functions.

6.1 Adapting the Model

Solving an MDP involves computing the expected value of every state and producing a policy that decides on the highest-valued option to take at each state. The quality of the policy is only as good as the accuracy of the underlying model at predicting user behavior. Using general wayfinding insight and a short pre-trial evaluation to define a model's initial transition probabilities should provide a reasonable first policy. The model can be improved by harnessing the MDP framework, which uses observed behaviors to adjust probabilities and converge upon the actual underlying model.

One challenge of interpreting a user's behavior given an option is deciding whether it reflects the user's likelihood of success with that specific option ("on-policy" learning in reinforcement learning terms) or the user's likelihood of success with all options of that type ("off-policy" learning) applicable to all situations where that type of option could be given. The former would be the correct approach if the reason for the user's behavior had to do with only the state she was in (such as facing a temporary obstruction), while the latter would be the correct approach if the reason was due to something more permanent (such as becoming better at identifying landmarks or experiencing worsened eyesight).

It may be impossible for the system to know how far-reaching the reasons for an observed behavior are, so a trade-off must be made to ensure that the model does not adapt too drastically without more observations. Care will also

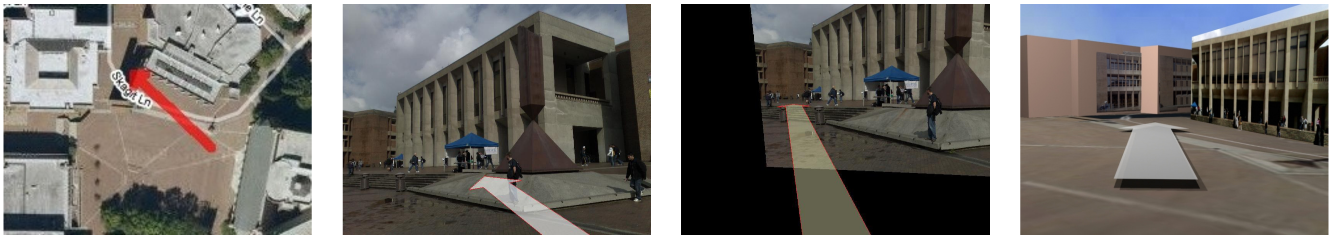

(a) A desired path segment shown on a satellite image

(b) A photo with overlaid arrow showing the desired path

(c) "Zoom-out" view that warps original photo to align to user's view

(d) View rendered using a lightweight 3-D model made from a small set of photos

Figure 5: Examples of different techniques to visually show a path involving landmarks.

have to be taken to determine how long the changes last. While the system should almost never immediately repeat a failed direction, it might be reasonable for the system to try that direction again in the future.

7. CONCLUSIONS

We have designed, prototyped, and evaluated a wayfinding assistance system that provides directions on a mobile device to individuals with cognitive impairments. Our prior user studies have informed the design of a decision-theoretic model for direction selection. Using the Markov Decision Process (MDP) framework, we are able to create customizable and adaptive user models that can determine the most appropriate sequence of directions to present to each individual. While there is still much to do to enhance the user model, results from the user study presented in this paper show that the system, as designed, can be capable of helping people find their way to their destination. As has been the case throughout our project, prioritizing further enhancements to our design will be informed by continued study and feedback with potential users.

8. ACKNOWLEDGMENTS

We thank Richard Anderson and Craig Prince for providing PC hardware for our user studies. This work was funded by the National Institute on Disability and Rehabilitation Research (NIDRR) Grant #H133A031739, a gift from Nokia Research, and a Microsoft Research Intelligent Systems for Assisted Cognition Award.

9. REFERENCES

[1] M. Bell, S. Reeves, B. Brown, S. Sherwood, D. MacMillan, J. Ferguson, and M. Chalmers. Eyespy: supporting navigation through play. In *Proceedings of CHI '09*, pages 123–132, New York, NY, USA, 2009.

[2] J. Boger, G. Fernie, P. Poupart, and A. Mihailidis. Using a POMDP Controller to Guide Persons With Dementia Through Activities of Daily Living. *Adjunct Proceedings of Ubicomp '03*, 2003.

[3] S. Carmien, M. Dawe, G. Fischer, A. Gorman, A. Kintsch, and J. F. Sullivan. Socio-technical environments supporting people with cognitive disabilities using public transportation. *ACM Transactions on Computer-Human Interaction*, 12(2):233 – 262, Jun 2005.

[4] Y.-J. Chang, S.-K. Tsai, and T.-Y. Wang. A context aware handheld wayfinding system for individuals with cognitive impairments. In *Proceedings of Assets '08*, pages 27–34, New York, NY, USA, 2008. ACM.

[5] J. Chung and C. Schmandt. Going my way: a user-aware route planner. In *Proceedings of CHI '09*, pages 1899–1902, New York, NY, USA, 2009.

[6] J. Goodman, P. Gray, K. Khammampad, and S. Brewster. Using Landmarks to Support Older People in Navigation. In S. Brewster, editor, *Proceedings of Mobile HCI 2004*, number 3160 in LNCS, pages 38–48. Springer-Verlag, Sept. 2004.

[7] H. Hile, R. Grzeszczuk, A. Liu, R. Vedantham, J. Kosecka, and G. Borriello. Landmark-Based Pedestrian Navigation with Enhanced Spatial Reasoning . In *Proceedings of Pervasive '09*. Springer-Verlag, 2009.

[8] A. L. Liu, H. Hile, H. Kautz, G. Borriello, P. A. Brown, M. Harniss, and K. Johnson. Indoor Wayfinding: Developing a Functional Interface for Individuals with Cognitive Impairments. In *Proceedings of Assets '06*, pages 95–102. ACM, 2006.

[9] A. L. Liu, H. Hile, H. Kautz, G. Borriello, P. A. Brown, M. Harniss, and K. Johnson. Informing the Design of an Automated Wayfinding System for Individuals with Cognitive Impairments. In *Proceedings of Pervasive Health '09*. IEEE, 2009.

[10] A. J. May, T. Ross, S. H. Bayer, and M. J. Tarkiainen. Pedestrian Navigation Aids: Information Requirements and Design Implications. *Personal Ubiquitous Computing*, 7(6):331–338, 2003.

[11] D. J. Patterson, L. Liao, K. Gajos, M. Collier, N. Livic, K. Olson, S. Wang, D. Fox, and H. Kautz. Opportunity knocks: a system to provide cognitive assistance with transportation services. In *Proceedings of Ubicomp '04*, 2004.

[12] E. Rukzio, M. Müller, and R. Hardy. Design, implementation and evaluation of a novel public display for pedestrian navigation: the rotating compass. In *Proceedings of CHI '09*, pages 113–122, New York, NY, USA, 2009.

[13] R. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112:181–211, 1999.

[14] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[15] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.