Evaluating Crowd Sourced Navigation for the Visually Impaired in a Virtual Environment

Wai L. Khoo, Greg Olmschenk, Zhigang Zhu, Tony Ro
Departments of Computer Science and Psychology
The City University of New York - The City College and the Graduate Center
New York, USA
wkhoo@gradcenter.cuny.edu, golmschenk@gradcenter.cuny.edu, zhu@cs.ccny.cuny.edu, tro@ccny.cuny.edu

***Abstract*—Crowdsourcing has been shown to be a powerful method for solving a variety of problems. In this paper, we introduce an approach for allowing a crowd to help navigate a visually impaired user to their destination in real-time. Furthermore, we experiment with several approaches in aggregating and feeding back crowd data to determine the optimal method. Our approach streams live video from the visually impaired user's mobile device to a crowd of sighted volunteers. Each crowd member is able to provide their opinion on how the user should proceed and our algorithm aggregates this into a single opinion that is sent as feedback to the user. In this paper, we first present the design and implementation of our crowd sourced navigation system, including webapp design and two aggregation algorithms: averaging and league leader approaches. A virtualized user (avatar) is also developed for more controlled and repeatable testing of aggregation approaches with multiple crowd sourced volunteers. We also tested two navigation modes: a real-user controlled avatar, and a program controlled avatar. Experimental results are provided with the two aggregation methods and the two navigation modes. Our results show that we do not have significant difference between the aggregation methods and the program controlled avatar performed better than a real-user controlled avatar.**

***Index Terms*—crowd sourcing; mobile; visually impaired; virtual environment**

I. Introduction

With the technology advances in sensors and mobile computing, more and more research and development efforts have been directed at assisting in navigation for visually impaired people (VIP). These include wayfinding using a cell phone camera with fiducial marks [1, 2], automatic cross-walk, bus, and traffic signal detection [3, 4, 5], stereo vision and/or RGB-D sensors for blind navigation [6, 7, 8], and label and signage reading [9, 10]. However none of these have been robust enough for practical use by blind users. There is still a long way to go to achieving a wearable vision system comparable to the Google cars with advanced sensors and high computing capacities. Alternatively, crowdsourcing assistance has been studied for various applications for the blind and visually impaired, elderly, and people in need, such as video annotation, label reading, and assisted navigation, which offer promise for real applications.

In this project, we propose a smartphone based crowd sourced navigation solution with a focus in evaluation in a virtual environment. To get from point A to point B, a VIP can use their smartphone to stream video of the view in front of them to an online portal where a group of "crowd volunteers" can direct them to their destination. Algorithm developments include developing a smartphone app to enable such service, a user-interface of the online portal, and an aggregation algorithm that transforms a set of volunteers' directions to a single response for the VIP. Before we conduct real users testing, we propose to use virtual environments to simulate VIP navigation while testing the crowd's ability in assisted guidance. Not only does virtual environments give us a more controlled study (with regards to person safety and inherited devices limitations), but also allows us to establish ground truth data for future comparisons, fine-tuning the aggregation algorithm, and troubleshoot the online portal.

In Section II we discuss related research that has been done on crowd sourcing and assisting the visually impaired. Section **??** describes our system design and Section IV describes our virtual environment design. In Section V, we examine the results from our early testing and provide qualitative analysis. Finally, we present our conclusions in Section VI.

II. Related Work

Many groups have used crowdsourcing to assist blind users. BlindSquare [11] is MIPsoft's GPS navigation software for the iPhone and iPad. It differs from other navigation applications by using crowd sourced data; it uses Foursquare for points of interest and OpenStreetMap for street information. Jeff Bigham of CMU pioneered the work on label reading using crowdsourcing in his Interactive Crowd Support system, VizWiz [12] and his team has collected thousands of images sent by visually impaired people. Researchers at Smith-Kettlewell Eye Research Institute use crowdsourcing for movie annotation for the blind [13]. However, the current research and services focus on static image queries which are inadequate for real-time assisted navigation.

So far the technologies in vision recognition are not reliable enough for successful applications of automated navigation. On the other hand, smartphone video streaming over Internet has been a mature technology, and humans are far more reliable than machines in recognizing situations for daily navigation and reading. Such real time video streaming services include Skype, Google Hangouts, WebEx, etc. We have also set up a preliminary crowd navigation testing site exhibiting the possibility of using real-time video streaming to assist a blind user in navigating by live feedback from volunteers online.

Crowdsourcing has proved to be an effective way to collect large amounts of labels for many machine learning tasks [14, 15]. A key element in crowdsourcing is how to aggregate the noisy labels. Popular choices include average aggregation, majority voting, and minimax entropy based approaches [16, 17, 18]. We plan to further tailor these existing techniques to address the unique challenges in crowd-assisted navigation, such as the smoothness of label reliability of the volunteers, the contextual information of video frames, etc.

Our proposed work is different from existing work in a number of ways: First, in time constraints: we want to return the end user the 'best-effort' answer within a tight time window, best in real-time. Therefore, both the response from the crowd and the consequent processing needs to be done in real-time or near real-time. Second, a safe, controlled, "ground-truth" virtual evaluation, based on our previous work in evaluating multimodal sensors [19] will be incorporated. For this purpose, a virtual smart phone will be implemented. Finally, data aggregation incorporating human behavioral study results will be used. The crowdsourcing data will provide a rich repository for human behavioral study, which the results of this study will in turn incorporate in our data aggregation algorithms.

III. System Designs

We designed a crowdsourcing approach to multimedia data sharing and services to the navigation of visually impaired. The goal of the work is to provide crowd services that are user accessible (especially for visually impaired), flexible (with friendly HCI and APIs for the ease of plugging in new apps to motivate online volunteers for their services), and efficient (near real time response, and a balanced workload between mobile phone, the back end system, and the different types of users). In our research, we use the onboard sensors of a COTS smartphone (iPhone or Android Phone), such as camera, compass, GPS, and accelerometer, to assist the navigation of a blind user. The basic function of the mobile computing is to stream the video and other sensory information to the crowd server so that volunteers can use the information to provide service. In assisted navigation for the blind, volunteers send back their feedback via voice or typing and the crowd program combines the results to provide the final feedback to the blind user, through voice, vibration, or the combination of them, depending on what the tasks are.

In cases where there are more than one volunteer, each of them might possibly give a different instruction to the blind user. Some of the instruction might also be from machine vision algorithms that provide direction information. Here, we want to aggregate all the available instructions into a single one that will be returned to the blind user. Among others, we need to consider the expertise, the reliability and the reputation (low reputation might indicate noisy, or even malicious volunteer, which we want to filter) in the aggregation process. In addition, we also need to consider asynchronocity of different volunteers - the frequency of instruction update.

In addition to simply providing a service to the visually impaired, this work presents several avenues for future research. Most notably is the integration of other vision algorithms. Most notably, vision algorithms can easily be considered as another member of the crowd providing additional feedback. This algorithm helps guide the user, but also the algorithm itself can now easily be tested for accuracy against human volunteers. Along with, the on-line process and data aggregation in the above and service evaluation with users, an offline analysis will in turn help better tailor our context-aware human-computer interfaces and further improve the online analysis tasks.

We have set up a simple crowd navigation testing site (http://crowd-navigation.appspot.com/) using the Google App Engine platform in conjunction with a media server, showing that it is possible to use real-time video streaming (we use TokBox [http://tokbox.com/]) to assist a blind user to navigate by real-time feedback from a volunteer online. In the following, we will discuss two main components: wedapp design and data aggregation.

Towards answering these questions, we begin to isolate the problems by simulating a blind user (virtual avatar) navigating in a maze while streaming the view in front of the avatar. The streamed data will be sent to the online portal where a group of volunteers will issue directional commands to the virtual avatar. The virtual environments will allow us to dynamically vary the number of volunteers and observe how frequent volunteers update their instruction.

For all experiments, we record time to completion and accuracy (*i.e.,* did the subject bump into any obstacles?), in addition to qualitative information. A sighted person will accompany each visually impaired user for safety.

A. Webapp Design

To effectively utilize the information from the crowd, a system had to be developed which would allow the instructions from the crowd to be examined, aggregated, and feedback to the user in a timely manner. Furthermore, the system must be easy and interesting to use for the volunteers to be likely to use it.

To accomplish this, we have developed a webapp through Google App Engine. Users are able to log into the webapp using a regular Google account. The visually impaired users then have the option to create a uniquely identified video stream which we refer to as a "room". When sighted volunteers log onto the service, they can enter any of the existing rooms and provide instructions. For any given room, all the instructions from all the users are collected and aggregated via various methods (Section III-B).

Each of the crowd volunteers can then be "graded" on their input and given "points" for doing a good job or docked for malicious behavior. These points can be used in later instances to weight each user's feedback.

Having developed the application on Google App Engine presents several advantages. Among the most notable is that Google handles the management of accounts. This is important

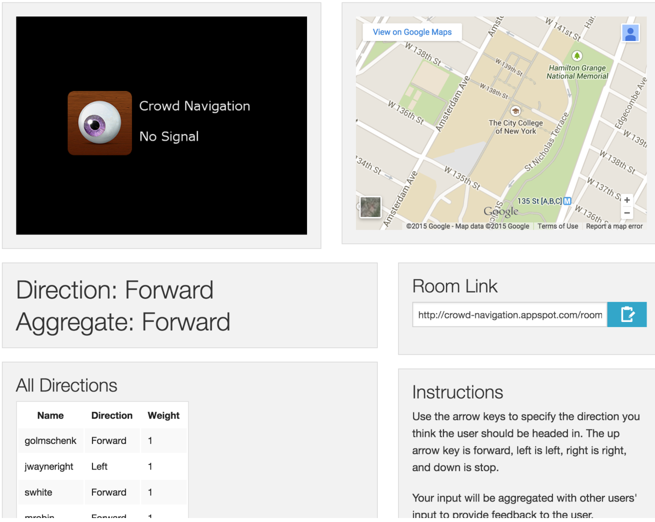

Figure 1: A screenshot of our webapp at crowd-navigation.appspot.com

because one major concern is the presence of malicious users. The use of Google accounts means that it's easier to determine which users are those who made an account specifically for malicious purposes and it is easy to ban such users. In addition, it's much easier to keep track of those users who have been most helpful and award their instructions more weight. App Engine also always allows for powerful server scaling. Should the service become popular in the future or even should there just be a surge in traffic, using App Engine allows our application to continue working without problems.

The webapp members of the crowd are presented with an interface (Figure 1) designed to allow them to best assist the VIP. A panel displays the video being streamed from the VIP's phone and location data from the phone can also be used to give a GPS position for the user on a map. The crowd member's feedback is displayed along with all other crowd members' instructions and the aggregate that is sent to the VIP.

In the current iteration of the webapp, the crowd members are only able to choose between four directions as feedback: forward, left, right, and stop. This limitation in the level of feedback is specifically chosen so that our testing can be simplified and variables can be better controlled. Once the usefulness and limitations of the system are better known, we intend to extend its features and capabilities.

On a lower level, implementation of the webapp uses a high replication non-relational database to store all user account and room data. Time sensitive data such as the guidance feedback given by the volunteers is passed through a cache on top of the database. The video streaming itself uses WebRTC to broadcast the user's footage to all crowd members.

B. Aggregation

One of the major concerns and areas of focus for our study is how the user feedback is aggregated. With a crowd of users providing instruction, we have to be careful how this information is relayed to the visually impaired user.

The naïve approach would just be to simply relay every instruction given from the crowd directly back to the visually impaired user. This of course would lead to an overwhelming amount of feedback, possibly conflicting with each other. Many of the crowd members may have different plans as to how the user should proceed and the constant changing of the instruction will be no help at all.

A more reasonable choice would be to take the aggregation of the instructions given from the crowd and send that back to the user. This way, only the primary opinion comes through to the user. There is of course the issue of how a single feedback message should be calculated from the alternatives provided by multiple volunteers. One option is to assume that all crowd members' instructions are valid. However, this may not always be true as some instructions may have been submitted after a time delay that makes them no longer relevant to the current situation. Instead, the average over a given time interval relative to the time of the user's request should be considered. Every piece of feedback will be based on the directions given in the relevant time interval. Of course, this raises its own issues as the time length and delay now plays a major role. Too short of a time interval will result in many cases where there is either no input during the interval or a single user's input is the only one considered. This leads to a problem that is almost identical to the naive approach above. On the other hand, too long of an interval means that the visually impaired user will not receive feedback until significantly after it is requested, and likely needed.

Another alternative for the aggregation is the use of a legion leader [20]. This approach still gathers all the instructions over an interval of time, but does not to send back a voted instruction to the visually impaired user. Instead, all the given instructions are considered and the crowd member who mostly closely matched the overall opinion of the crowd is chosen as the "leader" for the next time interval. The leader is given complete control during that time interval and only the leader's instructions are returned to the visually impaired user. This approach has several advantages. First, the feedback is immediately sent to the user. When the leader enters a command, there is no need to wait for the end of a time interval to send it to the user. Also, there is no problem with conflicting plans for how to proceed. One crowd member is given complete control for a short period of time, so it doesn't matter if about half the crowd think the person should go right around a pole while the other half think the user should go left. Only the person who is currently the leader picks and by the time they are no longer in control, the best choice for the route around the obstacle will probably be decided.

IV. Virtual Environment

Before we conduct real users testings, we propose to use virtual environments (VEs) to simulate VIP navigation while testing the crowd's ability in assisted guidance. While we want to simulate a realistic environment that the crowd can navigate

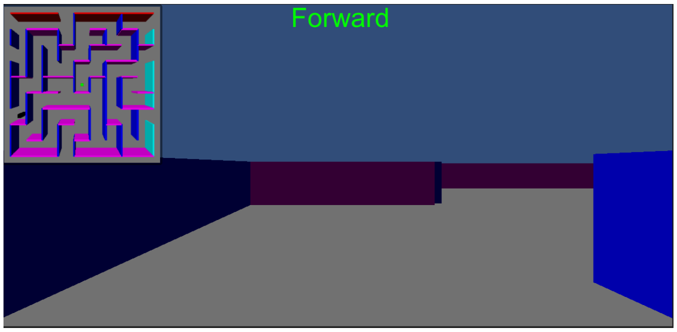

Figure 2: View of a VE.

a user in, for simplicity and a more controlled experiment, we decided to simulate random mazes instead. The purpose of this experiment is to verify the viability of the system and aggregation methods, therefore, all mazes are single floor and only contain one type of obstacles (*i.e.,* walls).

A. Environment Design

We use a game engine, Unity3D, to design five virtual environments as simple mazes with a single path from start to destination for each. Although the mazes are randomized, we had fixed it by storing the seed value, for repeatability. A view of a maze (from the viewpoint of the avatar, which is not visible) can be seen in Figure 2. This view is also streamed to the webapp (left panel of Figure 1) via a screen capturing software (*e.g.,* ManyCam). The VE also receives aggregated command from the webapp, which is displayed in the top central of the view, as well as text-to-speech. The command is received via standard HTTP GET from the crowd navigation testing site.

B. Navigation Modes

The virtual avatar can be controlled in one of two modes: 1) automatic; and 2) manual. In automatic mode, the avatar faithfully obeys the aggregated command received. If a "Forward" command is received, the avatar moves forward until a new (different) command is received. If a turning command (*i.e.,* left or right) is received, the avatar rotates in place at a constant speed in the corresponding direction until a new command is received. A "Stop" command can also be issued, which will stop the avatar from moving and rotating until a new command is received.

In manual mode, the avatar is controlled by a real user, who may or may not be VIP, via a joystick. In either case, the user will not be able to see the view and can only hear the command being spoken via text-to-speech. The user can have the command repeated if he or she missed it. The user may have taken the wrong turn, in this case, the crowd will have to correct the user.

C. Data Collection

For each maze and trial run, we collected the avatar's position and orientation. We also collected the command that was received from the webapp/crowd and the response the user made with regard to it (the automatic avatar always respond to the command correctly). Furthermore, we recorded how many times the avatar bumped into wall, whether it's contributed by the crowd or by the user. Lastly, we timed each run.

Data is also recorded on the webapp side. Data such as each individual crowd member's response, the aggregation method (*i.e.,* simple sum or legion leader), the aggregated response that was sent to the VE, the "leader" (only applicable for legion leader method), and timestamp. Table I summarizes the variables for which data were collected.

Independent	Aggregation method	Crowd size	
Dependent	Completion time	Number of errors	Shortest distance to destination, if timed-out

Table I: Independent and dependent variables.

V. RESULTS

We have conducted two sets of experiment. In the first experiment, we generated five different mazes for testing. 16 crowd volunteers participated in this set of tests. For each maze, we recorded the completion time and the number of times the crowd directed virtual agent came into contact with the walls. The ground truth times were obtained by the experimenters navigating the mazes locally (i.e. without going through the website aggregation) and without inputs from the crowd. Table II shows the crowd completion and ground truth times for all five mazes. The results show that the crowd completion time is significantly different from the ground truth time (Two-sample t-test, p=0.0009 at 5% significance level, df=8).

Figure 3 shows the trajectories of the crowd in mazes 1 – 3, using a simple sum method. While the avatar in mazes 1 and 3 moves exactly to the crowd's response, maze 2's avatar is controlled by another user. The user navigates the maze based on the crowd's feedback. This user was able to complete the maze without bumping into any walls. During this first set of tests, network issues resulted in blurry, broken video streaming. With the entire group of crowd members being on the same network, the issue caused significant problems in this set of tests, particularly in the later mazes. For these mazes, the wall contacts could be attributed to slow/delayed network and low resolution of the video stream. The end-of-experiment survey also reflects this opinion.

For the second experiment, we constructed mazes with larger wall/objects such that even when streamed in low resolution, the view is easy to see. We also worked to ensure proper network connections by streaming the view from a high-bandwidth location, slowing down the automatic avatar's speed, and increasing some objects' size (*e.g., walls*).

Maze #	1	2	3	4	5
Crowd time (s)	513.94	345.47	325.00	258.87	505.94
Truth time (s)	108.01	102.97	159.24	112.40	142.01

Table II: Crowd vs ground truth times in first experiment.

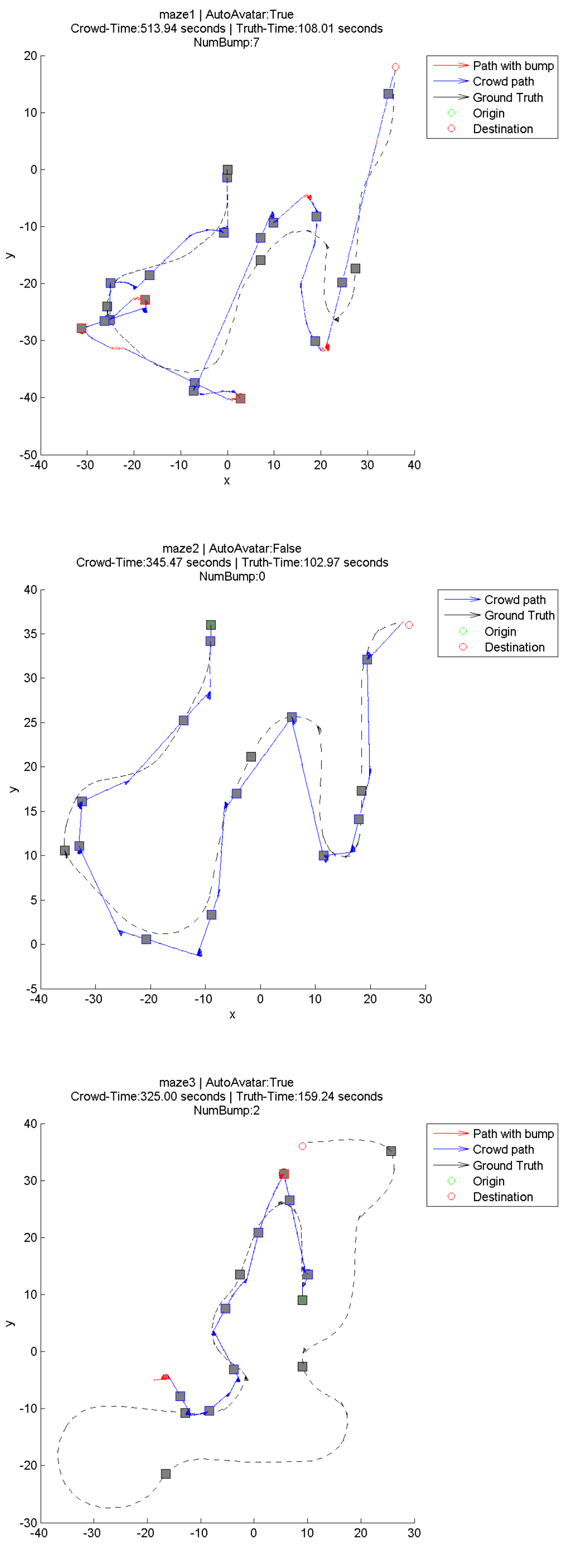

Figure 3: Top down view of mazes 1 – 3 in first experiment, showing crowd paths, paths with bump, ground truth paths, origin, destination, and 30-second interval (gray squares)

Maze #	1	2	3	4
Simple sum time (s)	221.64	180.27	292.86	322.79
Legion leader time (s)	219.65	182.41	263.50	228.89

Table III: Crowd times in second experiment.

Furthermore, this experiment also included the legion leader aggregation method. 11 volunteers participated in the crowd for this experiment. The number of volunteers in this experiment differ from the previous is because we want to see how the number of volunteers affect the aggregation methods and thus the performance of the user/avatar navigating through the mazes.

Table III shows the crowd completion times for both simple sum and legion leader aggregation methods. Maze 4's avatar is controlled by another user. The user navigates the maze based on the crowd's feedback. This user was able to complete the maze faster using the legion leader aggregation method. Although the result shows that crowd completion time for either aggregation method is not significantly different (two-sample t-test, p=0.432 at 5% significance level, df=6), the crowd time in Table II vs. simple average crowd time in Table III is significantly different (two-sample t-test, p=0.074 at 10% significance level, df=7). The number of mazes in this experiment differ from the previous is because of a subject withdrawal from the experiment.

Figure 4 shows the trajectories of the crowd in mazes 1 – 3, using a legion leader method. The avatars in mazes 1 to 3 moves exactly to the crowd's response. The overall trajectories in this experiment are much smoother than those in Figure 3. The crowd was also able to navigate the avatar to its destination in a shorter time. The overall feedback from users are positive and much improved from the first experiment. The crowd was able to clearly see the mazes and navigate the avatar without much of network delay.

VI. Conclusions and Discussions

This paper describes the design, implementation, and testing of crowd-based navigation system for the visually impaired. The system can be easily implemented at low-cost using a standard smartphone or tablet device that is capable of streaming data. In particular, an evaluation of the system is performed using virtual environments. A series of mazes are designed that allows the crowd to directly control the avatar or give feedback to a user, who in turn control the avatar. Although the crowd is consists of more than one member, only one aggregated response is fedback to the virtual environment and user. This type of crowd-based navigation may provide a useful form of assistance for individuals with visual impairments.

We have conducted preliminary experiments. The first experiment was conducted with prototype of mazes. User feedback indicated that the network connection was slow and thus, resulted in blurry, broken video streaming. With the second experiment, we addressed the above issues and the experience was more positive. We also shown that the second

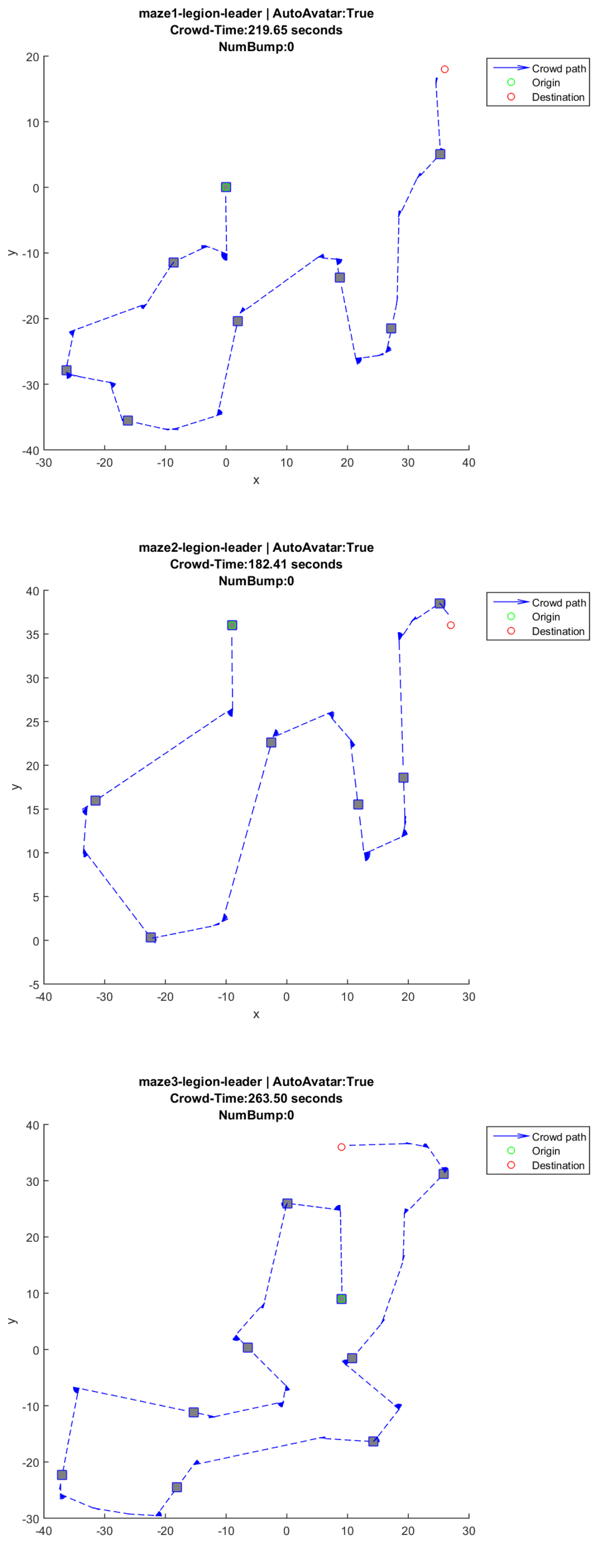

Figure 4: Top down view of mazes 1 – 3 in second experiment, showing crowd paths, paths with bump, ground truth paths, origin, destination, and 30-second interval (gray squares)

experiment has a better result, in term of completion time. Furthermore, we tested two aggregation methods in the second experiment: simple sum and legion leader. However, we do not have enough samples to prove one method is better than the other. Lastly, we had informally tested the smartphone app for crowd-based navigation where we have a blindfolded lab member walked from one room to another room on the same floor based on the aggregated responses provided by the crowd, where the crowd is consists of the rest of the lab members. We repeated this with different lab members.

We are planning to conduct more experiments, even large-scale ones, involving visually impaired people, and more realistic environments (both real and virtual) and conditions. We also plan to extend the experiments with a questionnaire investigating about the understandability of the feedback to the visually impaired, its timeliness, and the degree of confidence it provides to the user. The ultimate goal of this project is to provide an online service for the visually impaired. Specifically, we want to have a sustainable platform where we will have perpetual crowd members providing navigation directions to safely guide a visually impaired person to reach his/her destination. Other possible applications of this system can include, but not limited to, scenarios where people may get temporarily blind (*e.g.,* firefighters surrounded by smoke, or soldiers in the battlefield at night). Regardless, the main challenges of obtaining good video quality and critical mass crowd remain.

VII. Acknowledgments

This work is supported in part by National Science Foundation (EFRI Award # 1137172), VentureWell (formerly NCIIA, through Award # 10087-12), and a PSC-CUNY Award. We would like to thank Ms. Barbara Campbell and Mr. Masud Kalim at the NYS Commission for the Blind for the assistance in testing our apps and providing valuable comments and feedback. We also thank senior design students at the City College of New York in helping the crowd navigation tests.

References

[1] J. Coughlan, R. M, and H. Shen, "Cell phone-based wayfinding for the visually impaired," in *st Int. Workshop on Mobile Vision*, 2006.

[2] J. Coughlan and R. Manduchi, "A Mobile Phone Wayfinding System for Visually Impaired Users," *Assist technol Res Ser*, vol. 25, no. 2009, p. 849, 2009.

[3] V. N. Murali and J. M. Coughlan, "SMARTPHONE-BASED CROSS-WALK DETECTION AND LOCALIZATION FOR VISUALLY IMPAIRED PEDESTRIANS," *IEEE Int Conf Multimed Expo Workshops*, vol. 2013, pp. 1–7, Jul 2013.

[4] H. Pan, C. Yi, and Y. Tian, "A primary travelling assistant system of bus detection and recognition for visually impaired people," in *Multimedia and Expo Workshops (ICMEW), 2013 IEEE International Conference on*, July 2013, pp. 1–6.

[5] S. Wang, H. Pan, C. Zhang, and Y. Tian, "Rgb-d image-based detection of stairs, pedestrian crosswalks and traffic signs," *Journal of Visual Communication and Image Representation*, vol. 25, no. 2, pp. 263 – 272, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1047320313001934

[6] B. Chidester and M. Do, "Assisting the visually impaired using depth inference on mobile devices via stereo matching," in *Multimedia and Expo Workshops (ICMEW), 2013 IEEE International Conference on*, July 2013, pp. 1–6.

[7] E. Molina, A. Diallo, and Z. Zhu, "Visual noun navigation framework for the blind," *Journal of Assistive Technologies*, vol. 7, no. 2, pp. 118–130, 2013. [Online]. Available: http://dx.doi.org/10.1108/17549451311328790

[8] H. Tang, T. Ro, and Z. Zhu, "Smart sampling and transducing 3d scenes for the visually impaired," in *Multimedia and Expo Workshops (ICMEW), 2013 IEEE International Conference on*, July 2013, pp. 1–4.

[9] S. Wang, X. Yang, and Y. Tian, "Detecting signage and doors for blind navigation and wayfinding," *Network Modeling Analysis in Health Informatics and Bioinformatics*, vol. 2, no. 2, pp. 81–93, 2013. [Online]. Available: http://dx.doi.org/10.1007/s13721-013-0027-9

[10] C. Yi and Y. Tian, "Text extraction from scene images by character appearance and structure modeling," *Comput. Vis. Image Underst.*, vol. 117, no. 2, pp. 182–194, Feb. 2013. [Online]. Available: http://dx.doi.org/10.1016/j.cviu.2012.11.002

[11] "Blindsquare," http://blindsquare.com, accessed: March 2015.

[12] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, S. White, and T. Yeh, "Vizwiz: Nearly real-time answers to visual questions," in *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '10. New York, NY, USA: ACM, 2010, pp. 333–342. [Online]. Available: http://doi.acm.org/10.1145/1866029.1866080

[13] J. Miele, "Dvx: The descriptive video exchange," *The 27th Annual International Technology and Persons with Disabilities*, 2013.

[14] R. Snow, B. O'Connor, D. Jurafsky, and A. Y. Ng, "Cheap and fast—but is it good?: Evaluating non-expert annotations for natural language tasks," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, ser. EMNLP '08. Stroudsburg, PA, USA: Association for Computational Linguistics, 2008, pp. 254–263. [Online]. Available: http://dl.acm.org/citation.cfm?id=1613715.1613751

[15] P. Welinder, S. Branson, S. Belongie, and P. Perona, "The multidimensional wisdom of crowds," in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds., 2010, pp. 2424–2432.

[16] V. S. Sheng, F. Provost, and P. G. Ipeirotis, "Get another label? improving data quality and data mining using multiple, noisy labelers," in *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '08. New York, NY, USA: ACM, 2008, pp. 614–622. [Online]. Available: http://doi.acm.org/10.1145/1401890.1401965

[17] V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy, "Learning from crowds," *J. Mach. Learn. Res.*, vol. 11, pp. 1297–1322, Aug. 2010. [Online]. Available: http://dl.acm.org/citation.cfm?id=1756006.1859894

[18] D. Zhou, J. Platt, S. Basu, and Y. Mao, "Learning from the wisdom of crowds by minimax entropy," in *Advances in Neural Information Processing Systems (NIPS)*, December 2012. [Online]. Available: http://research.microsoft.com/apps/pubs/default.aspx?id=175659

[19] W. Khoo, J. Knapp, F. Palmer, T. Ro, and Z. Zhu, "Designing and testing wearable range-vibrotactile devices," *Journal of Assistive Technologies*, vol. 7, no. 2, pp. 102–117, 2013. [Online]. Available: http://dx.doi.org/10.1108/17549451311328781

[20] W. S. Lasecki, K. I. Murray, S. White, R. C. Miller, and J. P. Bigham, "Real-time crowd control of existing interfaces," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '11. New York, NY, USA: ACM, 2011, pp. 23–32. [Online]. Available: http://doi.acm.org/10.1145/2047196.2047200