COMBINING CROWDSOURCING AND MACHINE LEARNING TO COLLECT SIDEWALK ACCESSIBILITY DATA AT SCALE

FINAL PROJECT REPORT

by
Jon E. Froehlich
University of Washington

Sponsorship
PacTrans and the UW Allen School of Computer Science and Engineering

for

Pacific Northwest Transportation Consortium (PacTrans)
USDOT University Transportation Center for Federal Region 10
University of Washington
More Hall 112, Box 352700
Seattle, WA 98195-2700

In cooperation with U.S. Department of Transportation,
Office of the Assistant Secretary for Research and Technology (OST-R)

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The Pacific Northwest Transportation Consortium, the U.S. Government and matching sponsor assume no liability for the contents or use thereof.

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
	1723937	

4. Title and Subtitle Combining Crowdsourcing and Machine Learning to Collect Sidewalk Accessibility Data at Scale	**5. Report Date** June 9, 2021 **6. Performing Organization Code**
7. Author(s) and Affiliations Jon E. Froehlich 0000-0001-8291-3353 Allen School of Computer Science and Engineering University of Washington	**8. Performing Organization Report No.** 2019-S-UW-2
9. Performing Organization Name and Address PacTrans Pacific Northwest Transportation Consortium University Transportation Center for Federal Region 10 University of Washington More Hall 112 Seattle, WA 98195-2700	**10. Work Unit No. (TRAIS)** **11. Contract or Grant No.** 69A3551747110
12. Sponsoring Organization Name and Address United States Department of Transportation Research and Innovative Technology Administration 1200 New Jersey Avenue, SE Washington, DC 20590	**13. Type of Report and Period Covered** Final Report, May 2019 - 2021 **14. Sponsoring Agency Code**

15. Supplementary Notes

Report uploaded to: www.pactrans.org

16. Abstract

We are developing new data collection approaches that use a combination of remote crowdsourcing, machine learning, and online map imagery. Our newest effort, called *Project Sidewalk*, enables online crowdworkers to remotely label pedestrian-related accessibility problems by *virtually* walking through city streets in Google Street View. In 2019, we completed an 18-month deployment in Washington, D.C.: 1,150+ users provided over 200,000 geo-located sidewalk accessibility labels and audited 3,000 miles of D.C. streets. With simple quality control mechanisms, we found that minimally trained remote crowd workers could find and label 92 percent of accessibility problems in street view scenes, including *missing curb ramps*, *obstacles in the path*, *surface problems*, and *missing sidewalks*. For our PacTrans project, we proposed three threads of additional work. **(1) First**, we are deploying Project Sidewalk into three more cities, including two in the Pacific Northwest: Seattle, Washington, and Newberg, Oregon, to enable us to study and compare sidewalk accessibility factors across cities. **(2) Second**, to further scale our approach, we proposed new methods to automatically identify and classify sidewalk problems using deep learning techniques, which would be uniquely enabled by our large dataset. **(3) Finally**, we proposed new sidewalk accessibility models and interactive visualization tools to give stakeholders—from citizens to transit authorities—new understanding of their city's accessibility.

17. Key Words		18. Distribution Statement	
A: Transportation; E: Planning and design; J: Environment			
19. Security Classification (of this report)	**20. Security Classification (of this page)**	**21. No. of Pages**	**22. Price**
Unclassified.	Unclassified.	30	N/A

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized.**

SI* (MODERN METRIC) CONVERSION FACTORS

APPROXIMATE CONVERSIONS TO SI UNITS				
Symbol	**When You Know**	**Multiply By**	**To Find**	**Symbol**
		LENGTH		
in	inches	25.4	millimeters	mm
ft	feet	0.305	meters	m
yd	yards	0.914	meters	m
mi	miles	1.61	kilometers	km
		AREA		
in^2	square inches	645.2	square millimeters	mm^2
ft^2	square feet	0.093	square meters	m^2
yd^2	square yard	0.836	square meters	m^2
ac	acres	0.405	hectares	ha
mi^2	square miles	2.59	square kilometers	km^2
		VOLUME		
fl oz	fluid ounces	29.57	milliliters	mL
gal	gallons	3.785	liters	L
ft^3	cubic feet	0.028	cubic meters	m^3
yd^3	cubic yards	0.765	cubic meters	m^3
		NOTE: volumes greater than 1000 L shall be shown in m^3		
		MASS		
oz	ounces	28.35	grams	g
lb	pounds	0.454	kilograms	kg
T	short tons (2000 lb)	0.907	megagrams (or "metric ton")	Mg (or "t")
		TEMPERATURE (exact degrees)		
°F	Fahrenheit	5 (F-32)/9 or (F-32)/1.8	Celsius	°C
		ILLUMINATION		
fc	foot-candles	10.76	lux	lx
fl	foot-Lamberts	3.426	candela/m^2	cd/m^2
		FORCE and PRESSURE or STRESS		
lbf	poundforce	4.45	newtons	N
lbf/in^2	poundforce per square inch	6.89	kilopascals	kPa

APPROXIMATE CONVERSIONS FROM SI UNITS				
Symbol	**When You Know**	**Multiply By**	**To Find**	**Symbol**
		LENGTH		
mm	millimeters	0.039	inches	in
m	meters	3.28	feet	ft
m	meters	1.09	yards	yd
km	kilometers	0.621	miles	mi
		AREA		
mm^2	square millimeters	0.0016	square inches	in^2
m^2	square meters	10.764	square feet	ft^2
m^2	square meters	1.195	square yards	yd^2
ha	hectares	2.47	acres	ac
km^2	square kilometers	0.386	square miles	mi^2
		VOLUME		
mL	milliliters	0.034	fluid ounces	fl oz
L	liters	0.264	gallons	gal
m^3	cubic meters	35.314	cubic feet	ft^3
m^3	cubic meters	1.307	cubic yards	yd^3
		MASS		
g	grams	0.035	ounces	oz
kg	kilograms	2.202	pounds	lb
Mg (or "t")	megagrams (or "metric ton")	1.103	short tons (2000 lb)	T
		TEMPERATURE (exact degrees)		
°C	Celsius	1.8C+32	Fahrenheit	°F
		ILLUMINATION		
lx	lux	0.0929	foot-candles	fc
cd/m^2	candela/m^2	0.2919	foot-Lamberts	fl
		FORCE and PRESSURE or STRESS		
N	newtons	0.225	poundforce	lbf
kPa	kilopascals	0.145	poundforce per square inch	lbf/in^2

*SI is the symbol for the International System of Units. Appropriate rounding should be made to comply with Section 4 of ASTM E380. (Revised March 2003)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ADA:	Americans with Disabilities Act
AI:	Artificial intelligence
CSCW:	Computer-supported cooperative work
GSV:	Google Street View
HCI:	Human-computer interaction
MI:	Mobility impairments
ML:	Machine language
PacTrans:	Pacific Northwest Transportation Consortium
POI:	Point of interest
ResNet:	Residual neural network
WSDOT:	Washington State Department of Transportation

ACKNOWLEDGMENTS

We thank the Project Sidewalk team, including Michael Saugstad, Manaswi Saha, Ather Sharif, PaariGopal, Aroosh Kumar, Michael Duan, and many others.

EXECUTIVE SUMMARY

We are developing new data collection approaches that use a combination of remote crowdsourcing, machine learning, and online map imagery. Our newest effort, called *Project Sidewalk*, enables online crowdworkers to remotely label pedestrian-related accessibility problems by *virtually* walking through city streets in Google Street View. Rather than pulling solely from local populations, our potential pool of users scales to anyone with an Internet connection. In 2019, we completed an 18-month deploymentin Washington, D.C.: 1,150+ users provided over 200,000 geo-located sidewalk accessibility labels and audited 3,000 miles of D.C. streets. With simple quality control mechanisms (*e.g.*, majority vote), we found that minimally trained remote crowdworkers could find and label 92 percent of accessibility problems instreet view scenes, including *missing curb ramps*, *obstacles in the path*, *surface problems*, and *missing sidewalks*.

Building on this momentum, our original PacTrans proposal enumerated three threads of additional work:

1. **First**, we intended to deploy Project Sidewalk into three more cities, including two in the Pacific Northwest: Seattle, Washington, and Newberg, Oregon. These additional deployments would enable us to study new crowdsourcing workflows and user interfaces, study new methods for inferring data quality, and address questions about urban accessibility not previously possible. For example, what factors correlate with urban accessibility (*e.g.,* real estate pricing, socio-economics, zoning)? How does urban accessibility compare across cities and what are the key contributing factors?
2. **Second**, to further scale our approach, we proposed new methods to automatically identify and classify sidewalk problems using deep learning techniques, which would

be uniquely enabled by our large dataset. Key questions here included how does the choice of deep learning model and input features impact performance? How can we use these machine learning techniques to validate user-supplied labels or even find accessibility problems in street view panoramas automatically?

3. **Finally**, we proposed examining new sidewalk accessibility models and interactive visualization tools to give stakeholders—from citizens to transit authorities—new understanding of their city's accessibility.

In this report, we summarize progress along each thread and describe key contributions.

CHAPTER 1. PROJECT GOALS AND RESEARCH THREADS

Sidewalks significantly impact the mobility and quality of life of millions of Americans (National Council on Disability, 2007; U.S. Census Bureau, 2012). In this report, we describe new, scalable methods for collecting data on sidewalk accessibility that use machine learning, crowdsourcing, and online map imagery as well as new interactive visualizations aimed at providing novel insights into urban accessibility. As with our previous research (*e.g.*, (Hara, et al., 2016; Hara, et al., 2013; Saha et al., 2019)), we intended to work closely with key stakeholders, including local governments and transit departments, mobility-impaired individuals and caretakers, and walkability advocates to help shape and evaluate the design of our tools. While our proposed techniques and tools should work anywhere with *OpenStreetMaps* and available streetscape imagery (*e.g.*, *Google Street View*, *Mapillary*), two of our three immediate deployment targets were cities in the PacTrans region: Newberg, Oregon, and Seattle, Washington.

Our proposed work has the potential to change how citizens, urban planners, accessibility advocates, and local governments understand and analyze the accessibility of sidewalk infrastructure, whether to help plan, manage, and track improvements or to inform travel and living decisions. Accessible sidewalks not only benefit people with disabilities, but they are also generally safer and more user-friendly for all pedestrians (McMillen et al., 1999). For example, people carrying packages, pulling luggage, or pushing children in strollers are also affected by inaccessible streets and sidewalks. The overarching goal of our work is to transform the way in which sidewalk accessibility information is collected and visualized to ultimately improve pedestrian infrastructure and how people move about a city.

Below, we first provide research background before detailing a specific problem statement. We then describe the three core research threads of our original proposal.

1.1. Research Background and Problem Statement

Over 30 million U.S. adults have a physical disability that impacts their ambulatory activity (U.S. Census Bureau, 2012). Of these, nearly half report using an assistive aid such as a wheelchair (3.6 million) or a cane, crutches, or walker (11.6 million) (U.S. Census Bureau, 2012). Despite comprehensive civil rights legislation for Americans with disabilities (*e.g.*, (3rd Circuit, 1993; United States Department of Justice Civil Rights Division, 1990)), many city streets, sidewalks, and businesses in the U.S. remain inaccessible (Hu, n.d.; National Council on Disability, 2007)—including in the Pacific Northwest (Friedman, 2018; Gutman, 2017). The problem is not just that sidewalk accessibility fundamentally affects *where* and *how* people travel in cities but also that there are few, if any, mechanisms to determine accessible areas of a city. Indeed, the *National Council on Disability* noted that it could not find comprehensive information on the "degree to which sidewalks are accessible" across the U.S.

The lack of street-level accessibility information can have a significant negative impact on the independence and mobility of citizens. For example, in our own formative interviews with wheelchair users, we uncovered a prevailing view about navigating to new areas of a city: "*I usually don't go where I don't know [about accessible routes]*" (Participant C, congenital polyneuropathy) (Hara et al., 2016).

Although maintenance issues such as buckled or cracked sidewalks can pose significant accessibility challenges, so too do larger, more permanent infrastructural issues such as poorly placed utility poles or fire hydrants or the lack of curb ramps at intersections. These issues are significant. In a precedent-setting U.S. court case in 1993, the court ruled that the "*lack of curb* *cuts is a primary obstacle to the* smooth integration of those with disabilities into the commerce of daily life" and "without curb cuts,people with ambulatory disabilities simply cannot navigate the city" (3rd Circuit, 1993).

Traditionally, sidewalk assessment has been conducted via in-person street audits (SDOT Blog, 2017; Stollof & Barlow, 2008; Streets Wiki, n.d.), which are labor intensive and costly, or via citizen call-in reports, which are done on a reactive basis. Although some cities offer sidewalk availability information online or statuses about sidewalk repair information (*e.g.*, through government 311 databases (The District of Columbia, n.d.)), these solutions are neither comprehensive nor easily scalable. They are also top-downrather than bottom-up and are not focused specifically on collecting and providing accessibility information. Recent online tools like *Axsmap.com*, *Wheelmap.org*, and *AccessTogether.org* aim to address some of these problems by having volunteers with smartphones collect location-based accessibility information. While these efforts are important, they focus on assessing the accessibility of a place (point-of-interest or POI) rather than the accessible pathways to that place, and their value propositions are intrinsically tied to the amount and quality of data they collect. In a recent review of accessibility-oriented mapping sites, for example, Ding *et al.* (2014) found that most suffered from serious data sparseness issues. Only 1.6 percent of Wheelmap POIs included accessibility data. One key limiting factor is the reliance on local populations with physical experience of a place for data collection. While local users are likely to be reliable, the dependence on *in situ* reporting dramatically limits scalability—both *who* can supply data and *how much* data they can easily supply.

In contrast, we are developing new data collection approaches that use a combination of remote crowdsourcing, machine learning, and online map imagery. Our newest effort, called

Project Sidewalk[1], enables online crowdworkers to remotely label pedestrian-related accessibility problems by *virtually* walking through city streets in Google Street View (figure 1.1). Rather than pulling solely from local populations, our potential pool of users scales to anyone with an Internet connection. To train, engage, and sustain users, we apply basic game design principles such as interactive onboarding, mission-based tasks, and progress dashboards. In 2019, we completed an 18-month deployment in Washington, D.C.: 1,153 users provided over 255,000 geo-located sidewalk accessibility labels and audited 3,000 miles of DC streets (figure 1.2). With simple quality control mechanisms (*e.g.*, majority vote), we found that minimally trained remote crowdworkers could find and label 92 percent of accessibility problems in street view scenes, including *missing curb ramps*, *obstacles in the path*, *surface problems*, and *missing sidewalks*. To our knowledge, this was the largest and most granular open dataset of sidewalk accessibility data ever collected. As in past work (Badland, et al., 2010; Clarke, et al., 2010; Rundle, et al., 2011; Wilson et al.,2012), we also found high agreement between virtual audit data of pedestrian infrastructure and traditional, in-person audit data (Hara et al., 2015; Hara, et al., 2014), which furtherestablished our approach's viability.

[1] http://sidewalk.cs.washington.edu

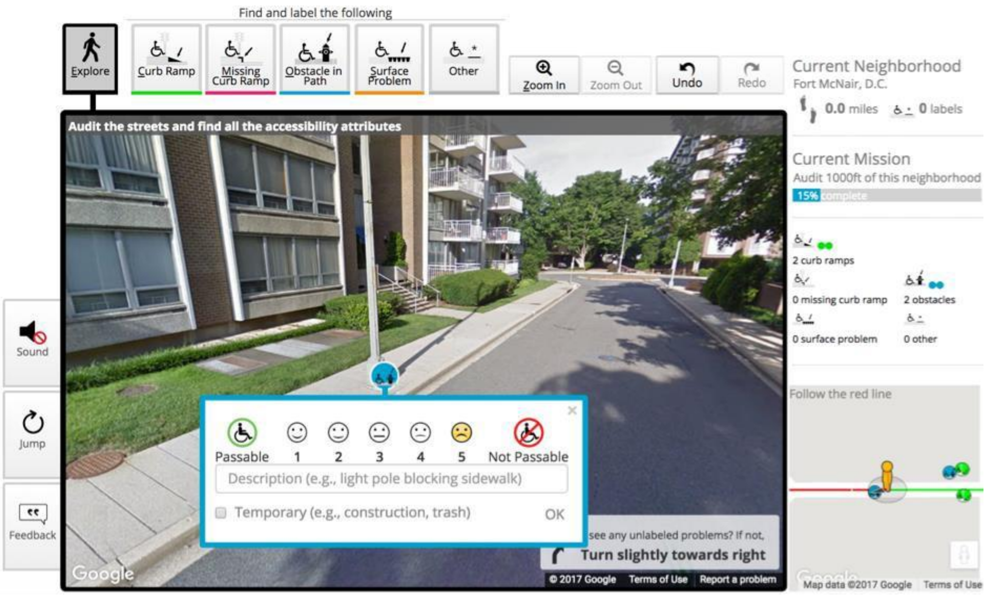

Figure 1.1. In Project Sidewalk, users are given missions to find and label accessibility problems in city neighborhoods by virtually walking in Google Street View (GSV). These labels are used to train computer vision algorithms for semi-automatic assessment and to provide new insights into and better transparency around urban accessibility.

Figure 1.2. In an initial 18-month deployment study of Project Sidewalk in Washington DC, we collected over 255,000 sidewalk accessibility labels, including curb ramps, missing curb ramps, sidewalk obstacles, and surface problems. Each dot above represents a geo-located accessibility label rendered at 50 percent translucency.

1.2 Proposed Work

Building on the above, our PacTrans proposal described three threads of additional work:

- **Expand Project Sidewalk into more cities.** First, we proposed deploying Project Sidewalk into three more cities, including two in the Pacific Northwest: Seattle, Washington, and Newberg, Oregon. These additional deployments would enable us to study new crowdsourcing workflows and user interfaces, study new methods for inferring data quality, and address questions about urban accessibility not previously possible. For example, what factors correlate with urban accessibility (*e.g.,* real estate pricing, socio-economics, zoning)? How does urban accessibility compare across cities, and what are the key contributing factors?

- **Explore deep learning methods to scale techniques.** Second, to further scale our approach, we proposed new methods to automatically identify and classify sidewalk problems using deep learning techniques, which would be uniquely enabled by our large dataset. Initial experiments have been promising: by utilizing transfer learning along with a deep residual neural network, we classified 24,000 pre-cropped images with 90.1 percent accuracy. Key questions here included how does the choice of deep learning model and input features impact performance? How can we use these machine learning techniques to validate user-supplied labels or even findaccessibility problems in street view panoramas automatically?
- **Explore and study how to visualize Project Sidewalk data.** Finally, we proposed new sidewalk accessibility models and interactive visualization tools to give stakeholders—from citizens totransit authorities—new understanding of their city's accessibility.

Below, we describe our progress across each thread. For organization, we allocate one chapter per thread.

CHAPTER 2. EXPANDING PROJECT SIDEWALK'S DEPLOYMENT

2.1 Project Sidewalk

Building on our successful D.C. pilot project and working with local and international partners, we have deployed Project Sidewalk in seven additional cities, including large urban centers such as Seattle, Wash., and Pittsburgh, Penn., as well as more rural areas like Newberg, Ore.. The full list of active cities comprises Seattle, Wash.,Newberg, Ore., Columbus, Ohio, Pittsburgh, Penn., Mexico City, Mexico, and San Pedro Garza García, Mexico. In total, 6,600 users have labeled nearly 540,000 sidewalk accessibility problems across 10,000 km of city streets and provided over 171,000 label validations (see figure 2.1, table 2.1, table 2.2)—to our knowledge, the largest open sidewalk accessibility dataset ever collected.

Our techniques and the collected data are making real-world impact. For example, our recently completed deployment in Newberg, Ore., resulted in 17,386 sidewalk accessibility labels from over 300 users (link), which were used to successfully advocate for and establish two new sidewalk repair programs by the Newberg City Council and the immediate authorization of $50,000 for repairs on city property. As all Project Sidewalk data are open, others have created their own interactive sidewalk tools such as one for Washington, D.C. by Barbera Moreno (https://bit.ly/SidewalksDC).

Figure 2.1. Example sidewalk accessibility features/problems identified and labeled by Project Sidewalk users. Label types include obstacles (blue circular label), surface problems (orange), and curb ramps (green). For readability, only one tag is shown per image (bottom-left corner), although often multiple labels were applied. All images were collected and categorized by Project Sidewalk users from the San Pedro, Mexico, deployment except for "fire hydrant," which was from Seattle, Washington.

Table 2.1. The five primary label types and corresponding tags used in Project Sidewalk. Multiple tags can be applied per label

Curb Ramps	Missing Curb Ramps	Sidewalk Obstacles	Surface Problems	Missing Sidewalks
Narrow, points into traffic, missing friction strip, steep, not enough landing space, not level with street	Alternate wheelchair route present, no alternate wheelchair route, unclear if curb ramp needed	Pole, tree, vegetation, trash/recycling can, parked car, parked bike, sign, garage entrance, stairs, street vendor, height difference, narrow, fire hydrant	Bumpy, cracks, grass, narrow sidewalk, uneven/slanted, very broken, height difference, brick	Ends abruptly, street has a sidewalk, street has no sidewalks, gravel/dirt road, shared pedestrian and car space

Table 2.2. Project Sidewalk's current dataset. To calculate label accuracy, we divide the number of *agree* validation judgments by the sum of agree and disagree judgments. The Washington DC pilot did not include crowdsourced validations and therefore, has "N/A" for those cells. Users are primarily volunteers, but some are also paid workers from Amazon Mechanical Turk.

	Users	Audit Distance	Curb Ramps	Missing Curb Ramps	Sidewalk Obstacles	Surface Problems	Missing Sidewalk	Total Labels	Total Labels Validated	Label Accuracy
Washington DC	1,395	5,482 km	150,680	19,792	22,264	8,964	45,395	247,095	N/A	N/A
Seattle, WA	3,052	2,300	51,140	26,970	7,524	17,642	26,535	129,811	68,053	82.5%
Newberg, OR	242	229	4,283	1,939	916	1,992	6,827	15,957	7,332	88.9%
Columbus, OH	380	318	11,552	875	2,768	4,401	5,778	25,374	9,299	86.7%
Pittsburgh, PA	68	169	5,315	586	2,983	2,310	2,073	13,267	1,865	87.2%
San Pedro, MX	1,105	1,285	4,006	20,090	48,080	20,949	7,167	100,292	24,197	84.1%
Azcapotzalco, MX	358	93	1170	1,141	1,590	2,119	293	6,313	3,784	89.1%
Totals	**6,600**	**9,875 km**	**226,976**	**70,252**	**84,535**	**56,258**	**93,775**	**531,796**	**114,530**	83.9%

Perhaps most exciting, because Project Sidewalk provides a low-cost sidewalk auditing approach, we have been contacted by cities that otherwise lack resources to perform their own assessments. For out Mexican-based deployments, we have been working with Liga Peatonal—a Mexico-based NGO dedicated to safe and accessible pedestrian infrastructure. Working with Liga Peatonal and local Mexican governments, we have developed and released a Spanish version of Project Sidewalk in two pilot cities: Azcapotzalco in Mexico City and San Pedro Garza García. Thus far, over 1,400 users have mapped 1,400 km of sidewalks and provided 106,000+ accessibility labels. The San Pedro government is using our data to understand inaccessible infrastructure, to examine correlates with pedestrian injuries and fatalities (in Mexico, over 44 percent of traffic-related deaths involve pedestrians), and to develop the municipality's new urban master plan with a focus on improved accessibility.

Project Sidewalk's labeling ontology was derived from accessible sidewalk standards (U.S. Access Board, 2013; US Department of Justice, 2010) and includes five primary label types and 35 tags. The label types are *curb ramps*, *missing curb ramps*, *sidewalk obstacles*, *surface problems*, and *missing sidewalks*. Each label can contain a severity assessment (a 1 to 5 scale in which 5 is an impassable barrier for a wheelchair user), an optional open-text description, and one or more label-specific tags. For example, *surface problems* can be tagged with eight additional descriptors, including *bumpy*, *cracks*, and *narrow* (table 2.1). All labels include additional metadata such as the image date, the labeling timestamp, validation information, and geo-location (lat/long).

2.2. Comparing Sidewalk Accessibility across Cities

Project Sidewalk data afford both large-scale *quantitative* analyses, such as examining what areas of a city have the most frequent and severe sidewalk accessibility problems, and complementary qualitative analyses, such as what these sidewalk problems look like (figure 2.1). We just applied for follow-up PacTrans funding to study these questions. However, as preliminary work, we performed an initial analysis of the *frequency* of sidewalk accessibility labels and their *severity* across five Project Sidewalk cities (table 2.3). As data are still being collected, this analysis is preliminary. We found that sidewalk problems were both more frequent and rated as more severe in our two Mexican cities than in the three U.S. cities (Seattle, Columbus, and Newberg). For example, *sidewalk surface problems* occurred 2.5 times per 100 meters in Mexico City and San Pedro in comparison to 0.6, 1.5, and 0.9 in Seattle, Columbus, and Newberg, respectively. These problems were not only more frequent, but they were also rated as more severe (3.6 in Mexico *vs.* 2.9, 2.1, and 2.7). By examining the labeled imagery, in San Pedro we found significant problems that would fundamentally impede wheelchair navigation, such as narrowness, physical obstacles such as vendor stands or storage, and large height differences between the street and sidewalk seam. In comparison, the poorest rated curb ramps were still problematic but less serious, such as lacking flares, landing space, or tactile strips, and/or were positioned into street traffic.

Table 2.3. An example analysis of the frequency of accessibility features and their quality across five Project Sidewalk cities. For label frequency (left), a lower number is better (except for *curb ramps*, for which higher is better). Here, San Pedro performed worst, with the fewest number of *curb ramps* and the most *missing curb ramps*, *sidewalk obstacles*, and *surface problems*. For label severity (right), quality is rated from 1 to 5, in which 5 is worst. The two Mexican cities performed worse across all primary label types.

	Label Frequency (per 100m)					Average Label Severity (Std dev)				
	Curb Ramp	Missing Curb Ramp	Missing Sidewalk	Obstacle	Surface Problem	Curb Ramp	Missing Curb Ramp	Missing Sidewalk	Obstacle	Surface Problem
Seattle, WA	2.1	1.6	1.5	0.3	0.6	1.5 (0.7)	3.8 (1.0)	4 (0.8)	3.2 (1.1)	2.9 (0.9)
Columbus, OH	4.2	0.3	1.8	1.2	1.5	1.4 (0.7)	3.8 (1.2)	4.1 (1.1)	2.2 (1.4)	2.1 (1.0)
Newberg, OR	1.9	0.9	3	0.4	0.9	1.5 (0.7)	3.9 (1.0)	3.9 (0.9)	3.1 (1.1)	2.7 (1.0)
Mexico City., MX	1.1	1.2	0.3	1.8	2.4	2.8 (1.4)	4.7 (0.6)	4.6 (0.8)	4.1 (1.0)	3.6 (1.2)
San Pedro., MX	0.6	2.2	0.4	4.1	2.5	2.8 (1.4)	4.4 (0.9)	4.5 (0.9)	4 (0.9)	3.6 (1.1)

CHAPTER 3. DEEP LEARNING FOR SIDEWALK QUALITY ASSESSMENT

Our overarching goal is to develop hybrid crowd+AI systems that quickly and accurately map and assesssidewalks from online imagery. Initially, we are focusing on two artificial intelligence (AI)-based tasks: first, **to automatically find and** classify sidewalk accessibility problems in street panoramas and satellite imagery and, second, **to auto-validate** crowdsourced human labels. As no machine language (ML) model is perfect, part of our proposed effort will be to explore methods that effectively combine humans+AI (e.g., through AI-assisted sidewalk labeling or feedback about the inferred accuracy of a label). Our deep learning experiments and results are described in detail by Weld et al. (2019). We summarize them below.

3.1. ResNet Machine Learning Model

Our current ML approach uses a modified residual neural network (ResNet) implemented in PyTorch. Typically, ResNet-based neural networks take images (pixels) as input. In our case, however, we want to leverage contextual knowledge about the sidewalk and the sidewalk label's position and geography.

These non-image features include the angle between the label point and the street axis, the distance to the nearest intersection, and the distance and bearing of the label to the city's urban center. We also plan to explore additional features such as census-tract information, real-estate pricing, and neighborhood zoning—each of which may correlate with sidewalk presence and condition and thus, improve our model.

To incorporate both sets of features (image and non-image), we currently perform a two-stage process. First, we feed a 224x224x3 image vector through a series of convolutional layers derived from ResNet to obtain a 512x1 vector—essentially treating intermediate ResNet layers as a method for dimensionality reduction. We then append an additional vector of non-image features (currently 12x1) to generate a final input vector of size 524x1. This combined vector is fed into the last network layer, which outputs a prediction vector of 5x1 in which output rows correspond to prediction confidence for each of the five current sidewalk condition classes: curb ramps, missing curb ramps, obstructions, surface problems, anda null class. To determine the final prediction, we compute the argmax of the prediction vector.

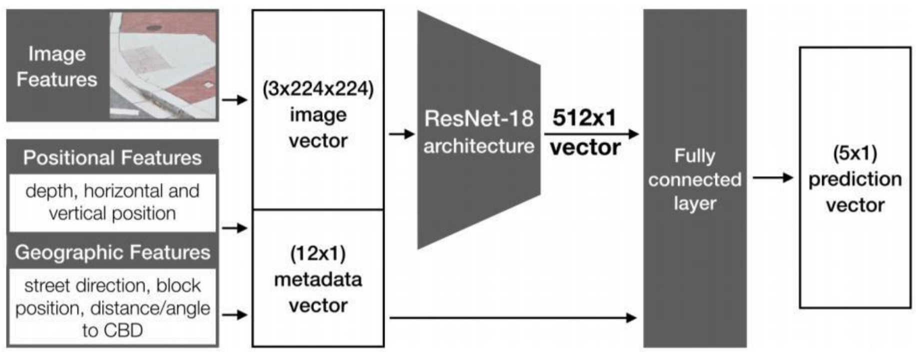

Figure 3.1. The structure of our modified ResNet-18 architecture, which, in addition to pixel-based imagery, incorporates extra features such as position in scene, depth information, and geographic data.

To further improve performance, we use transfer learning to initialize the ResNet model with weights learned from pre-training on the ImageNet image corpus. In preliminary experiments, this significantly improved our results—for example, the accuracy of auto-validating human-labeled curb ramps increased by over 60 percent.

3.2. Evaluation and Results

To evaluate our ResNet approach, we split our data into training, validation, and test datasets. For training and validation, we used 205,385 image-based sidewalk accessibility labels from 58,035 streetscape panoramas from the DC dataset. We randomly partitioned these panoramas and their corresponding labels into training, validation, and test sets following an

80/10/10 split. For testing, we manually created a ground-truth dataset by labeling a subset of the test dataset (224 of the 5,774 panoramas).

We examined both auto-validating and auto-labeling performance. While direct comparisons to previous work are difficult because of a lack of open datasets and software, our approach appears to substantially improve upon previous ML-based approaches (including our own previous work). For automatically validating pre-labeled crowdsourced scenes, the average precision was 81. Percent and recall was 77.2 percent. For auto-labeling streetscape panoramas, our performance dropped, as expected, to 47 percent precision and 41.2 percent recall. Unlike auto-validation in which our model is supplied a pre-cropped image around a human-supplied label, the auto-labeling task is more challenging: using a sliding window, our model searches through an entire panorama to find and classify a problem.

3.2.1. Performance as a Function of Training Set Size

To explore the effect of dataset size on performance, we trained our auto-validation model on increasingly large, randomly sampled subsets of our training dataset. Our results are shown in figure 3.2.Unsurprisingly, performance was positively correlated with training set size. With only 1,000 crops, our overall precision was 61 percent and recall was 63.9 precent, which improved to 79.7 percent and 80.4 percent with the full training set (213,000 crops). Interestingly, even at the maximum training set size, a plateau was not reached, particularly for the worst performing classes (surface problems and missing curb ramps). This suggests that even more training data would be beneficial. With our larger datasets, we would like to rerun this experiment.

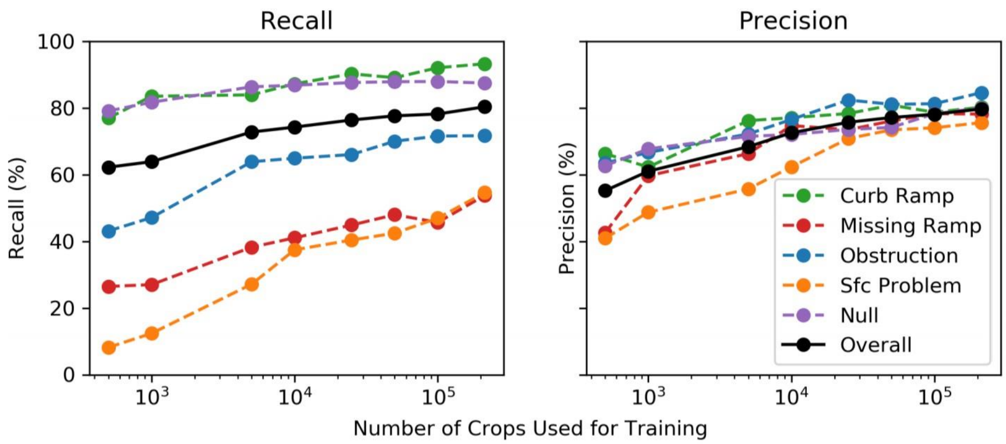

Figure 3.2. Performance overall and by feature type as the size of the training set increases. Note the log scale on the x axis

3.2.2. Exploring Cross-City Generalizability

While the results presented above were focused on Washington, D.C., ideally a model trained on one city's streetscape images would generalize to other cities. However, sidewalk infrastructure can vary in quality and design across geographic areas and neighborhood types (e.g., suburban residential vs. downtown commercial), impacting visual appearance. To examine the cross-city generalizability of our models, we used recent open datasets from two new Project Sidewalk deployment cities: Seattle, Wash. (a major U.S. city on the West Coast with 750,000 residents) and Newberg, Ore. (a small town outside of Portland, Ore., with 22,000 residents). For our experiments described below, we used a subset: the 9,535 labels created or validated by members of the Project Sidewalk team (to ensure maximum quality).

For each city, an additional 1,500 null labels were randomly sampled from the GSV panoramas. For each label, we used center cropping to transform points to image crops. Crops were randomly partitioned into train and test sets with an 80/20 split.

In total, we conducted four cross-city experiments; all used the auto-validation model but varied in training set composition. Each model was evaluated on the test set for its respective city (either Seattle or Newberg). The four models included the following:

1. **Baseline Model:** trained on D.C. only.
2. **D.C. + New City Model:** trained on both D.C. and the other city (either Seattle or Newberg).
3. **New City-Only Model:** trained on only Seattle or only Newberg without any D.C. data.
4. **(Best performing) New City-Only Model Initialized with D.C.:** same as New City-Only, but initialized with the weights from the D.C. model ("pre-trained") before training on the new city data.

The results are presented in figure 3.3**.** Both cities demonstrated similar trends in performance across the four models. The Baseline Model—which was trained on D.C.-only but tested on the new cities—resulted in 55.6 percent precision and 56.3 percent recall for Seattle and 46.0 percent precision and 48.6 percent recall for Newberg. Interestingly, some label types performed quite well, suggesting some uniformity across cities. For example, curb ramps achieved a recall of 83.8 percent in Seattle and 90.5 percent in Newberg (although precisions were at 40 percent for both cities). This high cross-city performance occurred perhaps because curb ramps were the only designed urban feature in our dataset, which likely resulted in visual and contextual consistency; the other label types were due to dilapidation (e.g., surface problems) or the lack of urban design (e.g., missing curb ramps). For the other three models, even training on a small amount of new city-specific data resulted in significant improvements over the baseline, with Seattle+D.C. improving from 55.6 percent to 71.9 percent for precision and from 56.4 percent to 78.4 percent for recall, and Newberg+D.C. improving from 46.0 percent to 60.3 percent for precision and from 48.6 percent to 75.9 percent for recall.

	Seattle		Newberg	
	Train	**Test**	**Train**	**Test**
Number of Panos	3,101	775	3,199	801
Curb Ramp	1,648	412	1,814	453
Missing Ramp	1,194	298	588	147
Obstruction	633	158	307	76
Surface Problem	1,000	250	446	111
Null	1,202	298	1,199	301
All Labels	**5,677**	**1,416**	**4,354**	**1,088**

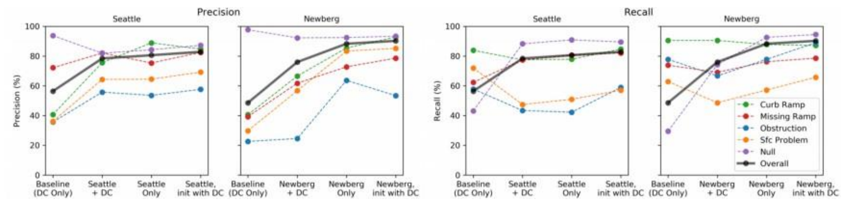

Figure 3.3. Cross-city generalizability datasets and results from four different cross-city training+testing experiments.

Overall, the best performing method was to use D.C. data to teach neural network weights that were then used to initialize ("pre-train") a model that was further trained on each individual city's data (either Seattle or Newberg). Our results were promising: the Seattle model pre-trained on D.C. achieved 76.2 percent precision and 82.8 percent recall, while the Newberg version achieved 80.6 percent precision and 90.2 percent recall. These results were competitive with the auto-validation model trained and tested on D.C. data, with overall 81.3 percent precision and 77.2 percent recall. The results suggested that our models can perform well with only a small amount of new training data per city.

3.3. Discussion

We discuss biases and potential dangers in automating sidewalk assessments, describe dataset limitations and how they may impact our results, enumerate future work relevant to machine learning,and reflect on possible uses for and impacts of automatically assessing sidewalks.

3.3.1. Biases in Automating Sidewalk Assessment

Any AI or ML system contains intrinsic norms, values, and biases. Ours is no different. These exist at many levels, from the data collection to the design and implementation of the machine learning approach itself. In this project, we leverage the Project Sidewalk dataset, which offers highly granular, geolocated sidewalk accessibility labels; however, the four label types (curb ramps, missing curb ramps, obstructions, and surface problems) do not comprehensively capture sidewalk accessibility. For example, no labels exist for crosswalks, accessible pedestrian signals (e.g., audio-based stoplights), stairs, or accessible public transit stops. Our ML approach may also introduce new biases. For instance, our system, though successful in a small number of cities, may not work in locations where sidewalk problems are less obvious and detectable or where streetscape imagery is not widely available. Failures in our ML model may incorrectly inform policymakers that certain sidewalks are accessible, hampering appropriate sidewalk transportation and funding. We emphasize that “on-the-ground” investigations should accompany the use of these ML tools to provide a mechanism for continued community involvement.

3.3.2. Dataset Limitations

Neural networks are notably data-hungry, as their expansive parameter space requires a large amount of training data. Our project is enabled by the size and richness of the crowdsourced Project Sidewalk dataset. As expected, however, crowdsourced data are noisy, which can undermine ML model training. The task of sidewalk assessment for accessibility problems is inherently subjective, and consistent labeling is difficult even for humans. For example, labelers may assign labels differently for the same accessibility problem. Similarly, some sidewalk areas are occluded from view within panorama images, making assessment impossible. While consistent, detailed labeling rules and training help, labelers still face ambiguity—such as, "Is this a surface problem or an obstruction?" "Are pedestrians intended to cross at this intersection?" Finally, in Project Sidewalk, crowdworkers are not asked to comprehensively label each visited GSV panorama; instead, their focus is on finding accessibility problems. So, once a problem is labeled, it need not be labeled in adjacent panoramas. This “under-labeling” per panorama may contribute to artificially high false negative rates in our model evaluations.

3.3.3. Future Work: Data and Computer Vision Methods

In this project, we trained models separately for the autovalidation and auto-labeling tasks, as this increased performance. For improved trainability and usability, future work should explore the development of a universal model that would be optimal for both tasks. Collecting additional data will also improve the performance of our models, as shown by our investigations of performance as a function of training set size. We will need more large, labeled sidewalk accessibility datasets from cities around the world, for which we will continue to reach out to potential communities and partners.

Mechanisms for improving human-labeled data quality remain an open question for Project Sidewalk. Also, with more detailed training data, it may become possible to train our system to produce more accurate and nuanced labels, such as the severity of a surface problem or the presence of friction strips on curb ramps. The additional data needed for these tasks are already being collected by Project Sidewalk,with crowdworkers rating the severity of problems on a 1 to 5 scale. Including these ratings in our model is a logical first step toward increased label detail. The performance of our neural network may transfer from one city to another to some degree, as shown by our Newberg and Seattle experiments, but will likely be lower because of differences between cities. The ability to transfer trained networks between cities could ease data collection requirements for new areas, but at the cost of potentially excluding communities with drastically different sidewalk infrastructure.

3.3.4. Impact on Sidewalk Accessibility

Our overarching long-term vision is to be able to automatically assess the accessibility of cities at scale within hours. This would put a powerful accountability tool in the hands of accessibility advocates, for example when monitoring the adherence of cities to ADA mandates. However, automation of any taskbrings up a number of complex challenges and ethical considerations that merit further discussion.

While our machine learning approach demonstrates a significant improvement over the state-of-the-art on image-based, automated assessment of sidewalk-level accessibility problems, it is important to assess how to use these results to meaningfully affect sidewalk accessibility. In addition to identifying concrete problems to be solved, the results could be aggregated into quantitative accessibility metrics on the street, neighborhood, or city level, as in the work that proposed AccessScore (Li, et al., 2018). This would allow users to make more informed choices about neighborhoods, housing, and transportation. We also see potential future applications to the problem of routing users with differing mobility, although this task is notably more difficult and requires more precision than our algorithm can yet provide. For example, even one missed obstruction could render a route impassible.

Our approach leverages only a portion of the data available through Project Sidewalk, which now include additional tags on labels such as the presence or absence of friction strips on curb ramps,differentiation between types of obstacles such as street poles or cars, and severity metrics.

Incorporating all of these details into automatic classifiers in a manner similar to this work could also result in more usable and accurate metrics for routing and other applications.

3.4 Conclusion

The overarching goal of our work is to develop fast and accurate sidewalk assessment methods using machine learning to help transform how city governments and citizens alike track, maintain, and use pedestrian infrastructure. In this project, we demonstrated a promising deep learning approach for auto-validating and auto-labeling sidewalks in streetscape imagery. Our ResNet models significantly improved upon the performance of previous automated systems and, in some cases, met or exceeded human labeling performance.

CHAPTER 4. EXPLORING SIDEWALK DATA AND VISUALIZATION REQUIREMENTS

Finally, in the last thread of our work, we explored data needs and visualization approaches for sidewalk accessibility data. We performed a multi-stakeholder analysis of the priorities, perspectives, and local decision-making around urban accessibility—specifically, pedestrian infrastructure—in three U.S. cities: Seattle, Washington, D.C., and New York. We conducted semi-structured interviews with 25 participants drawn from five stakeholder groups: (1) policymakers who develop city-wide accessibility policies and regulations, (2) department officials who implement and maintain these regulations (e.g., departments of transportation, an office on aging), (3) accessibility advocates who work to change ineffective policies, (4) people with mobility impairments (MI individuals) who have some form of mobility disability and directly experience (in)accessible environments; and (5) caregivers who are friends, family members, or professionals that care for MI individuals.

The semi-structured interview had two-parts: a formative component, which asked about perspectives of, approaches for, and decision-making processes about urban accessibility, and a design probe component, which examined reactions to envisioned urban accessibility analysis and visualization tools (see figure 4.1).

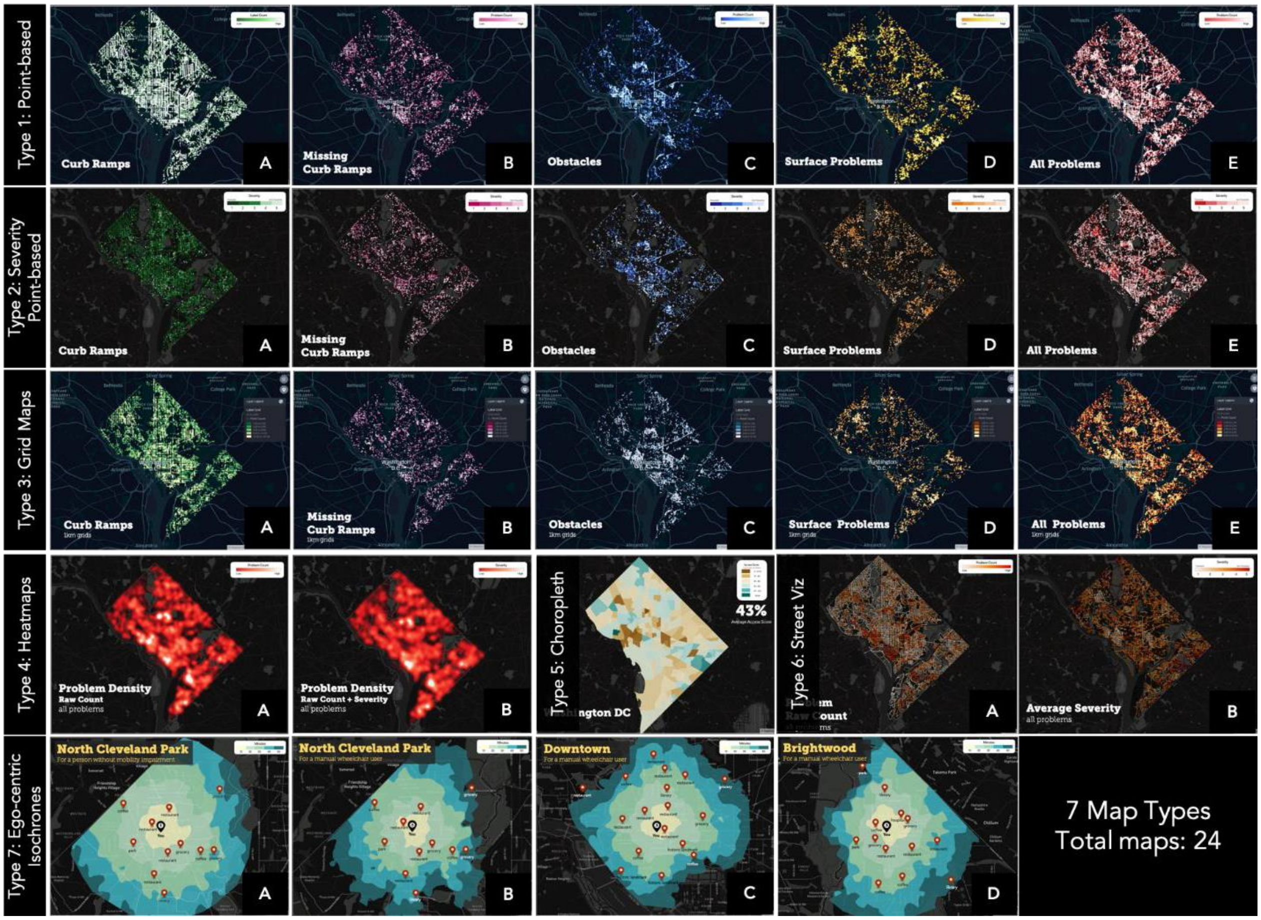

Figure 4.1. Example mockup sidewalk accessibility visualizations developed for our formative interviews.

We sought to address the following research questions:

- **RQ1**: Across stakeholder groups, what are the information needs and challenges for assessing and making decisions about urban accessibility and the role of data and technology in these practices?
- **RQ2**: How do stakeholder groups communicate and interact together to assess priorities and make decisions?
- **RQ3**: What are the future design opportunities to improve existing assessment, data analytics, and visualization methods?

This work is ongoing. We published an initial paper entitled *Urban Accessibility as a Socio-PoliticalProblem: A Multi-Stakeholder Analysis* (Saha et al., 2020), which described results from the first part of the interview (summarized below). We are now analyzing results from the second part of the interview (the visualization component). These latter results have not yet been published.

4.1. Summary of Findings

Using iterative qualitative coding, we identified and presented three high-level themes related to data and technology practices, decision-making, and challenges that impede accessible infrastructure development. Our findings highlighted the technological barriers in assessing urban accessibility as well as the socio- political barriers to infrastructure development. For the former, we identified disparities among groups in data and tool access. For example, policymakers had the least data/tool access, while advocates had insufficient tools to fit their needs. For the latter, we found that the presence of numerous actors, organizations, and their conflicting interests complicated decision-making and made accountability toward accessibility improvements difficult. Combined with limited funding and public disinterest, political will to bring change was also affected.

Our work contributes to the growing computer-supported cooperative work / human-computer interaction (CSCW/HCI) literature on urban governance and civic systems using multi-stakeholder analysis as a method (Asad, et al., 2017; Corbett & Le Dantec, 2018; Karusala, et al., 2019). Using this approach, we extend previous work (Larkin et al., 2015) by presenting the first U.S.-focused study that brought multiple perspectives to the task of understanding accessible infrastructure development processes. Our contributions include better understanding of the following:

1. The current practices and challenges of working within the socio-political realities of a civic ecosystem,
2. The role of technology in supporting and potentially undermining existing practices,
3. A *civic interaction space* that lays out the roles of and interactions between stakeholders in the civic decision-making structure, and
4. Points of future technological interventions in the form of data-driven assessment and civic engagement tools for improving accessibility through planning, advocacy, and policymaking.

4.2. Discussion

Within the CSCW and HCI literature (Aragón, et al., 2020; Boehner & Disalvo, 2016), civic technology has been positioned as a platform for open, collaborative government and community action, facilitating civic conversations and collaborative decision-making practices (Crivellaro et al., 2019; Foth, et al., 2013; Vlachokyriakos et al., 2016). Establishing trust between stakeholders is at the core of a successful civic engagement model (Grimmelikhuijsen, 2012; Harding, et al., 2015). In our findings, *lack of accountability* and *civic participation* were two significant issues for accessible infrastructure development. Trust strongly interconnects both issues: increasing accountability leads to increased trust in the government, which further reinforces and encourages civic participation. Below, we elaborate on these two issues within urban accessibility and present open questions for the CSCW community working on civic technology.

Lack of accountability. Because of the decentralized nature of accessibility improvements, the seamless blend between private and public spaces in urban centers, and the contributions of current governmental policies, individuals, and agencies, a clear understanding of *who* is responsible for accessible infrastructure is lacking. In U.S. cities, departments of transportation manage street and sidewalk infrastructure on public land; however, commercial building entrances and indoor spaces are the purview of private businesses, and sidewalks adjacent to residences are the property owner’s responsibility. These interdependencies, although core to urban life, create conflict and obscure accountability (Voida, et al., 2014). As one participant stated when describing tensions between a privately owned transit agency and their governmental organization: “*[it’s] city vs. private vs. federal.”* How can civic technology better surface these tensions and allow private citizens and governmental agencies to track and assess accessibility progress and, ultimately, increase accountability?

Lack of civic participation. The issue of perceived public disinterest by policymakers and department officials can also impact infrastructure development. For example, without voter interest in transportation levies, policymakers have difficulty funding large transportation projects through taxes. Public disinterest is often the result of being unaware of inaccessible environments or a lack of perceived personal impact. This suggests a need for *wider awareness* among communities about accessibility and the importance of civic participation. The current engagement practices related to 311 services are largely volunteer based, and representation of citizen voices is often inadequate. A successful approach to bring wide-scale policy change has been disability activism (Campbell & Oliver, 2013; Hahn, 1985; Kitchin & Wilton, 2003) such as the 1990 ADA civil rights legislation (Civil Rights Division United States Department of Justice, 1990) and the 2019 block-the-box legislation in Washington[2]. The success of such initiatives

---

1 [2] ‘Don’t block-the-box’ legislation was passed via House Bill 1793, which permits Seattle to use camera enforcement to fine motorists who block crosswalks and bus lanes (“Don’t Block the Box Bill Passes State Legislature | The Urbanist,” n.d.).

leads us to ask, how can civic technology support new practices that strengthen the collective voice of the people to drive change?

CHAPTER 5. REFERENCES

3rd Circuit, C. of A. (1993). *Kinney v. Yerusalim, 1993 No. 93-1168.* Retrieved from http://www.access-board.gov/prowac/yerusalim.htm

Aragón, P., Flores-Saviaga, C., Garcia, A. A., Dantec, C. A. Le, & Saldivar, J. (2020). Civic Technologies: Research, Practice and Open Challenges. *CSCW 2020 Workshop*.

Asad, M., Le Dantec, C. A., Nielsen, B., & Diedrick, K. (2017). Creating a Sociotechnical API: Designing City-Scale Community Engagement. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, 2295–2306. https://doi.org/10.1145/3025453.3025963

Badland, H. M., Opit, S., Witten, K., Kearns, R. A., & Mavoa, S. (2010). Can virtual streetscape audits reliably replace physical streetscape audits? *Journal of Urban Health : Bulletin of the New York Academy of Medicine*, *87*(6), 1007–1016. https://doi.org/10.1007/s11524-010-9505-x

Boehner, K., & Disalvo, C. (2016). Data, design and civics: An exploratory study of civic tech. *Conference on Human Factors in Computing Systems - Proceedings*, 2970–2981. https://doi.org/10.1145/2858036.2858326

Campbell, J., & Oliver, M. (2013). *Disability politics: Understanding our past, changing our future*. Routledge.

Civil Rights Division United States Department of Justice. *Americans with Disabilities Act of 1990*. ,(1990).

Clarke, P., Ailshire, J., Melendez, R., Bader, M., & Morenoff, J. (2010). Using Google Earth to conduct a neighborhood audit: Reliability of a virtual audit instrument. *Health & Place*, *16*(6), 1224–1229. https://doi.org/10.1016/J.HEALTHPLACE.2010.08.007

Corbett, E., & Le Dantec, C. A. (2018). The problem of community engagement: Disentangling the practices of municipal government. *Conference on Human Factors in Computing Systems - Proceedings*, *2018-April*. https://doi.org/10.1145/3173574.3174148

Crivellaro, C., Anderson, R., Lambton-Howard, D., Nappey, T., Olivier, P., Vlachokyriakos, V., … Wright, P. (2019). Infrastructuring public service transformation: Creating collaborative spaces between communities and institutions through HCI research. *ACM Transactions on Computer-Human Interaction*, *26*(3). https://doi.org/10.1145/3310284

Ding, C., Wald, M., & Wills, G. (2014). A Survey of Open Accessibility Data. *Proceedings of the 11th Web for All Conference*, 37:1--37:4. https://doi.org/10.1145/2596695.2596708

Don't Block the Box Bill Passes State Legislature | The Urbanist. (n.d.).

Foth, M., Agudelo, L. P., & Palleis, R. (2013). *Digital soapboxes: towards an interaction design agenda forsituated civic innovation*. 725–728. https://doi.org/10.1145/2494091.2495995

Friedman, G. (2018, May). Portland to settle with wheelchair users for $113 million in sidewalk fixes. *Oregon Live*. Retrieved from https://www.oregonlive.com/portland/2018/05/portland_to_settle_with_wheelc.html

Grimmelikhuijsen, S. (2012). Linking transparency, knowledge and citizen trust in government: an experiment. *International Review of Administrative Sciences*, *78*(1), 50–73. https://doi.org/10.1177/0020852311429667

Gutman, D. (2017, July 24). Settlement: Seattle to build thousands of sidewalk curb ramps over next 18 years. *The Seattle Times*. Retrieved from https://www.seattletimes.com/seattle-news/transportation/settlement-seattle-to-build-thousands-of-sidewalk-curb-ramps-over-next-18- years/

Hahn, H. (1985). Toward a politics of disability: Definitions, disciplines, and policies. *The Social Science Journal*.

Hara, K., Azenkot, S., Campbell, M., Bennett, C. L., Le, V., Pannella, S., … Froehlich, J. E. (2015). ImprovingPublic Transit Accessibility for Blind Riders by Crowdsourcing Bus Stop Landmark Locations with Google Street View: An Extended Analysis. *ACM Transactions on Accessible Computing (TACCESS)*, *6*(2), 5:1--5:23. https://doi.org/10.1145/2717513

Hara, K., Chan, C., & Froehlich, J. E. (2016). The Design of Assistive Location-based Technologies for People with Ambulatory Disabilities: A Formative Study. *Proceedings of the 2016 CHI Conference onHuman Factors in Computing Systems*, 1757–1768. https://doi.org/10.1145/2858036.2858315

Hara, K., Le, V., & Froehlich, J. (2013). Combining crowdsourcing and google street view to identify street-level accessibility problems. *Proceedings of the SIGCHI Conference on Human Factors inComputing Systems - CHI '13*, 631–640. https://doi.org/10.1145/2470654.2470744

Hara, K., Sun, J., Jacobs, D. W., & Froehlich, J. E. (2014). Tohme: Detecting Curb Ramps in Google StreetView Using Crowdsourcing, Computer Vision, and Machine Learning. *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, TBD.

Harding, M., Knowles, B., Davies, N., & Rouncefield, M. (2015). HCI, civic engagement & trust. *Conference on Human Factors in Computing Systems, 2015-April*, 2833–2842. https://doi.org/10.1145/2702123.2702255

Hu, W. (n.d.). For the Disabled, New York's Sidewalks Are an Obstacle Course - The New York Times. Retrieved September 1, 2018, from New York Times website: https://www.nytimes.com/2017/10/08/nyregion/new-york-city-sidewalks-disabled-curb-ramps.html

Karusala, N., Wilson, J., Vayanos, P., & Rice, E. (2019). Street-Level Realities of Data Practices in Homeless Services Provision. *Proceedings of the ACM on Human-Computer Interaction*, *3*(CSCW),1–23. https://doi.org/10.1145/3359286

Kitchin, R., & Wilton, R. (2003). Disability activism and the politics of scale. *The Canadian Geographer/LeG?Ographe Canadien*, *47*(2), 97–115. https://doi.org/10.1111/1541-0064.00005

Larkin, H., Hitch, D., Watchorn, V., & Ang, S. (2015). Working with policy and regulatory factors to implement universal design in the built environment: The Australian experience. *InternationalJournal of Environmental Research and Public Health*, *12*(7), 8157–8171. https://doi.org/10.3390/ijerph120708157

Li, A., Saha, M., Gupta, A., & Froehlich, J. E. (2018). Interactively Modeling And Visualizing Neighborhood Accessibility At Scale: An Initial Study Of Washington DC. *Poster Proceedings of ASSETS'18*, 444– 446. Retrieved from https://doi.org/10.1145/3234695.3241000

McMillen, B., Axelson, P. W., Chesney, D. A., Galvan, D. V., Kirschbaum, J. B., Longmuir, P. E., … Wong, K.

M. (1999). *Designing Sidewalks and Trails for Access. Part I of II : Review of Existing Guidelines andPractices Notice*. https://doi.org/10.1016/j.steroids.2007.12.028

National Council on Disability. (2007). *The Impact of the Americans with Disabilities Act: Assessing theProgress Toward Achieving the Goals of the ADA*. Retrieved from http://goo.gl/RkLL63

Rundle, A. G., Bader, M. D. M., Richards, C. A., Neckerman, K. M., & Teitler, J. O. (2011). Using Google Street View to Audit Neighborhood Environments. *American Journal of Preventive Medicine*, *40*(1),94–100. https://doi.org/10.1016/J.AMEPRE.2010.09.034

Saha, M., Chauhan, D., Patil, S., Kangas, R., Heer, J., & Froehlich, J. E. (2020). Urban Accessibility as a Socio-Political Problem: A Multi-Stakeholder Analysis. *Proceedings of the ACM on Human ComputerInteraction (PACM HCI); To Be Presented at CSCW2020*.

Saha, M., Saugstad, M., Maddali, H., Zeng, A., Holland, R., Bower, S., … Froehlich, J. E. (2019). ProjectSidewalk: A Web-based Crowdsourcing Tool for Collecting Sidewalk Accessibility Data at Scale. *Proceedings of CHI 2019*. Glasgow, Scotland: ACM.

SDOT Blog. (2017). Seattle Sidewalk Survey Update. Retrieved February 1, 2019, from https://sdotblog.seattle.gov/2017/08/28/seattle-sidewalk-survey-update/

Stollof, E. R., & Barlow, J. M. (2008). *Pedestrian Mobility and Safety Audit Guide*. Retrieved fromhttp://www.ite.org/PedAudits/AuditGuide.pdf

Streets Wiki. (n.d.). Walk Audit. Retrieved from http://streetswiki.wikispaces.com/Walk+Audit The District of Columbia. (n.d.). 311 Online http://311.dc.gov/.

U.S. Access Board. (2013). *Public Rights-of-Way Accessibility Guidelines*. Retrieved from https://www.access-board.gov/prowag/

U.S. Census Bureau. (2012). *Americans with Disabilities: 2010 Household Economic Studies*. Retrievedfrom http://www.census.gov/prod/2012pubs/p70-131.pdf

United States Department of Justice Civil Rights Division. (1990). *Americans with Disabilities Act of 1990, Pub. L. No. 101-336, 104 Stat. 328*.

US Department of Justice. (2010). *2010 ADA Standards for Accessible Design*. Retrieved from https://www.ada.gov/regs2010/2010ADAStandards/2010ADAstandards.htm

Vlachokyriakos, V., Crivellaro, C., Le Dantec, C. A., Gordon, E., Wright, P., & Olivier, P. (2016). Digital civics: Citizen empowerment with and through technology. *Conference on Human Factors in Computing Systems - Proceedings*, 1096–1099. https://doi.org/10.1145/2851581.2886436

Voida, A., Dombrowski, L., Hayes, G. R., & Mazmanian, M. (2014). Shared values/conflicting logics: working around e-government systems. *Proceedings of the SIGCHI Conference on Human Factors inComputing Systems*, 3583–3592. https://doi.org/10.1145/2556288

Weld, G., Jang, E., Li, A., Zeng, A., Heimerl, K., & Froehlich, J. E. (2019). Deep Learning for AutomaticallyDetecting Sidewalk Accessibility Problems Using Streetscape Imagery. *The 21st International ACMSIGACCESS Conference on Computers and Accessibility*, 196–209. https://doi.org/10.1145/3308561.3353798

Wilson, J. S., Kelly, C. M., Schootman, M., Baker, E. A., Banerjee, A., Clennin, M., & Miller, D. K. (2012).Assessing the Built Environment Using Omnidirectional Imagery. *American Journal of PreventiveMedicine*, *42*(2), 193–199. https://doi.org/10.1016/J.AMEPRE.2011.09.029