A Pilot Deployment of an Online Tool for Large-Scale Virtual Auditing of Urban Accessibility

Manaswi Saha[1], Kotaro Hara[1,2], Soheil Behnezhad[1], Anthony Li[1],
Michael Saugstad[1], Hanuma Maddali[1], Sage Chen[3], Jon E. Froehlich[1]
[1]University of Maryland; [2]Carnegie Mellon University, [3]University of Michigan
{manaswi, soheil, saugstad, hmaddali, antli, jonf}@umd.edu; kotaroh@cmu.edu; sagec@umich.edu

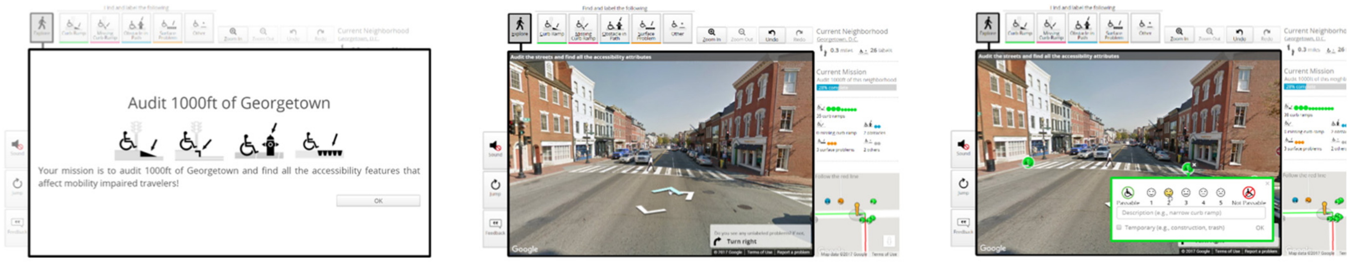

(a) Example mission screen (b) Virtually walk and explore city streets guided by mission map (c) Find and label accessibility features and problems

Figure 1. *Project Sidewalk* uses basic game design principles to engage online volunteers in virtually auditing urban accessibility in Google Street View.

ABSTRACT

Project Sidewalk is an online tool that allows anyone—from motivated citizens to government workers—to remotely label accessibility problems by virtually walking through city streets. Basic game design principles such as interactive onboarding, mission-based tasks, and stats dashboards are used to train, engage, and sustain users. We describe the current Project Sidewalk system, present results of a pilot public deployment with 581 users, and discuss open questions and future work.

CCS Concepts

• **Human-centered computing → Accessibility technologies**

Keywords

Urban accessibility; crowdsourcing; mobility impaired users; GIS

1. INTRODUCTION

Geographic Information Systems (GIS) such as Google Maps, Waze, and Yelp have transformed the way people travel and access information about the physical world. While these systems contain terabytes of data about road networks and points of interest (POIs), their information about *physical accessibility* is commensurately poor. GIS websites such as Axsmap.com, Wheelmap.org, and AccessTogether.org aim to address this problem by collecting location-based accessibility information provided by volunteers (*i.e.,* crowdsourcing). While these efforts are important and commendable, their value propositions are intrinsically tied to the amount and quality of provided data. In a survey, Ding *et al.* [3] found that most accessibility-oriented GIS sites suffered from serious data sparseness issues (*e.g.,* only 1.6% of the POIs in Wheelmap had data entered on accessibility). One key problem is that these sites rely on local populations with physical experience of a place, which dramatically limits who can supply data.

In contrast, we are exploring a different approach embodied in a new interactive tool called *Project Sidewalk*, which enables volunteers to contribute physical-world accessibility information by *virtually* visiting a location in Google Street View (Figure 1). Rather than pulling solely from a local population, our potential pool of users scales to anyone with an Internet connection and a web browser. Project Sidewalk extends previous work in virtual auditing tools, including *CANVAS* [1], *Spotlight* [2], *Bus Stop CSI* [5], and *Tohme* [7], which demonstrates that tool-mediated virtual audits of urban infrastructure have high concordance with traditional physical audits. However, these tools were not publicly deployed, were focused on small spatial regions, and relied on specialized user populations (*e.g.,* public health researchers [1, 2]) or paid crowdworkers [5, 7]. As Project Sidewalk is designed to engage and train non-expert volunteers in a complex task—that is, accessibility audits—our work is also strongly linked to and additionally informed by work in citizen science [4], which similarly seeks to engage non-experts in collecting and processing data through novel interactive tools (*e.g.,* Zooniverse.org).

In submission to ASSETS'17 poster track.

As initial work, our research questions are exploratory: how can we engage, train, and sustain volunteer users in virtual accessibility audits? What types of users participate (*e.g.,* GIS enthusiasts, government workers, concerned citizens, or perhaps those directly impacted by accessibility problems)? How do volunteers differ in labeling behavior and quality compared to experts? Is it feasible to rely purely on volunteers or must data be supplemented by paid crowd workers? To begin addressing these questions, we report on an initial pilot deployment of Project Sidewalk in Washington DC. We first describe the Project Sidewalk system before providing a high-level analysis of the data we have collected thus far. We close with a discussion of future work plans.

2. PROJECT SIDEWALK

Project Sidewalk culminates five years of research in virtual auditing interfaces for accessibility (see [6] for a history). We build on the labeling interfaces in [5, 7] but extend our system to: (i) interactively train volunteers about the tool and urban accessibility; (ii) incorporate basic game design principles such as missions and progress tracking to improve user engagement; and (iii) allow open-world exploring (similar to first-person video games) to increase street and sidewalk coverage. Such changes have required a substantial engineering effort: the Project Sidewalk open-source github repository has 2,155 commits from 20 team members and 118,207 lines of developed code in Scala, Javascript, and Python.

To use Project Sidewalk, visit http://projectsidewalk.io on a laptop or desktop (touchscreens are not currently supported). Upon clicking the '*Start Mapping*' button, new users are greeted by a multi-stage interactive tutorial to learn about the user interface and basic accessibility concepts. Once the tutorial is complete, users are auto-assigned a neighborhood in DC and are given their first mission (Figure 1a). Missions route users through specific streets in a neighborhood; as the user walks virtually along this path, they find and label street and sidewalk accessibility features (curb ramps) and problems (sidewalk obstacles, missing curb ramps, no sidewalks, *etc.*). After identifying a problem, the user is asked to provide a severity rating on a scale from '1' to '5' where '5' represents an impassable barrier for someone in a manual wheelchair. Once a mission is complete, users are greeted by a 'mission complete' screen along with a 'cheering' sound effect. Users are then assigned a new mission and the cycle continues. Users can choose to contribute anonymously or to register and login. Registered users can check their contribution activity on an analytics dashboard. Currently, however, there is no way to view or compare other users' contributions.

3. PILOT DEPLOYMENT

We launched a pilot deployment in September 2016, which was advertised through word-of-mouth and social media. Thus far, 581 users contributed 71,873 accessibility labels across 504 miles of DC streets (nearly 50% of all streets in the city)—see Figure 2. On average, users completed 3.6 missions (*SD*=13.8), which accounts for an average of 2.9 miles audited/user (*Med*=1.0; *SD*=8.8). As mission lengths vary, to examine auditing time, we calculated the average time to complete 1,000ft (~305 meters): 16.9 mins (*SD*=11.2). Of the 581 users, 315 (54.2%) chose to register and login. These registered users contribute a majority of labels: 58,887 *vs.* 12,986 for anonymous. In terms of the labels themselves, *curb ramps* comprise 64.3% of our data, followed by *no sidewalks* (17%), *missing curb ramps* (7.2%), *obstacles in path* (5.6%), and *surface problems* (4.9%)—see Table 1. These labeling statistics are raw counts and contain redundancies as users can label a problem multiple times from different viewpoints. Future post-processing is necessary to aggregate redundant labels into individual issues.

Labels	Total	Anon. Only	Avg. Severity (SD)	Labels/mile
Curb Ramp	46,213 (64.3%)	9,182	1.2 (0.5)	91.6
No Sidewalk	12,230 (17.0%)	1,293	N/A	24.3
Missing Curb Ramp	5,191 (7.2%)	1,285	3.8 (1.1)	10.3
Obstacle in Path	4,044 (5.6%)	673	3.1 (1.3)	8.0
Surface Problem	3,552 (4.9%)	465	2.3 (1.2)	7.0
Other	643 (0.9%)	16.1%	N/A	4.7
Total	71,873	12,986		1.3

Table 1. A breakdown of data collected in the pilot deployment. Label severity is a scale from '1' to '5' where '5' is worst.

We have also begun analyzing qualitative comments from email, social media, and our tool's 'feedback' form. Generally, responses have been positive: multiple city DOTs have inquired about deploying in their cities and ADA specialists have provided helpful suggestions. Reported problems include not always understanding *where* to apply labels (*e.g., "should driveways be labeled as curb ramps?*"), ambiguity between label types (*e.g.,* some problems could legitimately be labeled both an 'obstacle in path' and a 'surface problem), and technical issues (*e.g.,* 24% of our site visitors used unsupported browsers).

4. DISCUSSION AND CONCLUSION

While our pilot deployment demonstrates the potential of Project Sidewalk to train and use non-expert volunteers to virtually audit urban infrastructure for accessibility, there is much work to be

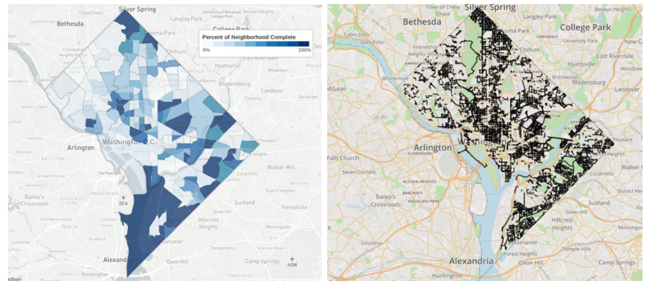

Figure 2. (a) A choropleth of Washington DC showing neighborhood completion rates (*Avg completion rate*=54.7%; *SD*=29.7%) and (b) a map showing the 504 miles of audited DC streets (in black).

done. We have three major goals: (i) to develop a robust, reusable platform for crowdsourcing urban accessibility data that can be deployed in any city in the world that has Google Street View; (ii) to increase scalability using machine learning (extending [7]); (iii) to support the development of novel accessibility-aware GIS applications that leverage Project Sidewalk's unprecedented data (http://projectsidewalk.io/api).

Towards (i), Ding *et al.*'s [3] survey of GIS accessibility sites focused on data completeness rather than quality—both are important. While we have not yet analyzed the accuracy of our labels, our prior work [6, 7] suggests 80-90% correctness. We are currently designing verification interfaces to help analyze and improve labeling accuracy. Once complete, we plan to launch a full deployment in Washington DC followed by 3-5 more U.S. cities, which will allow us to compare accessibility characteristics across regions. Towards (iii), we have begun building a suite of accessibility-aware GIS tools that leverage Project Sidewalk data such as *AccessScore*, which provides an at-a-glance view of neighborhood accessibility, *AccessRoute* that personalizes route suggestions based on a user's mobility level, and *AccessModel* that explores accessibility correlates with socio-economic, census tract, and other location-specific data. Finally, we are exploring hybrid volunteer + paid worker approaches. We estimate, for example, that auditing a city as large as DC with Project Sidewalk on Mechanical Turk at $7.25/hr (US minimum wage) would cost around $14,000 (including the 20% Amazon commission). In comparison, a traditional audit costs orders of magnitude more.

5. REFERENCES

[1] Bader, M.D.M., Mooney, *et al.* 2015. Development and deployment of the Computer Assisted Neighborhood Visual Assessment System (CANVAS) to measure health-related neighborhood conditions. *Health & Place*. 31, (Jan. 2015), 163–172.

[2] Bethlehem, J.R., Mackenbach, J.D., *et al.* 2014. The SPOTLIGHT virtual audit tool: a valid and reliable tool to assess obesogenic characteristics of the built environment. *International Journal of Health Geographics*. 13, 1 (Dec. 2014), 52.

[3] Ding, C., Wald, M. and Wills, G. 2014. A Survey of Open Accessibility Data. *Proc. of W4A* (NY, USA, 2014), 37:1--37:4.

[4] Haklay, M. 2013. Citizen Science and Volunteered Geographic Information: Overview and Typology of Participation. *Crowdsourcing Geographic Knowledge: VGI in Theory and Practice*. D. Sui, S. Elwood, and M. Goodchild, eds. Springer. 105–122.

[5] Hara, K., Azenkot, S., Campbell, M., *et al.* 2015. Improving Public Transit Accessibility for Blind Riders by Crowdsourcing Bus Stop Landmark Locations with Google Street View: An Extended Analysis. *ACM Trans. Access. Comput.* 6, 2 (Mar. 2015), 5:1--5:23.

[6] Hara, K. & Froehlich, J. 2015. Characterizing & visualizing physical world accessibility at scale using crowdsourcing, computer vision, & machine learning. *ACM SIGACCESS Newsletter*. 113, 13–21.

[7] Hara, K., Sun, J., Jacobs, D.W. and Froehlich, J.E. 2014. Tohme: Detecting Curb Ramps in Google Street View Using Crowdsourcing, Computer Vision, and Machine Learning. *Proc. of UIST'14,* 189-204.