ALWAYS ANOTHER MOUNTAIN

STORIES OF PERSISTENCE, PUSHBACK, & THE PURSUIT OF IMPACT

Jon E. Froehlich
Professor, Computer Science
Core Faculty, Urban Design & Planning
Associate Director, CREATE & PacTrans

University of Chicago
Oct 16, 2026

MAKEABILITY LAB
DESIGN USE BUILD
CREATE
Center for Research and Education on Accessible Technology and Experiences
PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE & ENGINEERING
UNIVERSITY of WASHINGTON

2026

CHI'26 BARCELONA



MAKEABILITY LAB

MY PHD STUDENTS



Arnavi Chheda-Kothary
UW CS



Aditya Gunturu
UW CSE



Kotaro Hara
UMD CS
Now: Prof at SMU



Liang He
UW CS
Now: Prof at UT-Dallas



Jared Hwang
UW CSE



Dhruv Jain
UW CS
Now: Prof at Michigan



Seokbin Kang
UMD CS
Now: Director, Zoom



Jaewook Lee
UW CSE
Soon: Apple



Chu Li
UW CS
On academic job market



Matthew L Mauriello
UMD CS
Now: Prof at Delaware



Manaswi Saha
UW CSE
Now: Accenture Labs



Lee Stearns
UMD CS
Now: APL



Xia Su
UW CS
Now: Google



Daniel Campos Zamora
On academic job market

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

- Human-AI
- Urban Computing
- Visualization
- Accessibility



Daniel Campos Zamora

UW CS

On academic job market

- Robotics
- Fabrication
- Sensing Techniques
- Sustain./Accessibility

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

Human-AI

Urban Computing

Visualization

Accessibility



Daniel Campos Zamora

UW CS

On academic job market

Robotics

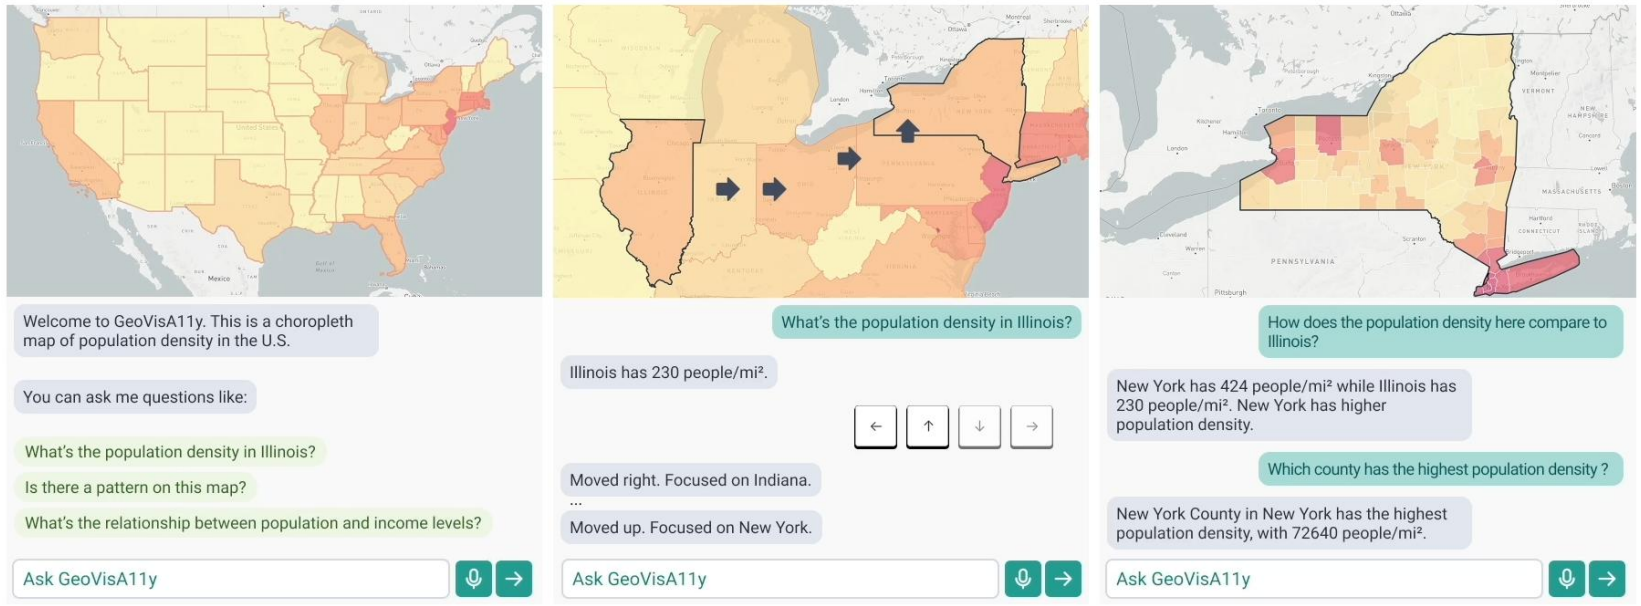
Fabrication

Sensing Techniques

Sustain./Accessibility

GeoVisA11y

An AI-based Geovisualization Question-Answering System for Screen-Reader Users



Li *et al.*, GeoVisA11y, CHI'26 Best Paper 🏆

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

- Human-AI
- Urban Computing
- Visualization
- Accessibility

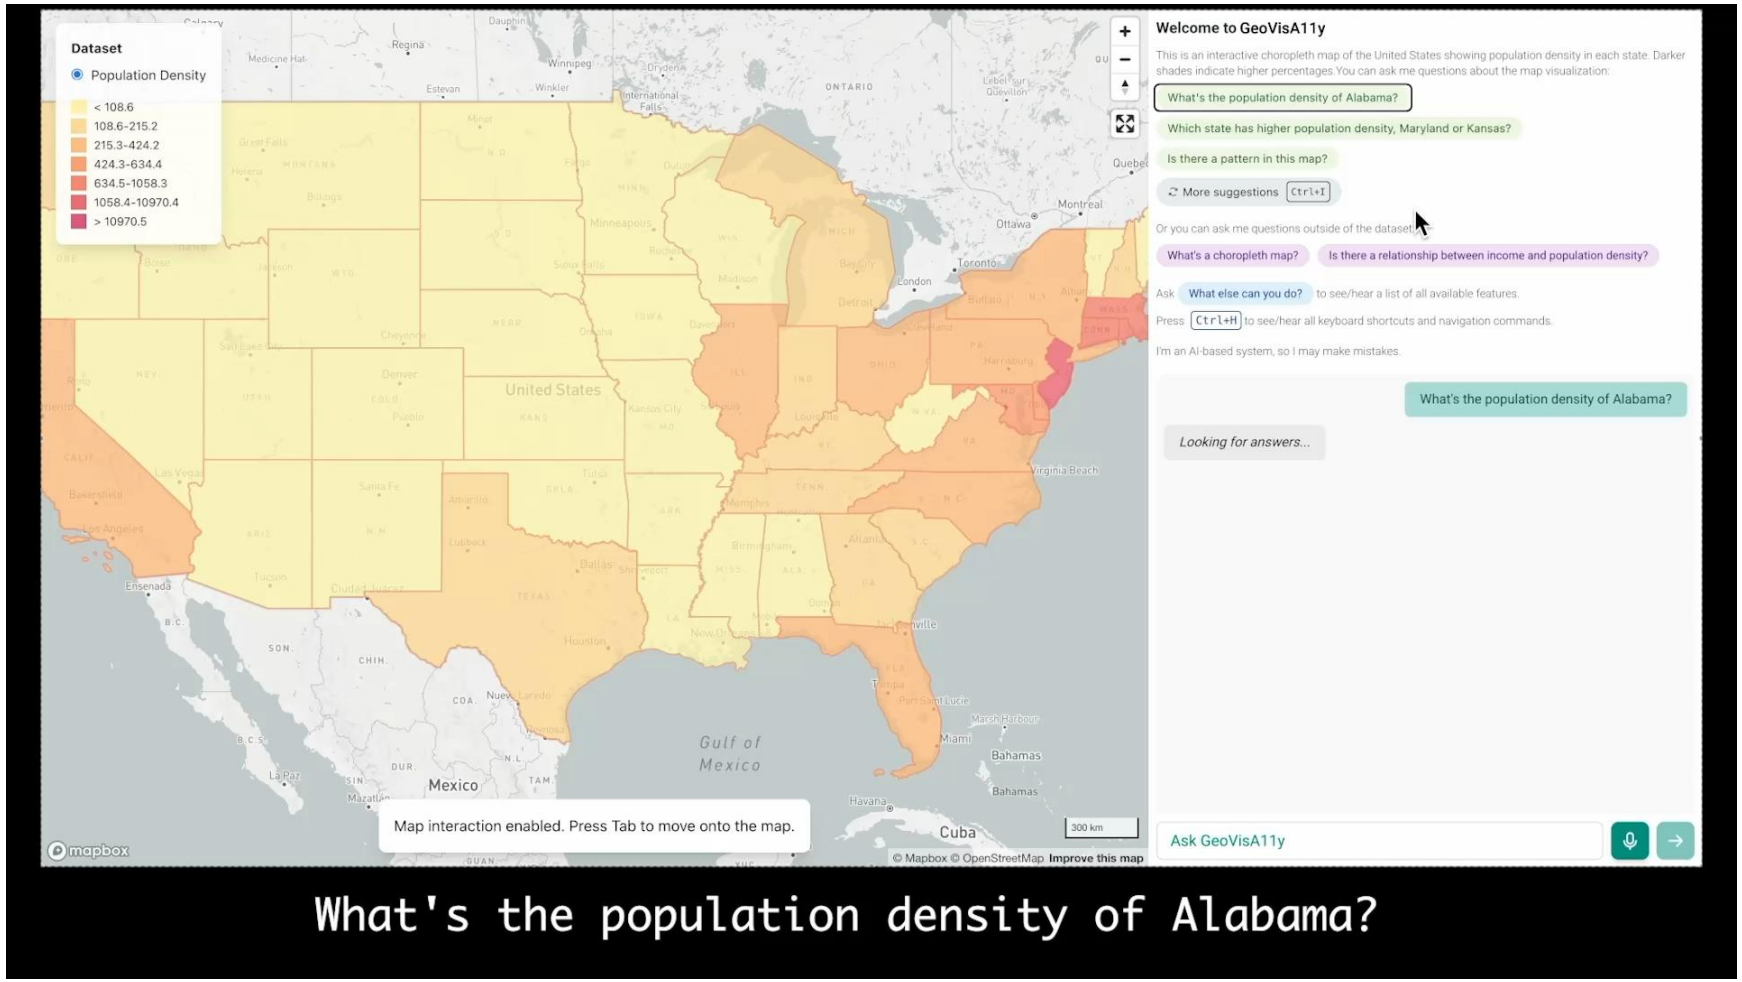


Daniel Campos Zamora

UW CS

On academic job market

- Robotics
- Fabrication
- Sensing Techniques
- Sustain./Accessibility



Li *et al.*, GeoVisA11y, CHI'26 Best Paper 🏆

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

Human-AI

Urban Computing

Visualization

Accessibility



Daniel Campos Zamora

UW CS

On academic job market

Robotics

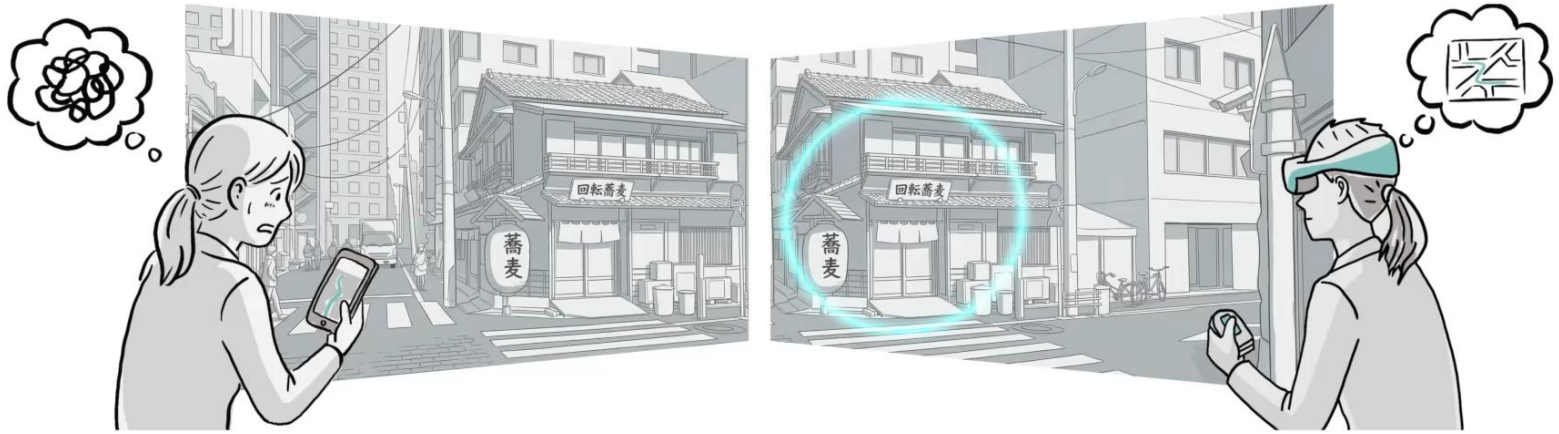
Fabrication

Sensing Techniques

Sustain./Accessibility

LandmarkLens

Predicting and Presenting Effective Landmarks for Mixed-Reality Urban Exploration



Chu Li, Yotam Sechayk, Jared Hwang, Jon E. Froehlich, Takeo Igarashi

University of Washington & University of Tokyo

Li *et al.,* LandmarkLens, UIST'26

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

- Human-AI
- Urban Computing
- Visualization
- Accessibility

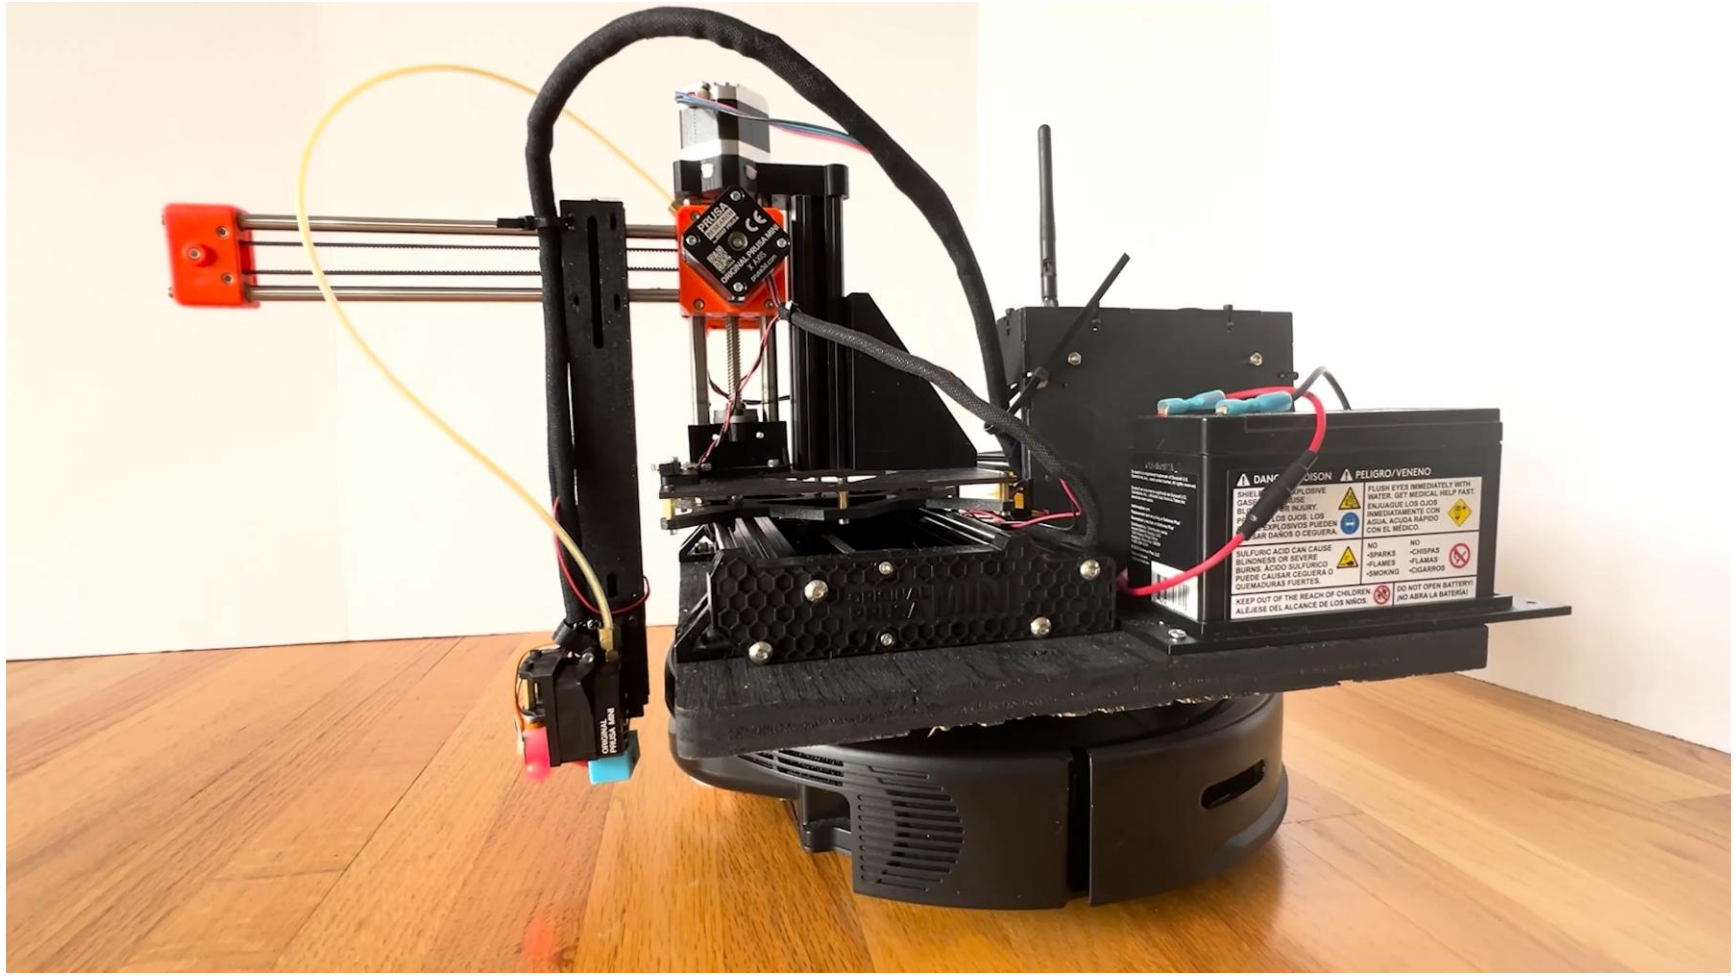


Daniel Campos Zamora

UW CS

On academic job market

- **Robotics**
- **Fabrication**
- **Sensing Techniques**
- **Sustain./Accessibility**



Campos Zamora *et al.,* MobiPrint, UIST'24

MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Chu Li

UW CS

On academic job market

- Human-AI
- Urban Computing
- Visualization
- Accessibility



Daniel Campos Zamora

UW CS

On academic job market

- **Robotics**
- **Fabrication**
- **Sensing Techniques**
- **Sustain./Accessibility**



Towards Rapid Fabrication of Custom Tactile Surface Indicators for Indoor Navigation

Daniel Campos Zamora
Allen School of Computer Science
University of Washington, USA
danielcz@cs.uw.edu

Liang He
Computer Graphics Technology
Purdue University, USA
lianghe@purdue.edu

Jon E. Froehlich
Allen School of Computer Science
University of Washington, USA
jonf@cs.uw.edu

Figure 1: Samples of tactile surface indicators (TSIs) created via rapid prototyping methods (left) such as 3D printing, vacuum forming, CNC milling, and heat transfer melting—(right) which were evaluated through a qualitative study with a blind cane user to validate their tactile saliency as well as an O&M specialist.

Abstract

Tactile surface indicators (TSIs) provide ground-based tactile cues to help pedestrians who are blind or low-vision safely and independently navigate different environments. For example, TSIs can serve as warnings for hazards (*e.g.*, edge of a subway platforms) and directional guides (*e.g.*, a route through a mall). In this exploratory work, we examine how digital fabrication technologies such as 3D printing, CNC milling, vacuum forming, and heat transfer melting can enable the production of custom TSIs. To compare different fabrication approaches, we designed and evaluated a series of prototypes with varied surface materials and design features (*e.g.*, bump height). We then solicited feedback on our ideas and fabricated TSIS via two initial qualitative evaluations: one with a blind cane user and another with an Orientation and Mobility (O&M) specialist. Our initial findings demonstrate that digital fabrication processes—primarily 3D printing and CNC milling—can produce salient and useful TSIs, and indicate interest in our approach and how highly customized, rapidly fabricable TSIs could support navigation in reconfigurable indoor spaces.

CCS Concepts

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

Keywords

Tactile surface indicators, detectable warning surfaces, digital fabrication, 3D printing, blind and low vision

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

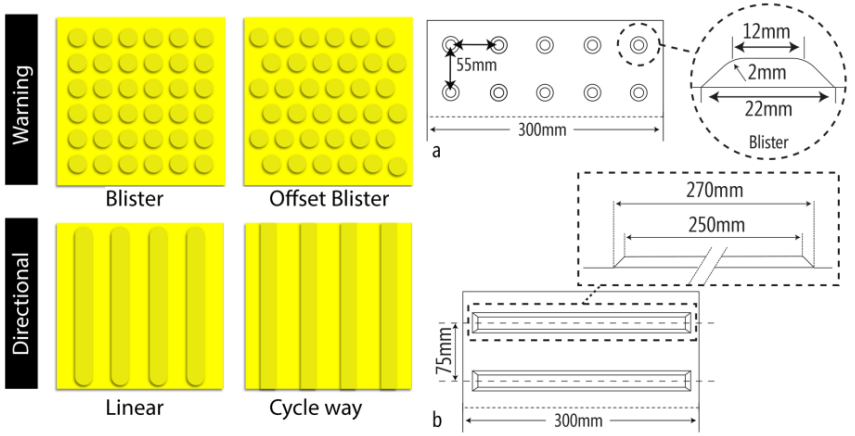
ASSETS '24, October 27–30, 2024, St. John's, NL, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0677-6/24/10
https://doi.org/10.1145/3663548.3688519

ACM Reference Format:
Daniel Campos Zamora, Liang He, and Jon E. Froehlich. 2024. Towards Rapid Fabrication of Custom Tactile Surface Indicators for Indoor Navigation. In *The 26th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '24), October 27–30, 2024, St. John's, NL, Canada.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3663548.3688519

Campos Zamora *et al.*, ASSETS'24





MAKEABILITY LAB

MY GRADUATING PHD STUDENTS



Human-AI

Urban Computing

Visualization

Accessibility

Chu Li

UW CS

On academic job market



Robotics

Fabrication

Sensing Techniques

Sustain./Accessibility

Daniel Campos Zamora

UW CS

On academic job market



Campos Zamora *et al.,* HapticWidgets, *In Submission*


ALWAYS ANOTHER MOUNTAIN
STORIES OF PERSISTENCE, PUSHBACK, & THE PURSUIT OF IMPACT

SEVEN LESSONS

PhD Student — Professor

A MOUNTAIN
2008

A PIVOT
2006

A KNOCKDOWN
2009

A LAB
2011

A CULMINATION
2016

2026

CHI'26 BARCELONA




2008
VIEW FROM MY APARTMENT
BARCELONA CATHEDRAL
Photo by Jon E. Froehlich

2008

SHARED BICYCLING SYSTEMS



Froehlich *et al.*, IJCAI'09







NURIA OLIVER

Photos by Jon E. Froehlich

IJCAI'09

Froehlich, Neumann, Oliver

Sensing & Predicting the Pulse of the City through Shared Bicycling

Sensing and Predicting the Pulse of the City through Shared Bicycling

Jon Froehlich[1], Joachim Neumann[2], Nuria Oliver[2]

[1]Computer Science and Engineering
University of Washington
jfroehli@cs.washington.edu

[2]Telefonica Research, Barcelona, Spain
{joachim, nuriao}@tid.es

Abstract

City-wide urban infrastructures are increasingly reliant on network technology to improve and expand their services. As a side effect of this digitalization, large amounts of data can be sensed and analyzed to uncover patterns of human behavior. In this paper, we focus on the *digital footprints* from one type of emerging urban infrastructure: shared bicycling systems. We provide a spatiotemporal analysis of 13 weeks of bicycle station usage from Barcelona's shared bicycling system, called Bicing. We apply clustering techniques to identify shared behaviors across stations and show how these behaviors relate to location, neighborhood, and time of day. We then compare experimental results from four predictive models of near-term station usage. Finally, we analyze the impact of factors such as time of day and station activity in the prediction capabilities of the algorithms.

1 Introduction

Observing and modeling human movement in urban environments is central to traffic forecasting, understanding the spread of biological viruses, designing location-based services, and improving urban infrastructure. However, little has changed since Whyte (1980) observed in his "Street Life Project" that the *actual* usage of New York's streets and squares clashed with the original ideas of architects and city planners. A key difficulty faced by urban planners, virologists, and social scientists is that obtaining large, real-world observational data of human movement is challenging and costly (Brockman et al., 2006).

As websites have evolved to offer geo-located services, new sources of real-world behavioral data have begun to emerge. For example, Rattenbury *et al.* (2007) and Girardin *et al.* (2008) used geo-tagging patterns of photographs in Flickr to automatically detect interesting real-world events and draw conclusions about the flow of tourists in a city. In addition, as city-wide urban infrastructures such as buses, subways, public utilities, and roads become digitized, other sources of real-world datasets that can be implicitly sensed are emerging. Ratti *et al.* (2006) and González *et al.* (2008) used cellular network data to study city dynamics and human mobility. McNamara *et al.* (2008) used data collected from an RFID-enabled subway system to predict co-location patterns amongst mass transit users. Such sources of data are ever-expanding and offer large, under-explored datasets of physically-based interactions with the real world.

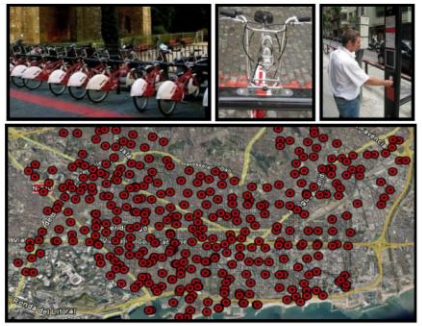


Figure 1. (top) A Bicing station; a close-up of a bicycle and parking slot; and a user at a station kiosk using RFID to check-out a bicycle; (bottom) The 390 Bicing stations distributed across the city of Barcelona, Spain.

In this paper, we introduce a novel source of real-world human behavioral data from a new type of urban infrastructure: shared bicycling systems. We show how station usage data from Barcelona's Bicing system (Figure 1) can be used to infer cultural and geographic aspects of the city and predict future bicycling station usage behavior, which corresponds to human movement in the city.

In particular, the main contributions of this paper are: (1) demonstrating the potential of using shared bicycling as a data source to gain insights into city dynamics and aggregated human behavior; (2) exploring the relationship between spatiotemporal patterns of bicycle usage and underlying city behavior and geography; and (3) studying patterns in bicycle station usage, including the prediction of usage patterns and an analysis of how factors such as the time of the day affect this prediction. In our analysis, we emphasize not just what the bicycling station usage data

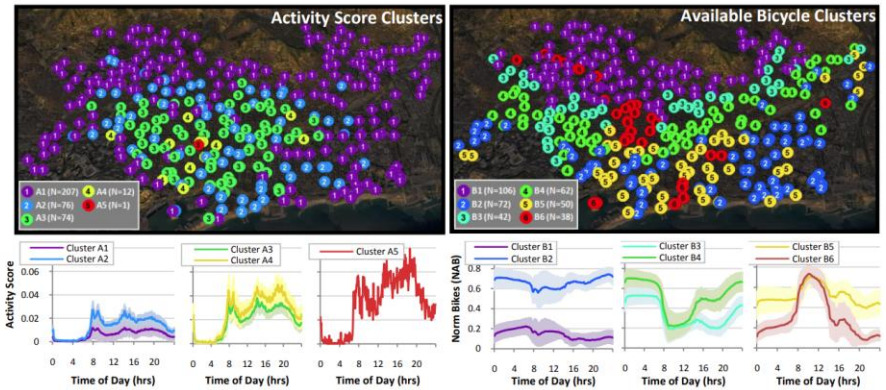


Figure 5. The five Activity Clusters created by progressively combining similar Activity Score station DayViews via dendrogram clustering.

Figure 6. The six Bicycle Clusters created by progressively combining similar available bicycle DayView clusters via dendrogram clustering.

Activity Clusters. The clustering algorithm returned five activity clusters (Figure 5), each with a similar three-pronged spike shape (see also Figure 3). The clusters generally become more active (from A1 to A5) as one moves from the outward edges of the city into downtown. The least active cluster, Cluster A1 (N=207), surrounds nearly the entire perimeter of Barcelona. The more active clusters (A3, A4, and A5) become noticeably more active as the day advances.

Bicycle Clusters. Our algorithm identified six bicycle clusters (depicted in Figure 6) with three classes of behavior: *outgoing* (Clusters B3 and B4), *incoming* (Clusters B5 and B6) and *flat* (Clusters B1 and B2). The *outgoing* clusters show a precipitous drop in available bicycles around 7-8AM as people leave for work, a slight rise at 2-3PM during lunch and a return to early morning levels by 10-11PM. These stations are spread around the edges of the downtown and midtown sections of the city. The *incoming* stations are located in high density commercial areas and along two major arterial routes: Rambla de Catalunya and Avinguda Diagonal. The incoming station shape is nearly inverse of the outgoing stations: people begin arriving around 7-8AM and begin leaving around 1-2PM. Many financial businesses in Barcelona open at 9AM and close around 2:30PM, which aligns well with the temporal patterns of these two clusters. Finally, clusters B1 and B2 have relatively *flat* usage patterns. Cluster B2 tends to have a high degree of available bicycles (on average, it is 66% full) whereas cluster B1 is just the opposite (15% availability). One reason for this discrepancy is likely due to Barcelona's topography: the city itself is built on a long incline. Stations located at the top of Figure 6 are between 80-110 meters above sea level versus those at the bottom, which are at 0-10 meters above sea level. People tend not to bicycle up to the higher altitudes, thus leaving those stations in Cluster B1 starving for bicycles.

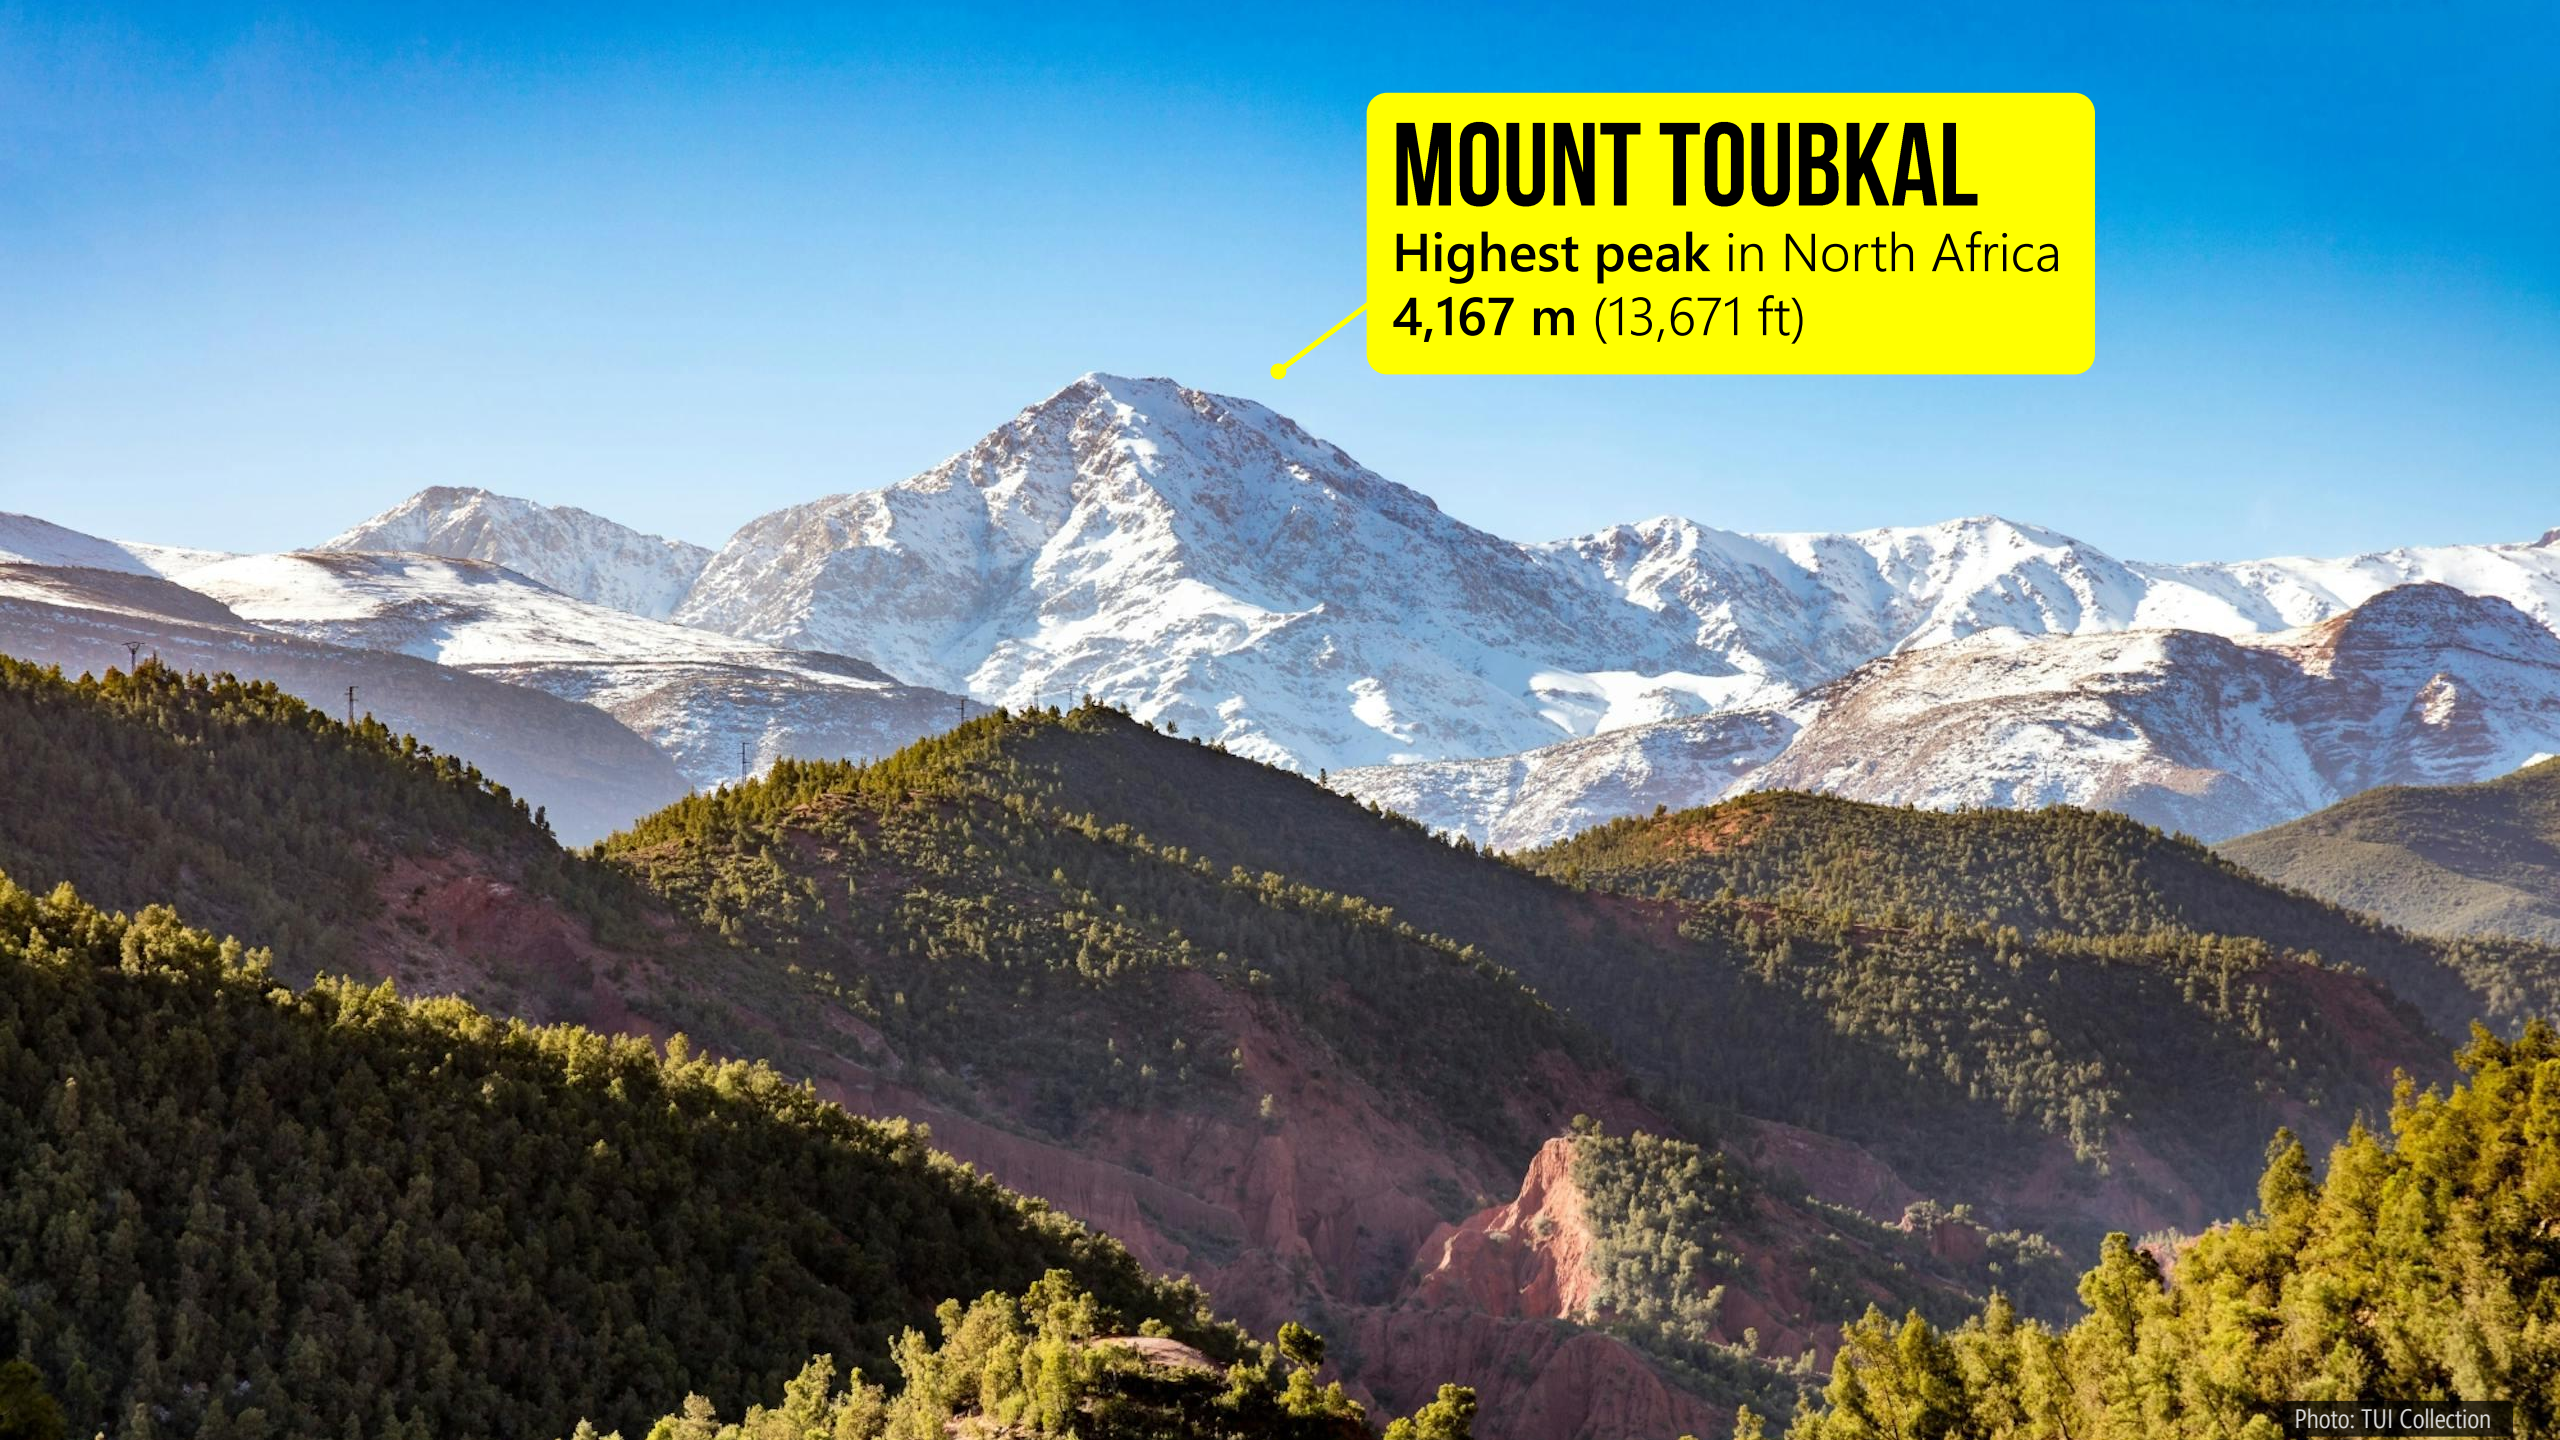
4 Station Behavior Prediction

We focus next on the prediction of station usage. In particular, we are interested in predicting the number of available bicycles at each station at a given time in the future. This work is related to traffic forecasting. Most approaches in traffic engineering rely on flow theory and incorporate queue-theoretic models (Vandaele *et al.*, 2000). Horvitz *et al.* (2005) took an alternative approach that is closely related to ours: they successfully modeled key traffic bottleneck areas using a Bayesian network and ignored the underlying flows. Similarly, we do not attempt to model individual bicycle movements in the city but rather focus on modeling Bicing station usage directly.

Predictive station usage models would (a) allow for more accurate load balancing of the stations; (b) assist urban planners and city officials by providing them with information about *expected* activity in the city; and (c) open the way to new mobile services for Bicing users. In the previously mentioned online survey, Bicing users identified *finding an available bicycle and parking slot* as the two most important problems in their Bicing experience (76% and 66% of respondents, respectively). Therefore, we are also interested in predicting the probability of finding a free bicycle and slot in a station at a given time in the future. Furthermore, our models shed light on some of the factors that influence the *predictability* of station usage behavior.

4.1 Models of Station Behavior

We have implemented four simple predictive models, including a Bayesian network (BN) to predict the availability of bicycles at each station. All models have three input parameters: (1) the current time t_0; (2) the last


MOUNT TOUBKAL
Highest peak in North Africa
4,167 m (13,671 ft)
Photo: TUI Collection

MOUNT TOUBKAL

DAY 1: HIKING TO BASE CAMP











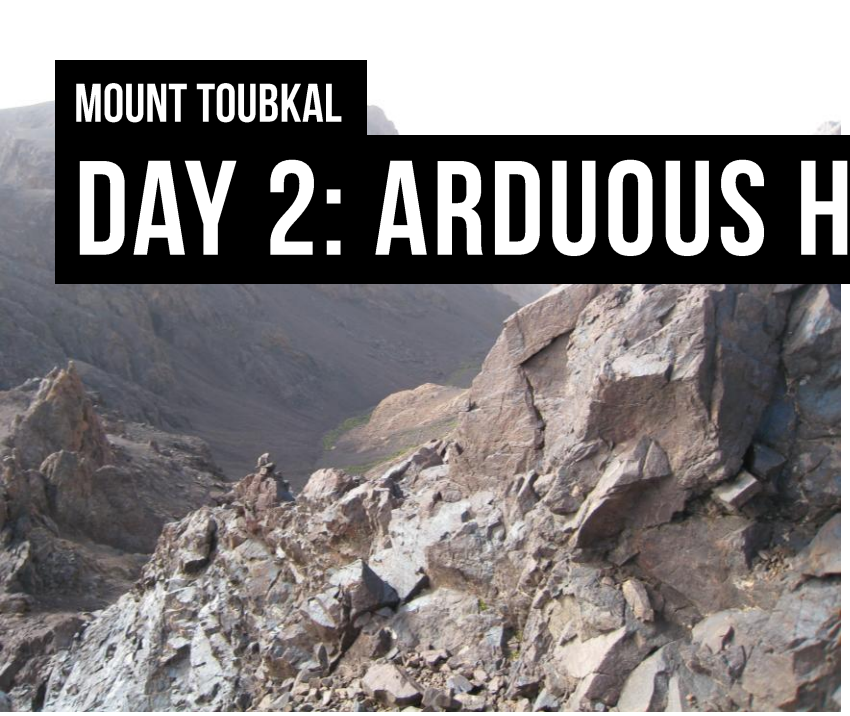
MOUNT TOUBKAL

DAY 2: SUNRISE DEPARTURE



MOUNT TOUBKAL

DAY 2: ARDUOUS HIKE













MOUNT TOUBKAL

DAY 2: THE PEAK



Photo by Jon E. Froehlich

MOUNT TOUBKAL

DAY 2: THE PEAK



Photo by Jon E. Froehlich

MOUNT TOUBKAL

DAY 2: THE PEAK



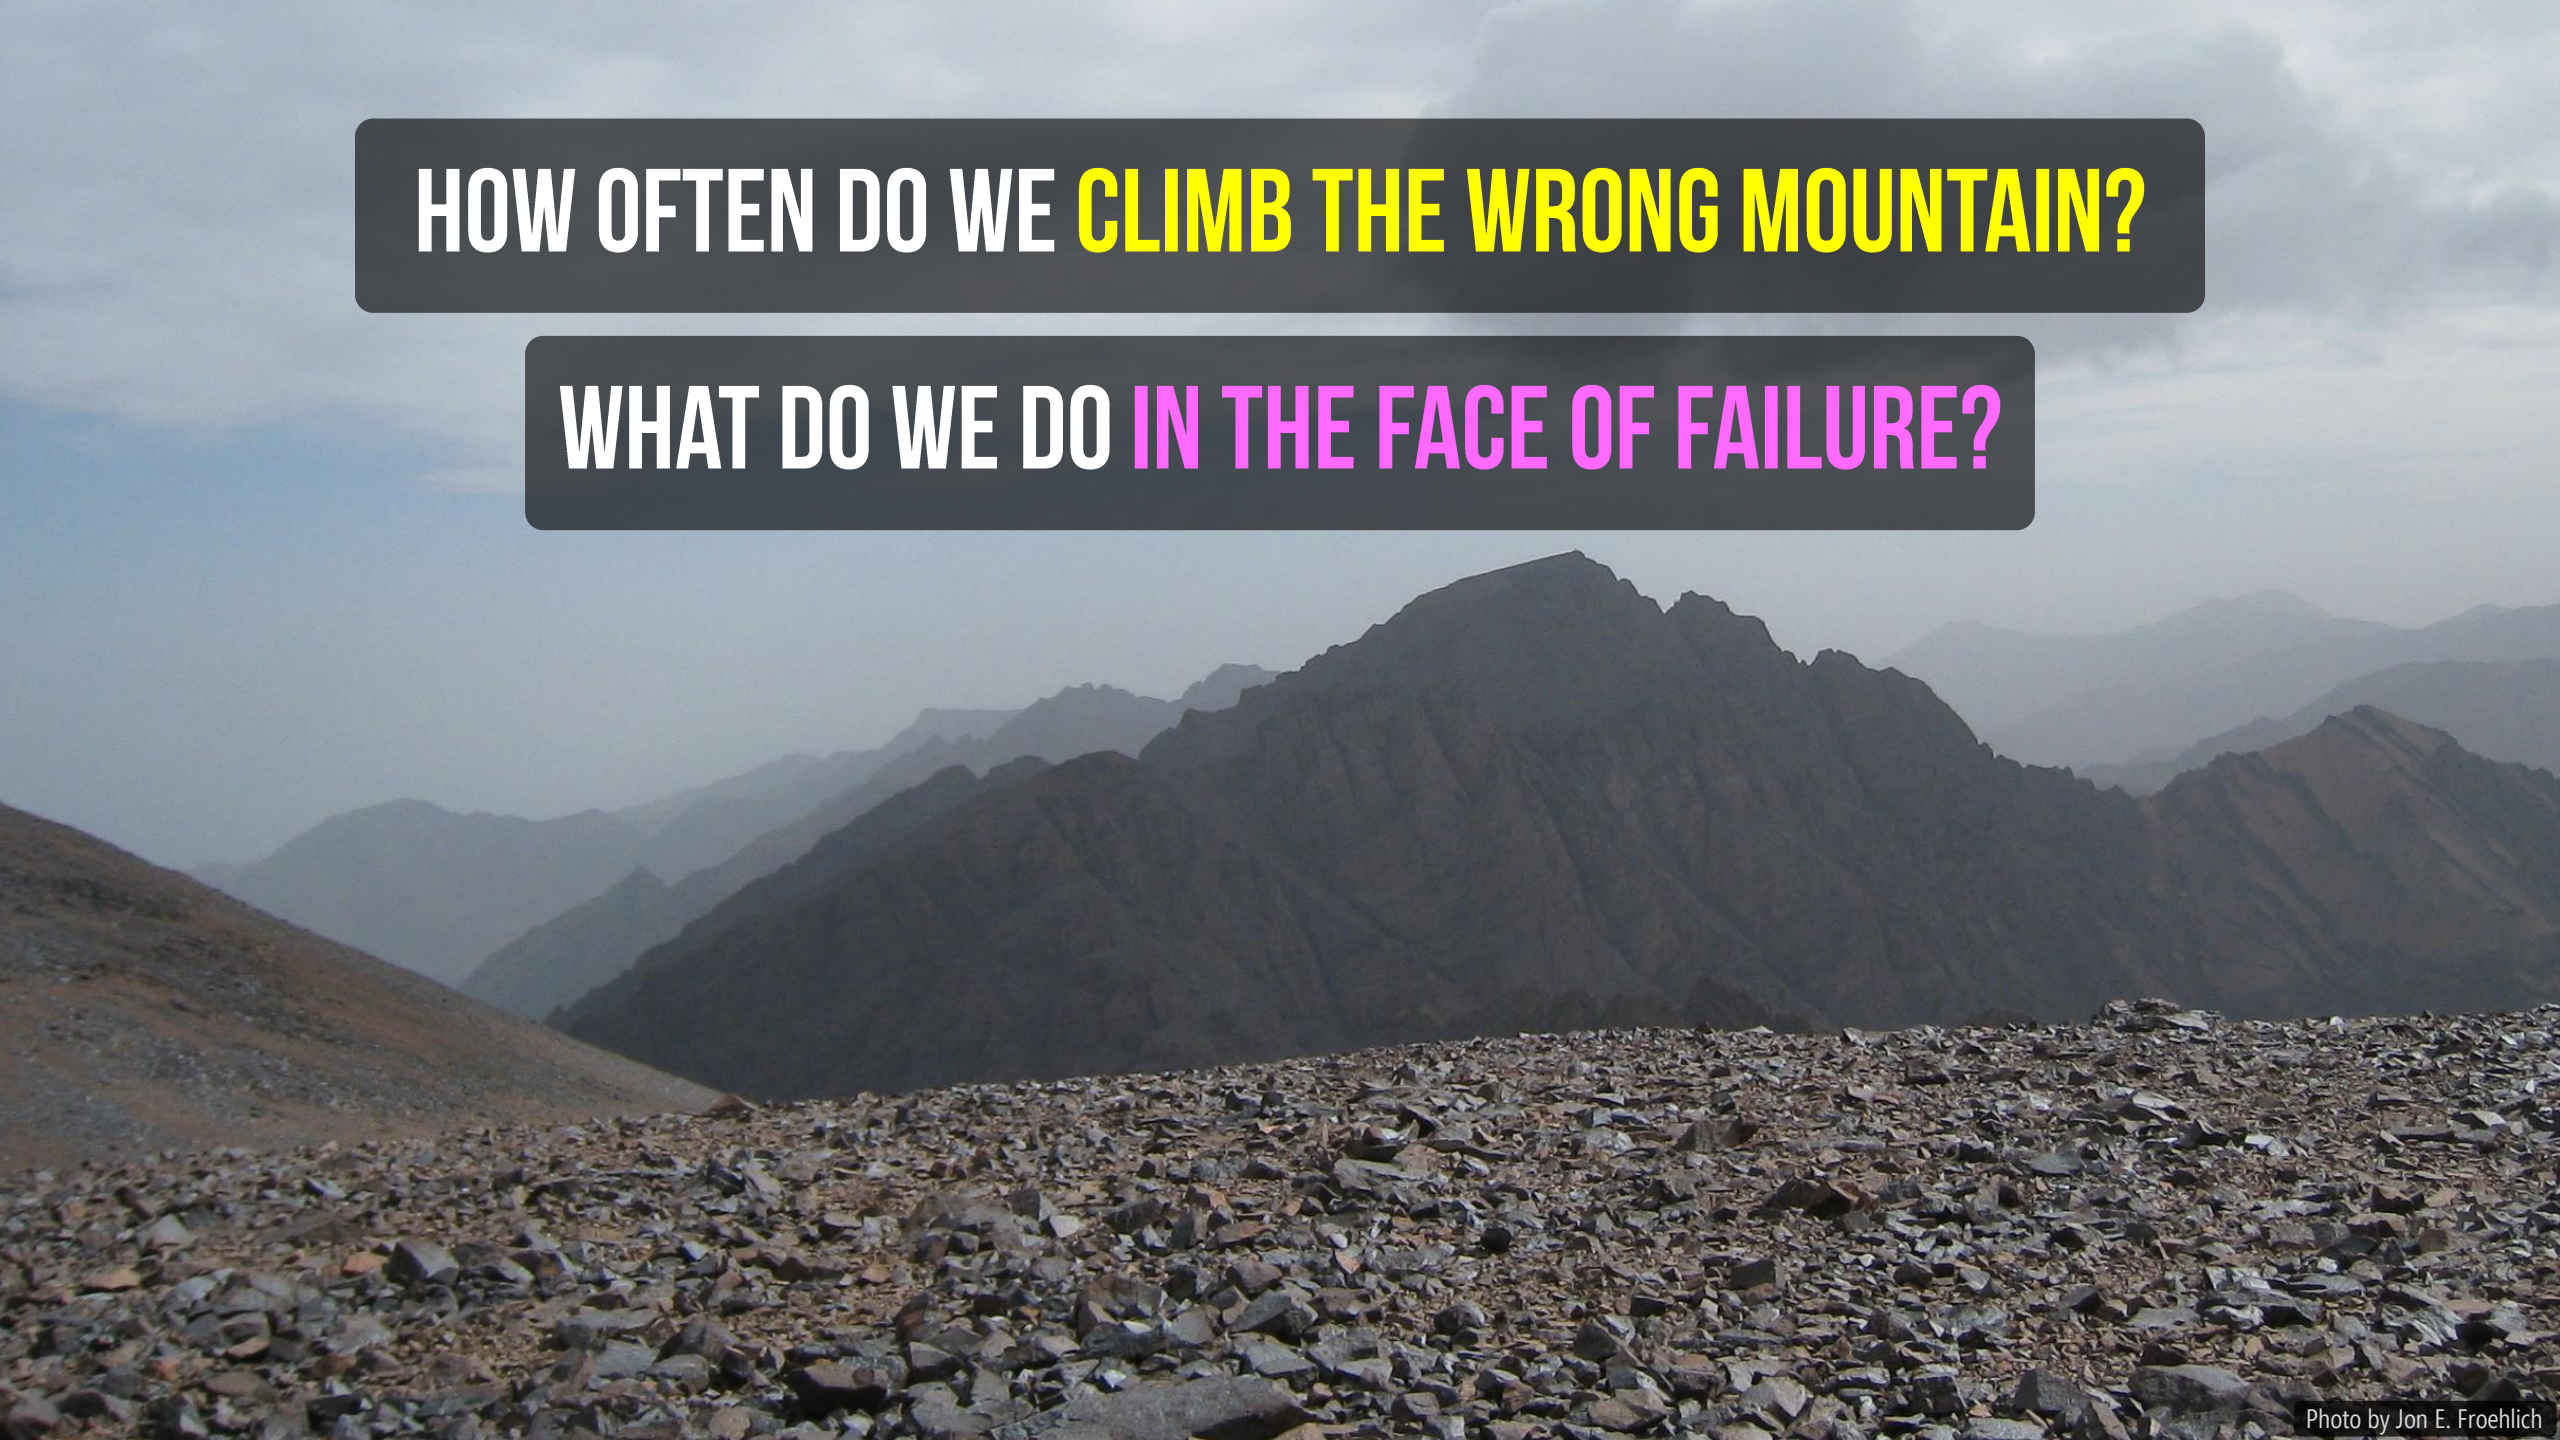
Photo by Jon E. Froehlich


MOUNT TOUBKAL
The Wrong
DAY 2: THE PEAK
NOT ON MOUNT TOUBKAL
Photo by Jon E. Froehlich

HOW OFTEN DO WE CLIMB THE WRONG MOUNTAIN?



Photo by Jon E. Froehlich


HOW OFTEN DO WE CLIMB THE WRONG MOUNTAIN?
WHAT DO WE DO IN THE FACE OF FAILURE?
Photo by Jon E. Froehlich

MOUNT TOUBKAL

DAY 3: BEGIN AGAIN



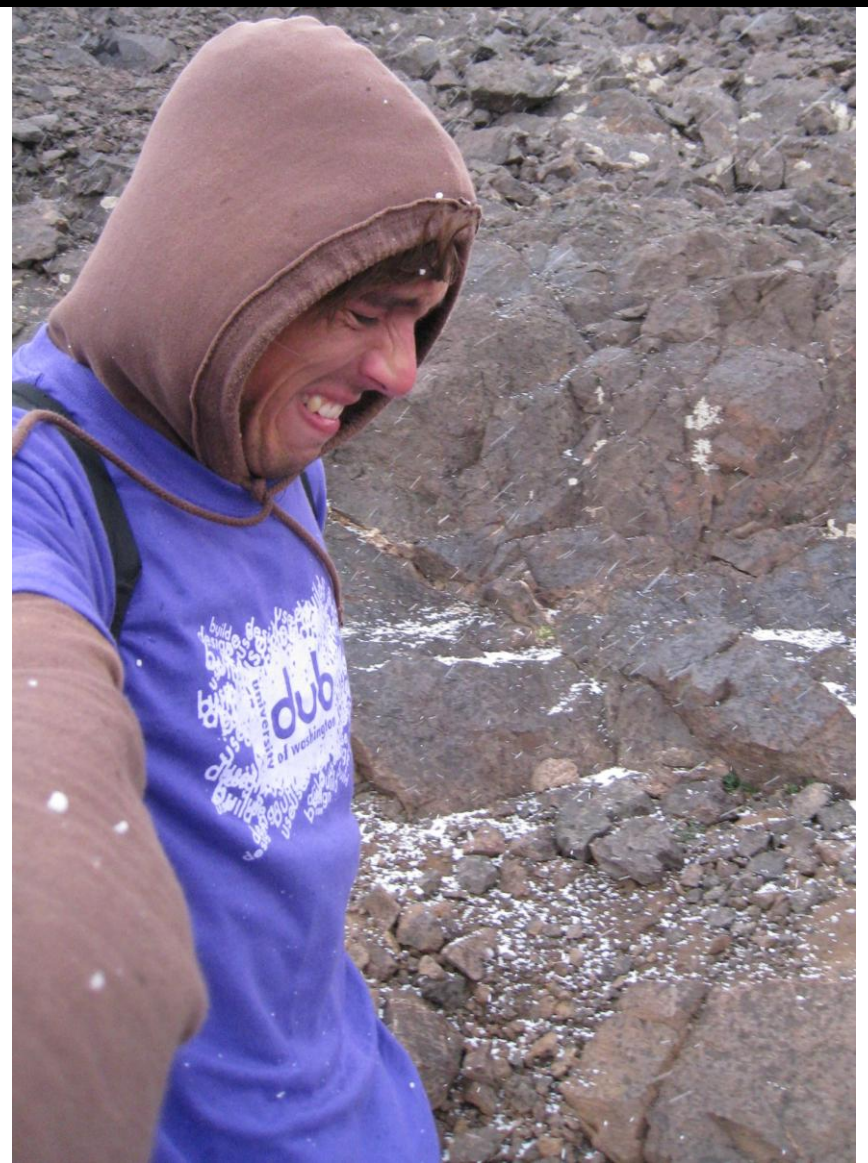
Photo by Jon E. Froehlich

MOUNT TOUBKAL

DAY 3: HIKING THE REAL MOUNT TOUBKAL











Photos by Jon E. Froehlich


Photo by Jon E. Froehlich


Photo by Jon E. Froehlich



Photo by Jon E. Froehlich


TOP OF MOUNT TOUBKAL
Photo by Jon E. Froehlich

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key

PCS

Submissions | Reviews | Account | sign out

Submissions

Make a new submission to | Society: SIGCHI | Conference/Journal | Track | Go

Submissions and associated information, including reviews, are removed after the conference takes place. Please ensure that you download all information, including reviews, before the conference takes place.

PAPER REJECTIONS

Existing Submissions

arch: | Help

Date	Status	Title	Venue
Jan 19, 2026	Not Accepted	Paper #1	DIS 2026 Papers
Dec 04, 2025	Not Accepted	Paper #2	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #3	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #4	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #5	CHI 2026 Papers

Oct 23, 2022	Not Accepted	Paper #23	ASSETS 2023 Technical Papers
Oct 23, 2022	Not Accepted	Paper #24	ASSETS 2023 Technical Papers
Oct 10, 2021	Not Accepted	Paper #25	UIST 2021 Papers
Jan 12, 2021	Not Accepted	Paper #26	CHI 2021 Papers
Jan 12, 2021	Not Accepted	Paper #27	CHI 2021 Papers
Jan 12, 2021	Not Accepted	Paper #28	CHI 2021 Papers
Jan 12, 2021	Not Accepted	Paper #29	CHI 2021 Papers
Oct 26, 2020	Not Accepted	Paper #30	ASSETS 2020 Technical Papers
Oct 26, 2020	Not Accepted	Paper #31	ASSETS 2020 Technical Papers
Apr 25, 2020	Not Accepted	Paper #32	CHI 2020 Papers
Jun 23, 2019	Not Accepted	Paper #33	DIS 2019 Papers
Feb 15, 2019	Not Accepted	Paper #34	IMWUT 2018 February 2019
	Not Accepted	Paper #35	CHI 2022 Papers
	Not Accepted	Paper #36	CHI 2022 Late-Breaking Work

Research.gov
Online Grants Management for the NSF Community

Welcome Jon Froehlich | Sign Out (Home) | My Profile | Contact | Help

My Proposals Results

GRANT REJECTIONS

Agency NSF Show 50 per page

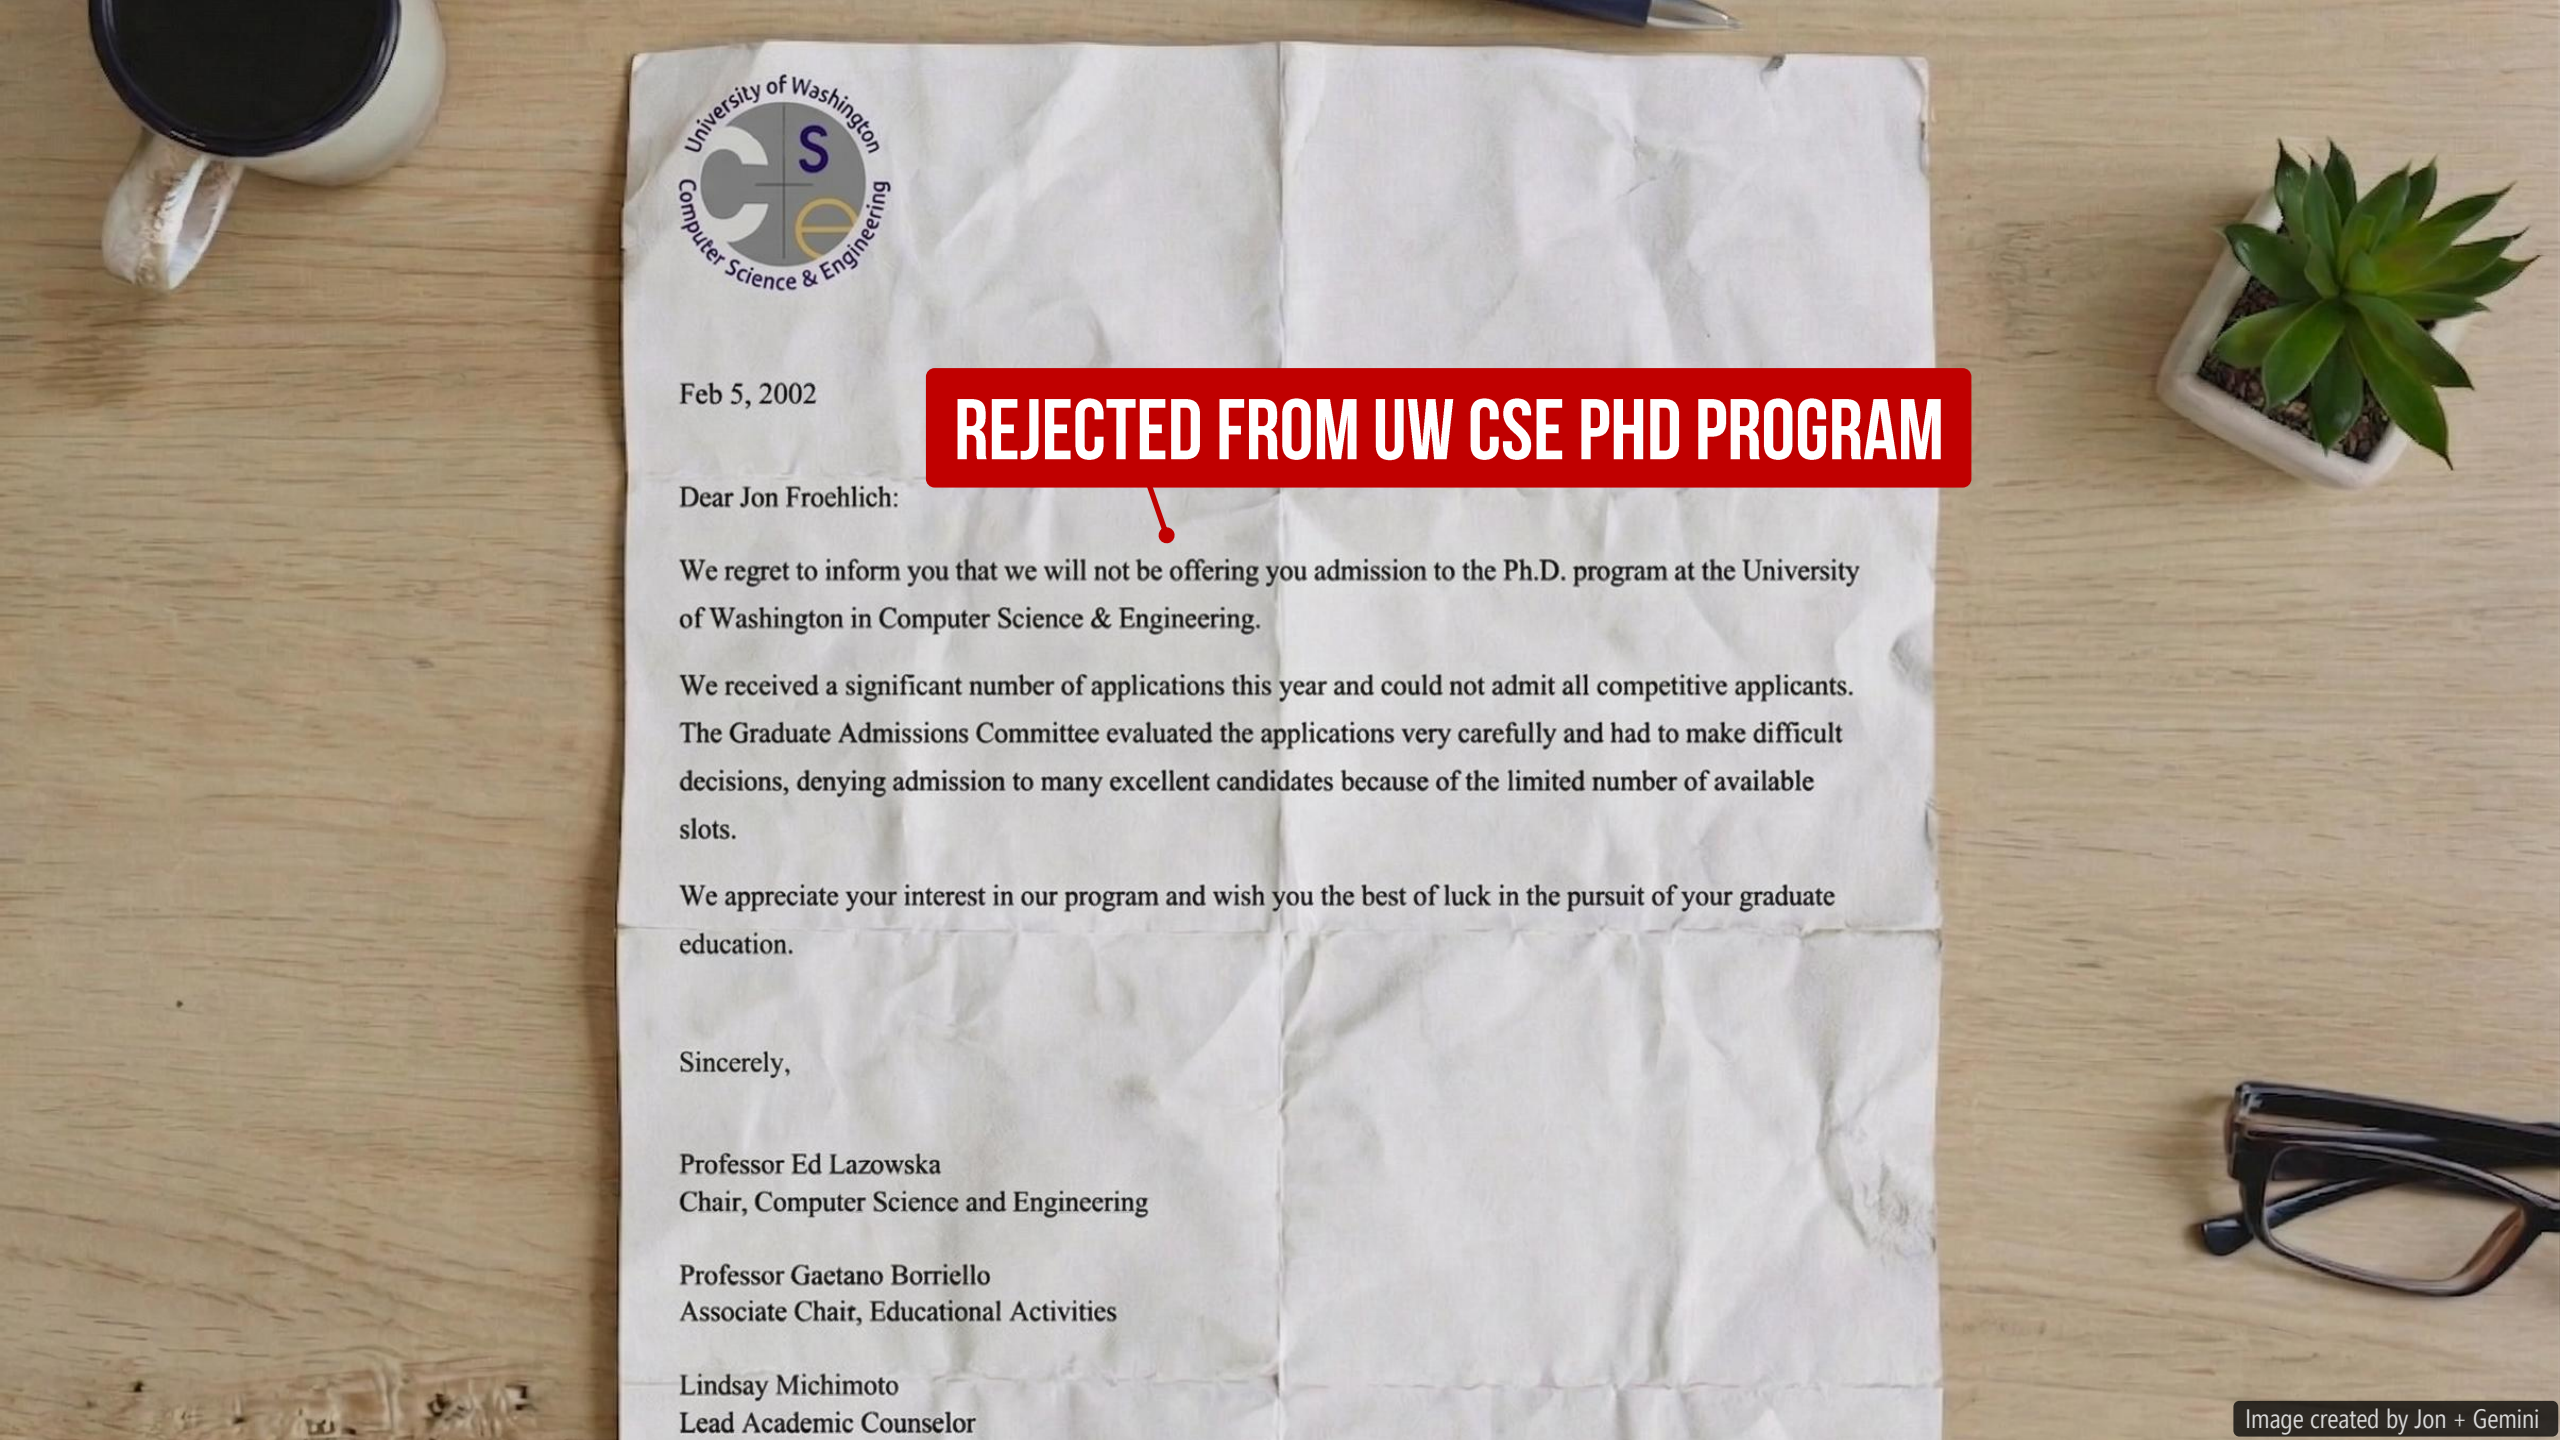
PAGE: 1 of 1

Agency	Agency Tracking Number	Grants.gov Tracking ID	Submitting Institution	Descriptive Title of Project	Status	Status Date	Received Date	Requested Amount
NSF			U of Washington View SAM Legal Business Name	NRT: Innovative Graduate Education and Research Trainee Program...	Declined	03/05/2024	09/06/2023	$3,000,000
NSF			Utah State University View SAM Legal Business Name	NSF Convergence Accelerator Track H: Universal Pathways: AI...	Declined	12/08/2023	08/29/2023	$4,999,987
NSF			U of Washington View SAM Legal Business Name	Collaborative Research: Frameworks: Cyberinfrastructure to Catalyze and Sustain...	Declined	08/06/2023	12/15/2022	$1,047,892
NSF			U of Washington View SAM Legal Business Name	RETTL: Collaborative Research: Making Circuit Prototyping Accessible for...	Declined	06/04/2021	01/20/2021	$425,000
NSF			U of MD, College Park View SAM Legal Business Name	CHS: Medium: Collaborative Research: Wearable Sound Sensing and...	Declined	06/21/2017	10/19/2016	$1,099,579
NSF			U of MD, College Park View SAM Legal Business Name	CAREER: MakerWear: Computational Clothing as a Design Platform...	Declined	12/17/2015	07/21/2015	$694,147
NSF			U of MD, College Park View SAM Legal Business Name	HCC: Small: Personalizing Touchscreen Interaction Through Large-Scale Online...	Declined	09/20/2013	12/17/2012	$496,812

PAGE: 1 of 1

Export: CSV | Excel | XML

Agency NSF Show 50 per page


REJECTED FROM UW CSE PHD PROGRAM
University of Washington
Computer Science & Engineering
Feb 5, 2002
Dear Jon Froehlich:
We regret to inform you that we will not be offering you admission to the Ph.D. program at the University of Washington in Computer Science & Engineering.
We received a significant number of applications this year and could not admit all competitive applicants. The Graduate Admissions Committee evaluated the applications very carefully and had to make difficult decisions, denying admission to many excellent candidates because of the limited number of available slots.
We appreciate your interest in our program and wish you the best of luck in the pursuit of your graduate education.
Sincerely,
Professor Ed Lazowska
Chair, Computer Science and Engineering
Professor Gaetano Borriello
Associate Chair, Educational Activities
Lindsay Michimoto
Lead Academic Counselor
Image created by Jon + Gemini

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key

Submissions and associated information, including reviews, are removed after the conference takes place. Please ensure that you download all information, including reviews, before the conference takes place.

Existing Submissions

PAPER REJECTIONS

arch: Help

Date	Status	Title	Venue
Jan 19, 2026	Not Accepted	Paper #1	DIS 2026 Papers
Dec 04, 2025	Not Accepted	Paper #2	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #3	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #4	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #5	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #6	CHI 2026 Papers
Jul 01, 2025	Not Accepted	Paper #7	VIS 2025 Full Papers
Apr 16, 2025	Not Accepted	Paper #8	ASSETS 2025 Technical Papers
Apr 09, 2025	Not Accepted	Paper #9	UIST 2025 Papers

Submissions and associated information, including reviews, are removed after the conference takes place. Please ensure that you [illegible]ion, including reviews, before the conference takes place.

Existing Submissions

PAPER REJECTIONS

arch: Help

Date	Status	Title	Venue
Jan 19, 2026	Not Accepted	Paper #1	DIS 2026 Pap
Dec 04, 2025	Not Accepted	Paper #2	CHI 2026 Pap
Dec 04, 2025	Not Accepted	Paper #3	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #4	CHI 2026 Papers
Dec 04, 2025	Not Accepted	Paper #5	
Dec 04, 2025	Not Accepted	Paper #6	CHI 2026 Papers
Jul 01, 2025	Not Accepted	Paper #7	VIS 2025 Full Papers
Apr 16, 2025	Not Accepted	Paper #8	ASSETS 2025 Technical Papers
Apr 09, 2025	Not Accepted	Paper #9	UIST 2025 Papers



CHU LI
My PhD Student

CHI'26 BEST PAPER

Research.gov
Online Grants Management for the NSF Community

Sign Out (Home) | My Profile | Contact | Help

My Proposals Results

Agency NSF Show 50 per page

PAGE: 1 of 1

Agency	Agency Tracking Number	Grants.gov Tracking ID	Submitting Institution	Descriptive Title of Project	Status	Status Date	Received Date	Requested Amount
NSF			U of Washington View SAM Legal Business Name	NRT: Innovative Graduate Education and Research Trainee Program...	Declined	03/05/2024	09/06/2023	$3,000,000
NSF			Utah State University View SAM Legal Business Name	NSF Convergence Accelerator Track H: Universal Pathways: AI...	Declined	12/08/2023	08/29/2023	$4,999,987
NSF			U of Washington View SAM Legal Business Name	Collaborative Research: Frameworks: Cyberinfrastructure to Catalyze and Sustain...	Declined	08/06/2023	12/15/2022	$1,047,892
NSF			U of Washington View SAM Legal Business Name	RETTL: Collaborative Research: Making Circuit Prototyping Accessible for...	Declined	06/04/2021	01/20/2021	$425,000
NSF			U of MD, College Park View SAM Legal Business Name	CHS: Medium: Collaborative Research: Wearable Sound Sensing and...	Declined	06/21/2017	10/19/2016	$1,099,579
NSF			U of MD, College Park View SAM Legal Business Name	CAREER: MakerWear: Computational Clothing as a Design Platform...	Declined	12/17/2015	07/21/2015	$694,147
NSF			U of MD, College Park View SAM Legal Business Name	HCC: Small: Personalizing Touchscreen Interaction Through Large-Scale Online...	Declined	09/20/2013	12/17/2012	$496,812

PAGE: 1 of 1

Export: CSV | Excel | XML

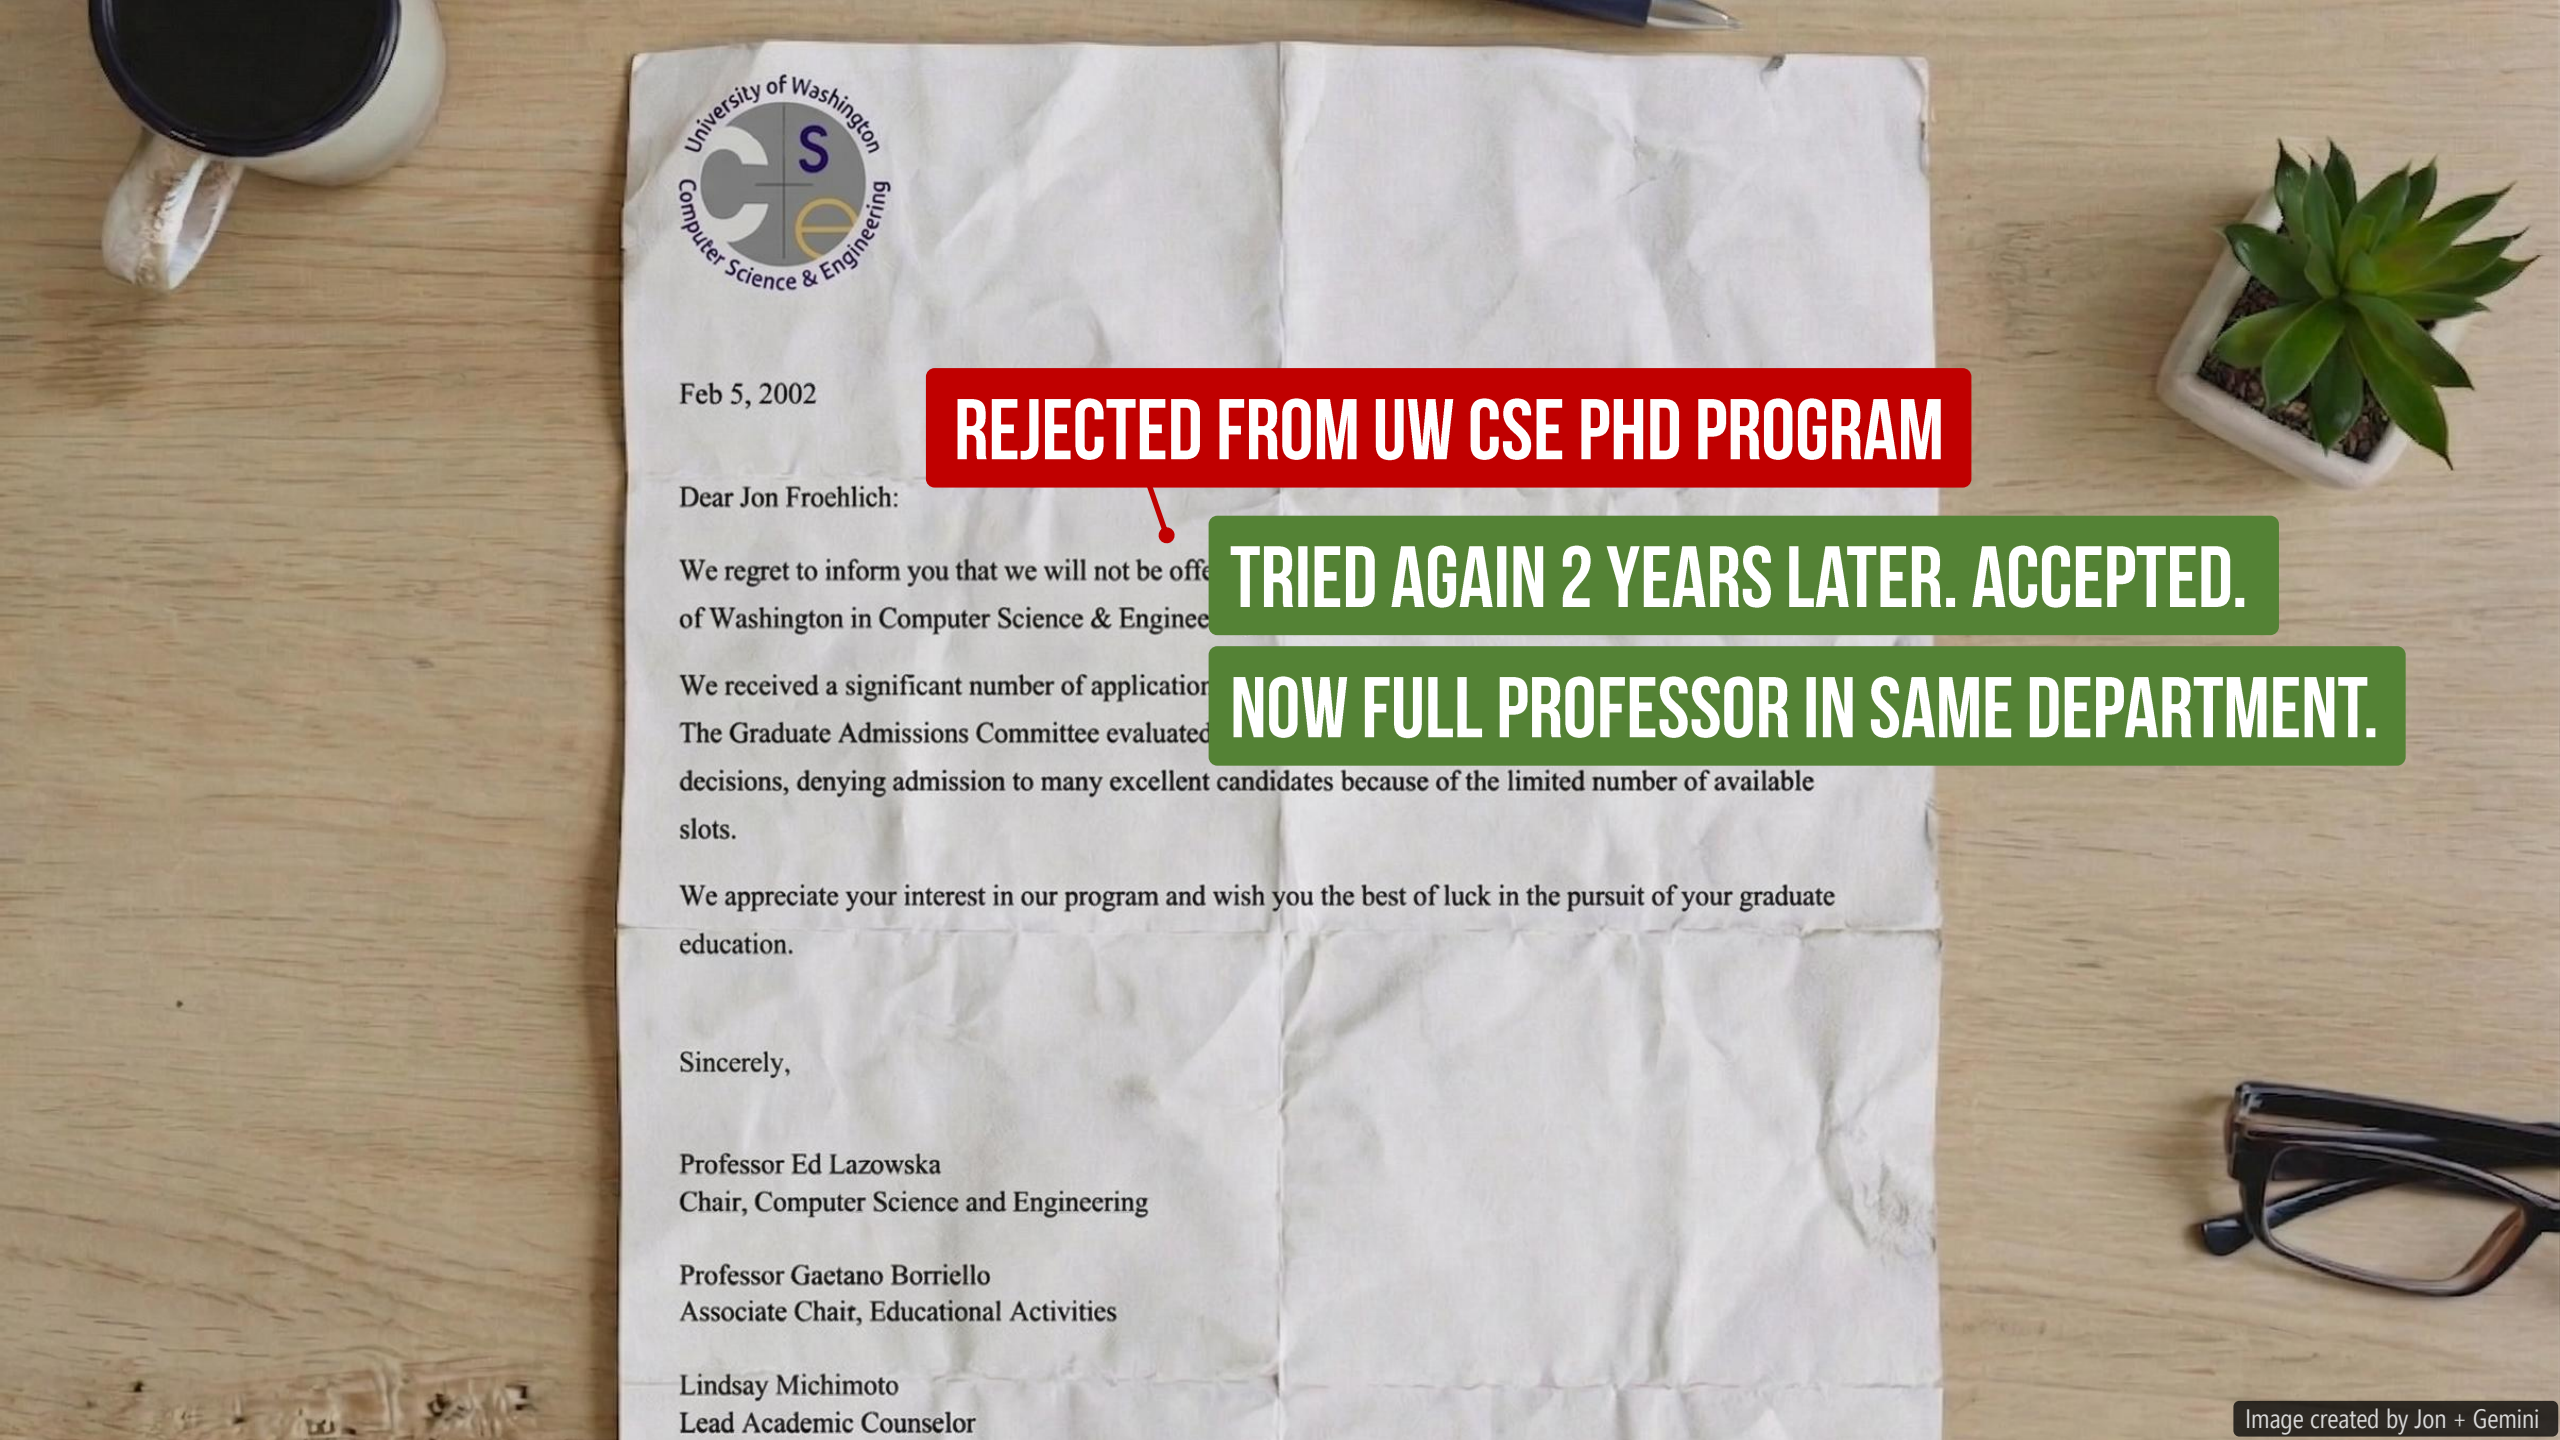
Agency NSF Show 50 per page

GRANT REJECTIONS

AWARDED ON RESUBMIT







University of Washington
Computer Science & Engineering
Feb 5, 2002
Dear Jon Froehlich:
We regret to inform you that we will not be offe
of Washington in Computer Science & Enginee
We received a significant number of applicatio
The Graduate Admissions Committee evaluated
decisions, denying admission to many excellent candidates because of the limited number of available
slots.
We appreciate your interest in our program and wish you the best of luck in the pursuit of your graduate
education.
Sincerely,
Professor Ed Lazowska
Chair, Computer Science and Engineering
Professor Gaetano Borriello
Associate Chair, Educational Activities
Lindsay Michimoto
Lead Academic Counselor
REJECTED FROM UW CSE PHD PROGRAM
TRIED AGAIN 2 YEARS LATER. ACCEPTED.
NOW FULL PROFESSOR IN SAME DEPARTMENT.
Image created by Jon + Gemini

SIGCHI SOCIETAL IMPACT

MY CSE OFFICE



JAMES LANDAY

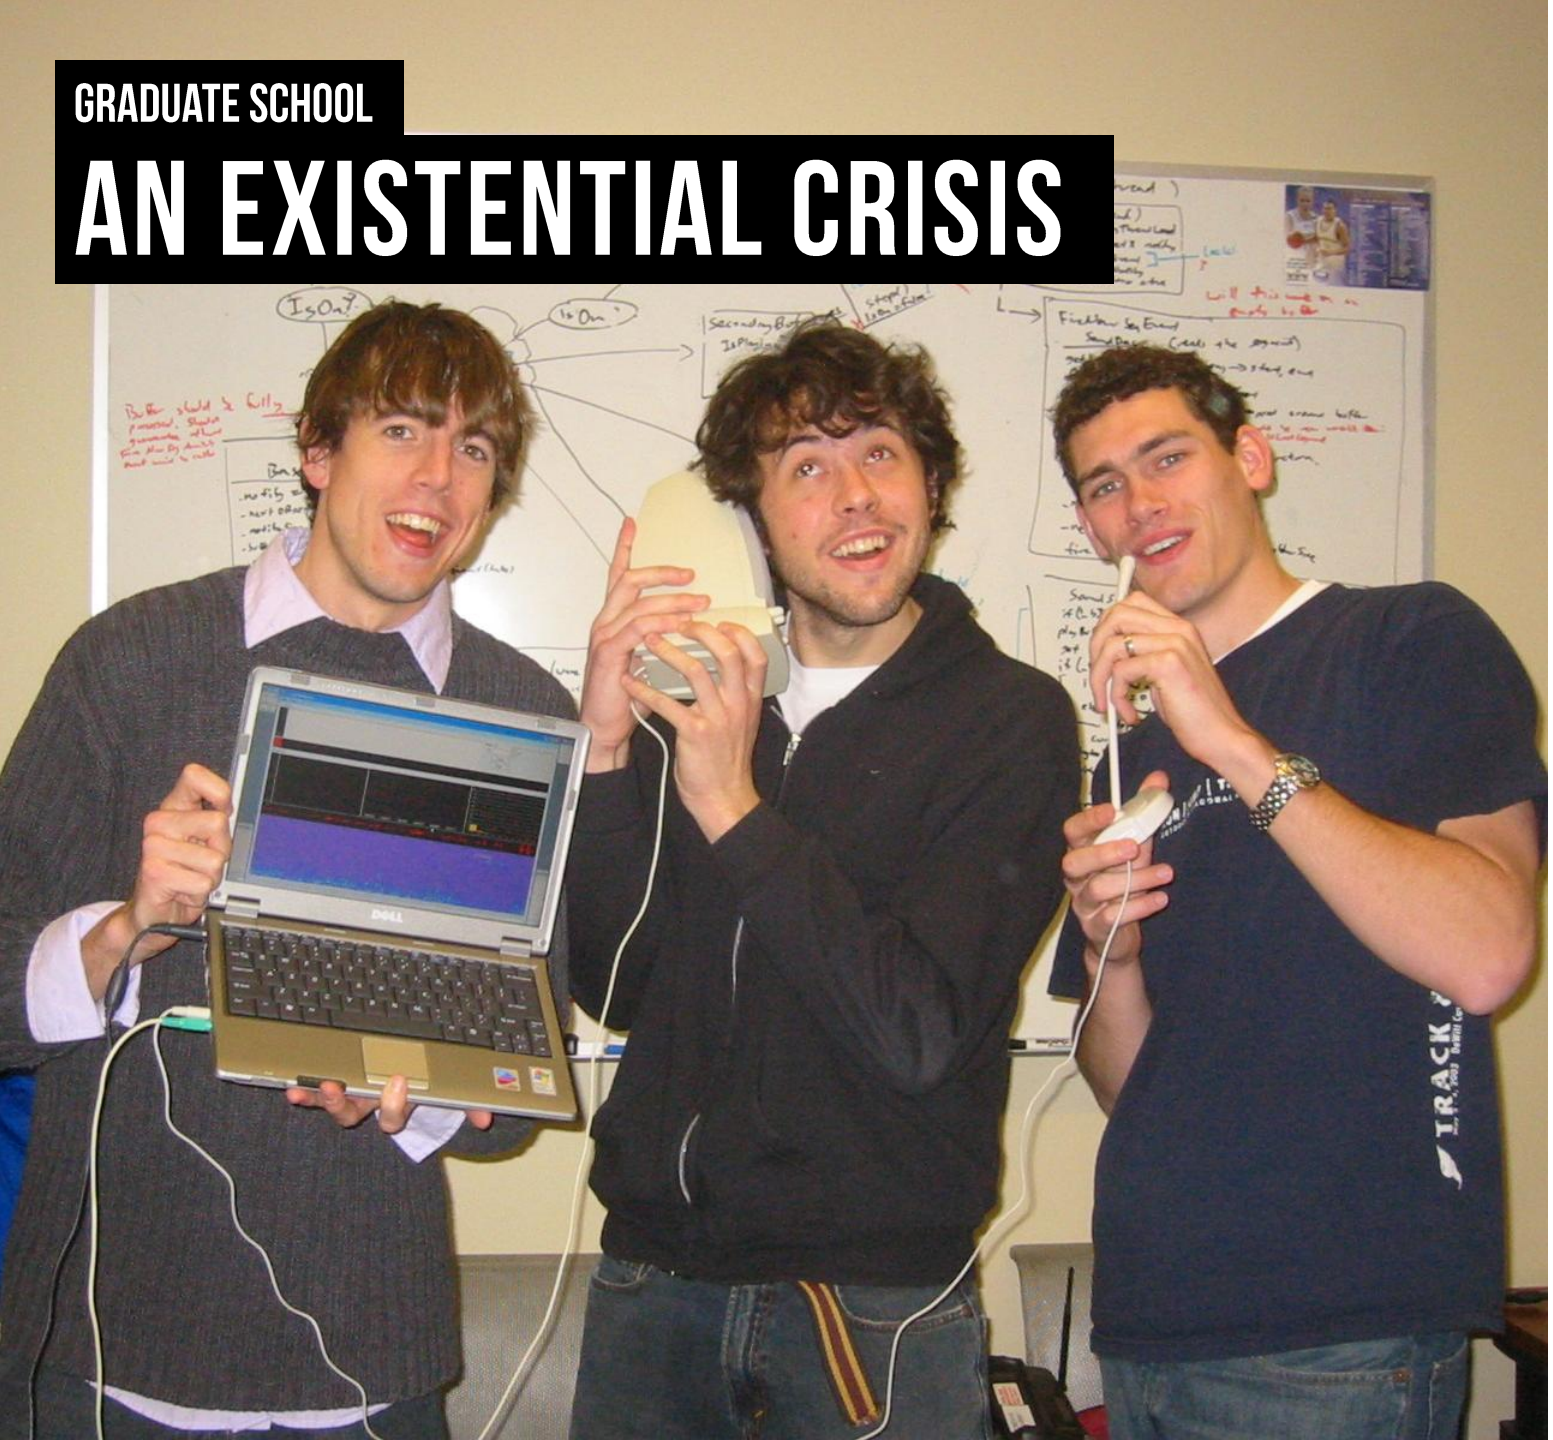
SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

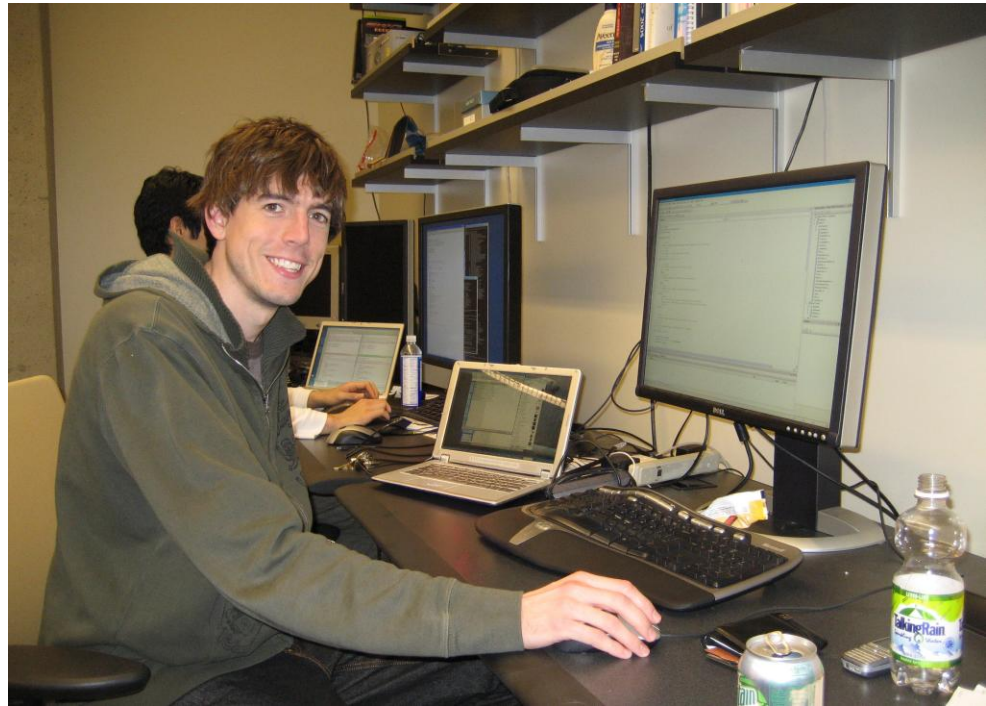
1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key

GRADUATE SCHOOL

AN EXISTENTIAL CRISIS







GRADUATE SCHOOL

NSF PANEL OF GRAD EXPERIENCE



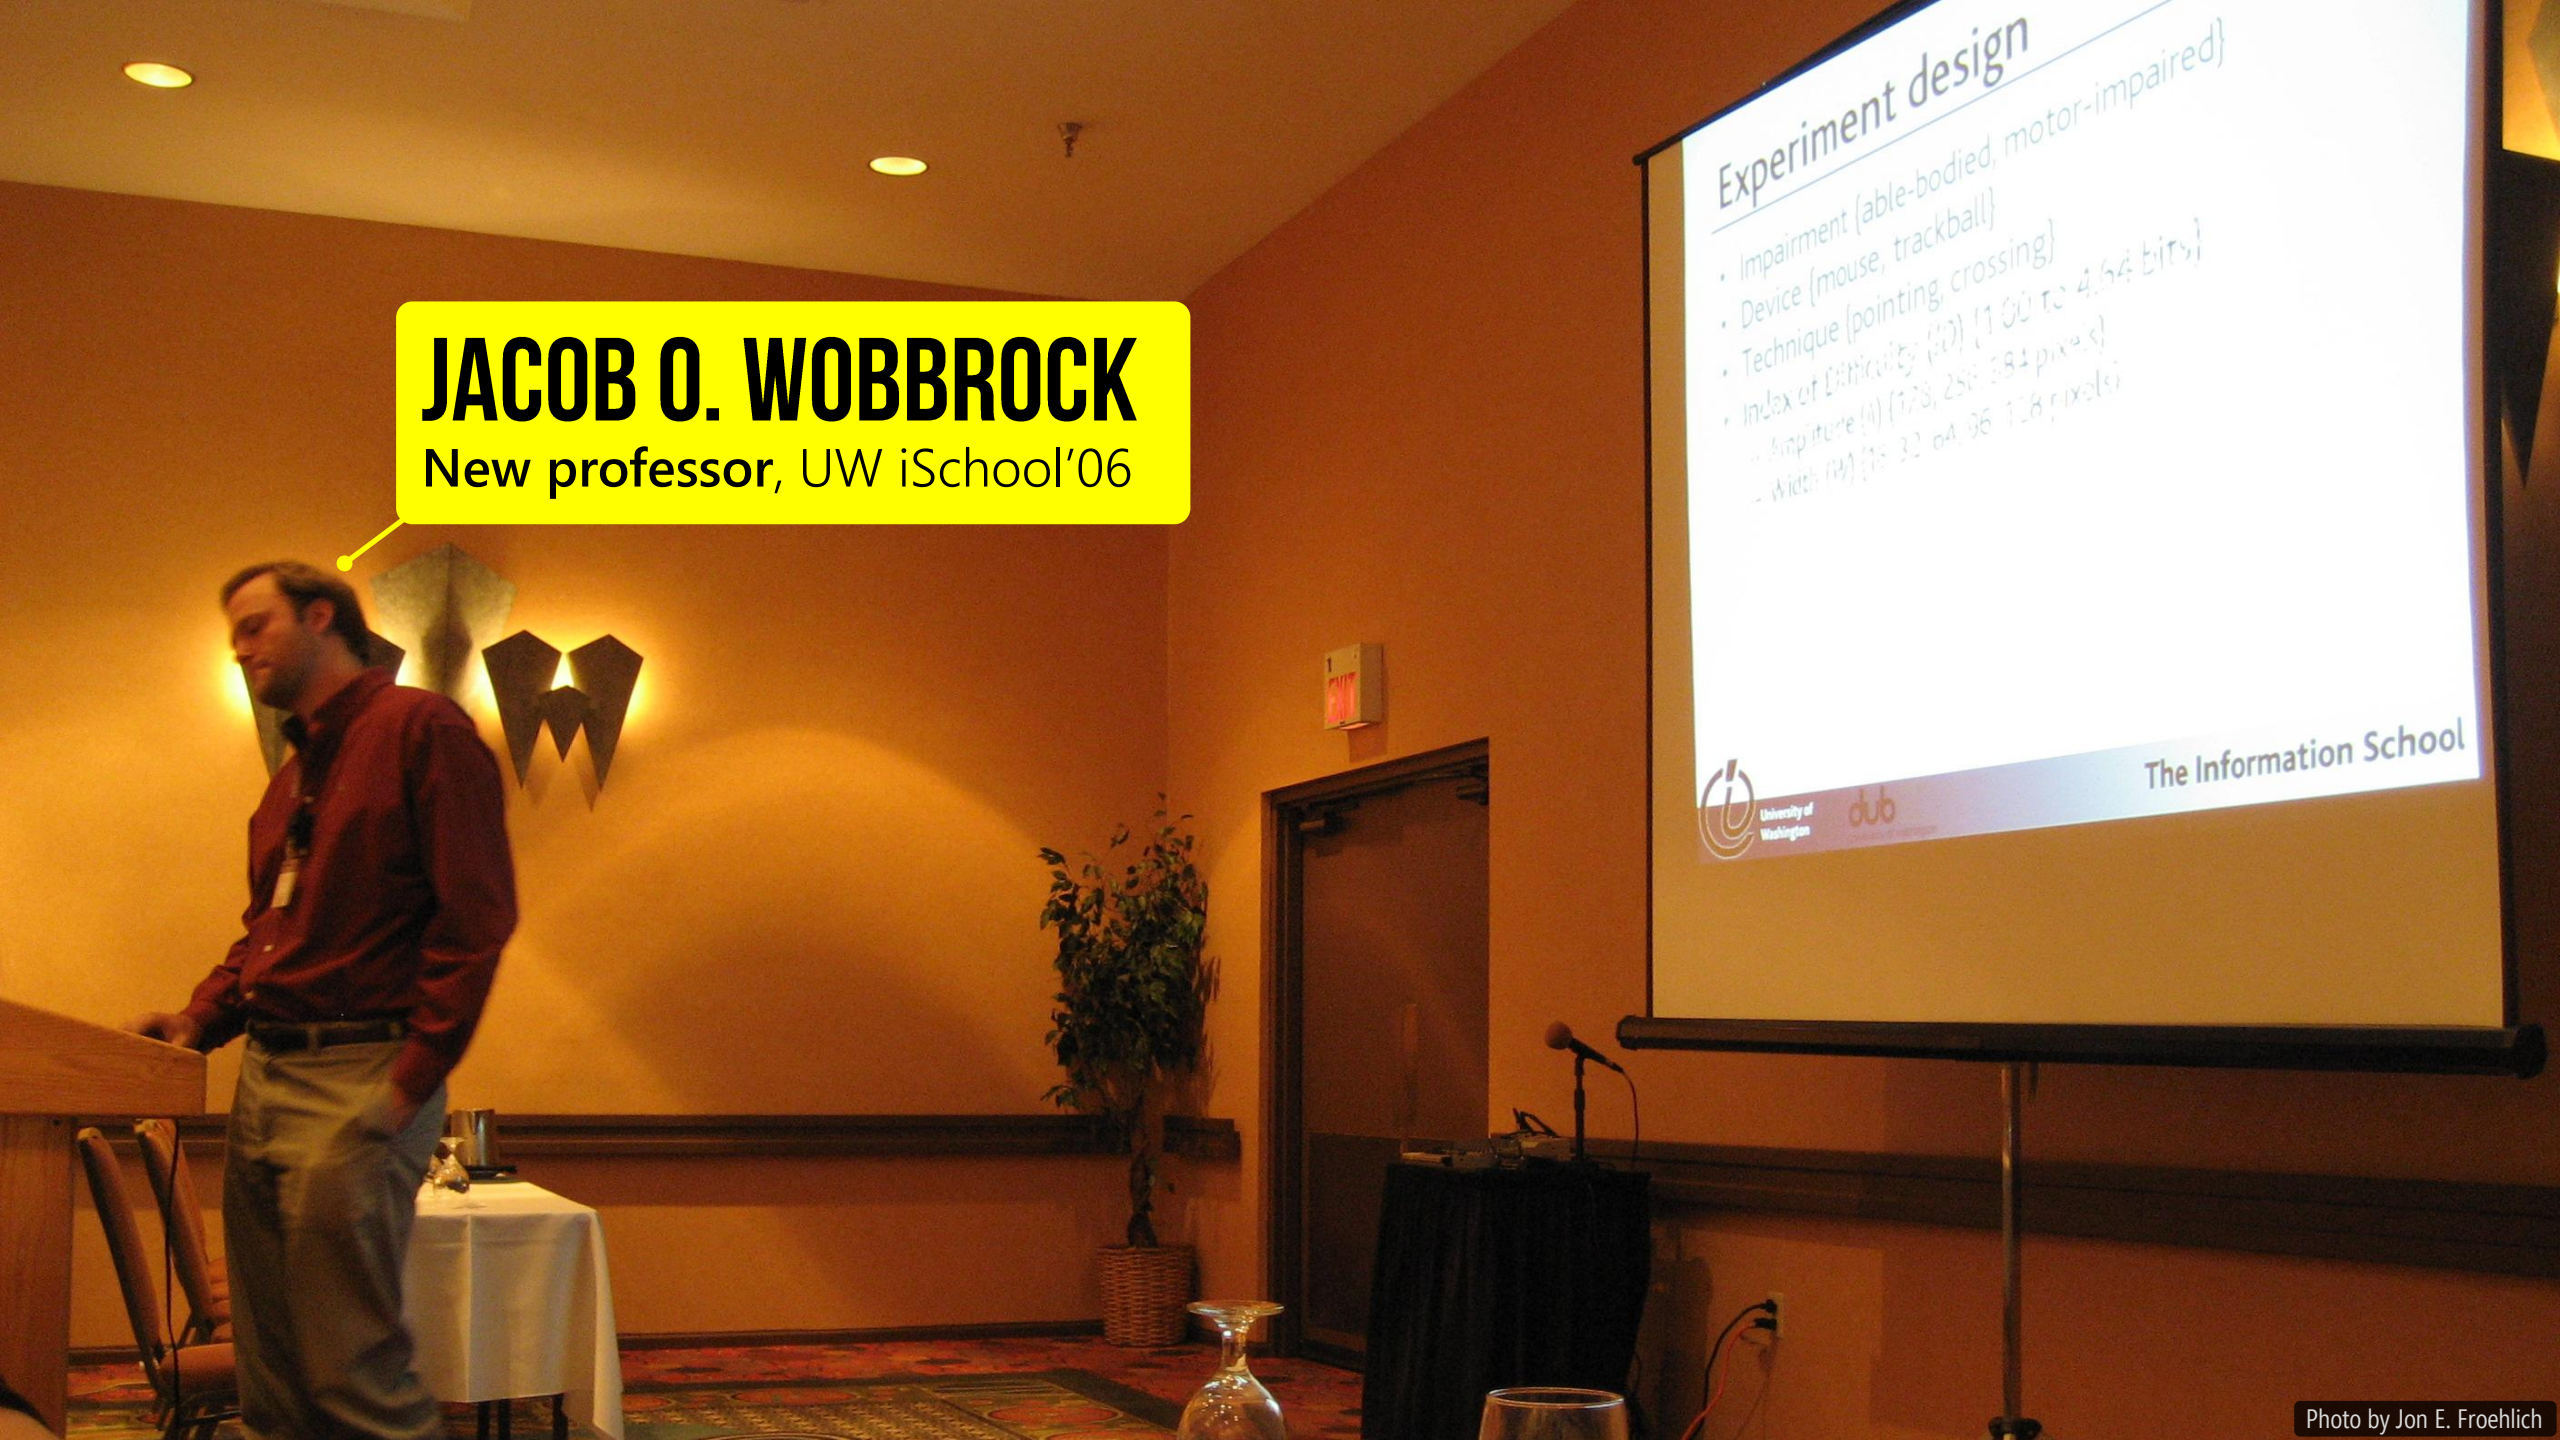
Image created by Jon + Gemini


HOW MANY OF YOU HAVE THOUGHT ABOUT DROPPING OUT?
CONFERENCE ROOM
Image created by Jon + Gemini


CONFERENCE ROOM
I NEEDED TO WORK ON SOMETHING THAT MATTERED
Image created by Jon + Gemini



Image modified by Gemini


JACOB O. WOBBROCK
New professor, UW iSchool'06
Experiment design
The Information School
Photo by Jon E. Froehlich


PDA TOUCH SCREENS
Require fly-in directly to target
Limited tactile feedback
Lift-action difficult
Video by Jon E. Froehlich

ASSETS'07

Froehlich, Wobbrock, Kane

Barrier Pointing: Using Physical Edges to Assist Target Acquisition on Mobile Device Touch Screens



SHAUN KANE

Barrier Pointing: Using Physical Edges to Assist Target Acquisition on Mobile Device Touch Screens

Jon Froehlich,[1] Jacob O. Wobbrock,[2] Shaun K. Kane[2]

[1]Computer Science and Engineering
University of Washington
The Paul Allen Center, Box 352350
Seattle, WA 98195-2840
jfroehli@cs.washington.edu

[2]The Information School
University of Washington
Mary Gates Hall, Box 352840
Seattle, WA 98195-2840
{wobbrock, skane}@u.washington.edu

ABSTRACT

Mobile phones and personal digital assistants (PDAs) are incredibly popular pervasive technologies. Many of these devices contain touch screens, which can present problems for users with motor impairments due to small targets and their reliance on tapping for target acquisition. In order to select a target, users must tap on the screen, an action which requires the precise motion of flying into a target and lifting without slipping. In this paper, we propose a new technique for target acquisition called *barrier pointing*, which leverages the elevated physical edges surrounding the screen to improve pointing accuracy. After designing a series of barrier pointing techniques, we conducted an initial study with 9 able-bodied users and 9 users with motor impairments in order to discover the parameters that make barrier pointing successful. From this data, we offer an in-depth analysis of the performance of two motor-impaired users for whom barrier pointing was especially beneficial. We show the importance of providing physical stability by allowing the stylus to press against the screen and its physical edge. We offer other design insights and lessons learned that can inform future attempts at leveraging the physical properties of mobile devices to improve accessibility.

Categories and Subject Descriptors

H.5.2. [**Information interfaces and presentation**]: User interfaces – *input devices and strategies*. K.4.2. [**Computers and society**]: Social issues – *assistive technologies for persons with disabilities*.

General Terms: Design, Experimentation, Human Factors.

Keywords: Target acquisition, touch screens, edges, corners, accessible interfaces, motor impairments, mobile phones, PDAs.

1. INTRODUCTION

In 2006, the number of mobile phone subscribers in the world surpassed 2.5 billion.[1] This is more than *twice* the number of PC users worldwide. The technology curve shows mobile phones and PDAs getting smaller and more powerful every year, providing features that extend beyond voice calls and text messaging. People with motor impairments, however, often find these devices difficult to use [3][14]. Their reduced size makes input challenging as the buttons are small, condensed and sometimes

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASSETS'07, October 15–17, 2007, Tempe, Arizona, USA.

Copyright 2007 ACM 978-1-59593-573-1/07/0010...$5.00.



(a) Fly-in directly to barrier target (b) Rest the stylus against the screen and stroke into the target.

Figure 1. Two ways the physical properties of a touch screen device can be leveraged to assist movement: (a) The screen's physical edge catches the stylus as it flies into a target. (b) The screen itself provides an additional flat plane to aid movement.

recessed. Prior work has shown that motor-impaired users may not have the physical strength to press hard physical buttons [14]. And certain device form factors (e.g., clamshell phones) require dexterity and strength just to open.

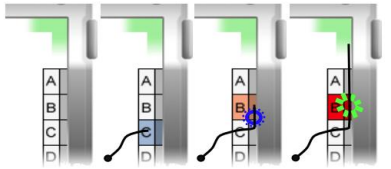
Many emerging high-end phones such as the Microsoft Windows Mobile Smartphone and the Apple iPhone contain touch screens. The iPhone features only one physical button, relying instead almost exclusively on touch screen interactions.[2] However, touch screens pose an additional set of challenges for mobile device accessibility: they lack tactile feedback and the physical stability available with hard buttons, and their interfaces often require dexterous motor control.

Still, touch screen devices offer rich potential for motor-impaired users. Myers et al. [14] demonstrated the use of Palm PDAs as viable alternatives to the mouse and keyboard for users with muscular dystrophy. Touch screens have the potential to support more fluid interactions than their button-based counterparts as they are *direct interaction* systems; that is, the point of stylus or finger contact is co-located with the screen's output. Thus, unlike traditional mobile phones, where target selection is accomplished *indirectly* via button presses and/or joystick movements, touch screen-based targets can be selected and manipulated directly.

Wobbrock et al. [24][25] demonstrated how physical edges could be used to assist motor-impaired users with touch screen-based text entry. Walker and Smelcer [21] demonstrated the benefits of impenetrable borders on targets to reduce mouse movement time. Our work expands into the realm of target acquisition on mobile

[1] http://www.wirelessintelligence.com/

[2] http://www.apple.com/iphone/



(a) *Normal* (b) *Hover* – Stylus over target *C* (c) *Select* – Stylus into barrier trough for *B* (d) *Confirm* – Rapidly move stylus towards screen corner

Figure 5. The "velocity stroke with corner confirmation" interaction. By eliminating the need for lift, the stylus is able to rest on the screen throughout a series of interactions; however, the user must reach the escape velocity before exiting the selected target.

3.2.2 *Velocity Stroke with Corner Confirmation*

This interaction is similar to the previous technique except for the confirmation step. Rather than lift, the user must stroke into a *confirmation corner* to confirm his/her selection (Figure 5). Confirmation corners allow the user to operate the device with the stylus continuously pressed against the screen. Once a target is selected and a confirmation corner is accessed, the user may begin another target acquisition *without lifting* the stylus.

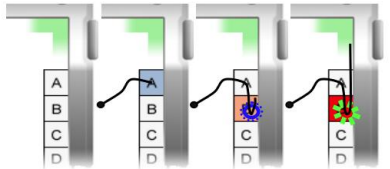
In order to distinguish between a confirmation corner stroke and a target selection stroke, a speed threshold ("escape velocity") is used as a mode switch. The user must maintain this escape velocity after exiting a selected target until the stylus reaches the confirmation corner. If the velocity is not maintained, the current selection is canceled or if the stylus is over another target's trough, this new target becomes selected. Currently, this threshold is set to 200 pixels/second, which was derived empirically with able-bodied testers. This value is slightly slower than the threshold used to ignore mouse clicks in *Steady Clicks* [18].

3.2.3 *Reverse Stroke with Corner Confirmation*

Because moving the stylus at escape speeds may be difficult for motor-impaired users, we developed an alternative method for selection. The "reverse stroke with corner confirmation" technique uses a direction reversal stroke along the target's edge to indicate a selection. For example, if the user is running south along the right edge of the screen, the user would reverse directions (by stroking north) within a target to select that target (Figure 6). The user must then move the stylus along the edge to a corner to confirm. Note that unlike the previous two barrier techniques, "reverse stroke with corner confirmation" does not use selection troughs—thus, its selectable area is the full size of the widget. We found this made it easier for subjects to reverse directions accurately. Note that the direction change can be made as slowly or quickly as desired; speed is not an issue.

4. INITIAL STUDY

In order to better understand barrier pointing, we conducted an initial study of 18 subjects, 9 of whom had motor impairments and 9 of whom did not. For each subject, we administered the 3 aforementioned barrier techniques in addition to the two baseline fly-in-and-tap conditions. This study was meant to provide both quantitative and qualitative insights into how barrier techniques may be used by motor-impaired and able-bodied users.



(a) *Normal* (b) *Hover* – Stylus over target *A* (c) *Select* – Reverse direction against edge while over target *B* (d) *Confirm* – Stylus continues in north direction towards corner

Figure 6. The "reverse stroke with corner confirmation" barrier interaction. For illustrative purposes, these figures demonstrate a reverse stroke slightly away from the edge. The ideal interaction, however, is simply running the stylus along the edge and reversing direction.

4.1 Method

4.1.1 *Subjects*

Eighteen subjects volunteered for the initial study: 9 able-bodied (AB) and 9 motor-impaired (MI). Two females participated in the AB group and 4 in the MI group. The average age was 23.9 (*SD*=2.4) and 37.7 (12.5) for the AB and MI groups, respectively. All subjects indicated that they used a desktop computer for at least an hour every day. Similarly, all subjects, with the exception of one MI participant, owned and used a mobile phone multiple times a day. Two MI subjects (MI1 and MI6) indicated that they currently own a touch screen mobile device. The MI subjects and their health conditions are listed in Table 1. Note that the motor impairments of our subjects covered a very broad range. This was by design, as we wanted to determine what types of users might benefit most from barrier pointing.

4.1.2 *Apparatus*

The initial study was conducted on two HTC Wizard Pocket PC Phones running Windows Mobile 5. One of the devices was outfitted with Velcro and attached to a clipboard for stability; the other was held in the hand. Participants were able to choose either device. The HTC Wizard is 5.3 ounces, 2.83"×4.25"×0.93" in size, and contains a 1.75"×2.25" single-touch touch screen with a QVGA resolution of 320×240 pixels. The touch screen sampling rate was measured empirically at ~100 Hz. The screen edge is 1.7 mm high. As the HTC Wizard ships with a small, ponderous stylus, we offered two alternatives: (1) a thick 5.75" stylus with a rubber grip, or (2) a 6.25" stylus, which was much lighter than (1) and about 40% smaller in diameter.

Subj	Sex	Age	Condition
MI1	M	48	Neuromuscular condition, low strength
MI2	F	55	Parkinson's, slight jitter, good motor control
MI3	M	38	Limited hand function, limited trunk balance
MI4	M	39	Tetraplegia (SCI C5), no use of triceps, pectorals, hands.
MI5	M	50	Limited coordination, episodic peripheral neuropathy
MI6	F	21	Cerebral Palsy (CP)
MI7	F	41	Degeneration of neck and spine
MI8	M	23	Spastic CP, lack of fine motor skills
MI9	F	24	Poor motor control as result of injury

Table 1. MI Study Participants


BARRIER POINTING
Screen supports motor action
Edges guide movement
Smoothing reduces errors
Video by Jon E. Froehlich

SIGCHI SOCIETAL IMPACT TALK

SOCIAL IMPACT AREAS



ACCESSIBILITY

HEALTH & WELLNESS

ENVIRONMENTAL SUSTAINABILITY

STEM EDUCATION

Icons from the Noun Project: Heart by Timothy Miller; Leaf by NeueDeutsche; Microscope by Maxim Kulikov; Access icon from http://accessibleicon.org

SIGCHI SOCIETAL IMPACT TALK

SOCIAL IMPACT AREAS



ACCESSIBILITY



HEALTH
& WELLNESS



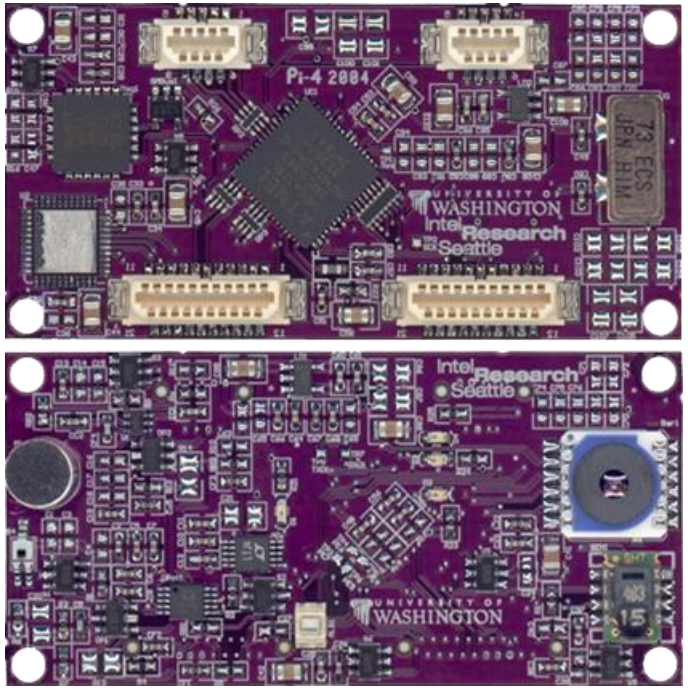
ENVIRONMENTAL
SUSTAINABILITY

STEM
EDUCATION

Icons from the Noun Project: Heart by Timothy Miller; Leaf by NeueDeutsche; Microscope by Maxim Kulikov; Access icon from http://accessibleicon.org

INTEL RESEARCH

MOBILE SENSING PLATFORM





Wearable activity tracking device

The Mobile Sensing Platform:

An Embedded Activity Recognition System

The MSP is a small wearable device designed for embedded activity recognition with the aim of broadly supporting context-aware ubiquitous computing applications.

Tanzeem Choudhury
Dartmouth College
Sunny Consolvo, Beverly Harrison, Jeffrey Hightower, Anthony LaMarca, Louis LeGrand, Ali Rahimi, and Adam Rea
Intel Research
Gaetano Borriello, Bruce Hemingway, Predrag "Pedja" Klasnja, Karl Koscher, James A. Landay, Jonathan Lester, and Danny Wyatt
University of Washington
Dirk Haehnel
Stanford University

Activity-aware systems have inspired novel user interfaces and new applications in smart environments, surveillance, emergency response, and military missions. Systems that recognize human activities from body-worn sensors can further open the door to a world of healthcare applications, such as fitness monitoring, eldercare support, long-term preventive and chronic care, and cognitive assistance. Wearable systems have the advantage of being with the user continuously. So, for example, a fitness application could use real-time activity information to encourage users to perform opportunistic activities. Furthermore, the general public is more likely to accept such activity recognition systems because they are usually easy to turn off or remove.

For systems implementing these applications to be practical, the underlying recognition module must detect a variety of activities that are performed routinely in many different manners by different individuals under different environmental conditions. This presents the challenge of building systems that can handle the real world's noisy data and complexities. Furthermore, deploying the systems imposes some important constraints. The deployment must protect the user's privacy as well as the privacy of those with whom the user comes in contact. The sensors must be lightweight and unobtrusive, and the machine-learning algorithms must be trainable without requiring extensive human supervision. These constraints have made robust recognition systems difficult to engineer.

Over the past four years, we've been building an automatic activity recognition system using on-body sensors. The Mobile Sensing Platform (MSP) tackles several of these design and deployment challenges. Moreover, we've carried out several real-world deployments and user studies, using the results to improve the hardware, software design, and activity recognition algorithms. The lessons learned have broad relevance to context-aware ubiquitous computing applications.

Activity recognition systems

Activity recognition systems typically have three main components:

- a low-level *sensing module* that continuously gathers relevant information about activities using microphones, accelerometers, light sensors, and so on;

Published by the IEEE CS ■ 1536-1268/08/$25.00 © 2008 IEEE

INTEL RESEARCH

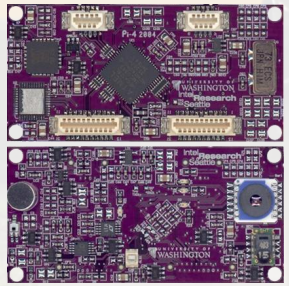
MOBILE SENSING PLATFORM



SMARTPHONE

Real-time Bluetooth communication.

MOBILE SENSING PLATFORM







Pager-size wearable device automatically tracks physical activities throughout day
[Lester, IJAI'05; Pervasive'06]

Photo by Jon E. Froehlich

CHI'08

Consolvo, *et al.*
Activity Sensing in the Wild: A Field Trial of UbiFit Garden



SUNNY CONSOLVO

Activity Sensing in the Wild: A Field Trial of UbiFit Garden

Sunny Consolvo[1, 2], David W. McDonald[2], Tammy Toscos[1], Mike Y. Chen[1], Jon Froehlich[3], Beverly Harrison[1], Predrag Klasnja[1, 2], Anthony LaMarca[1], Louis LeGrand[1], Ryan Libby[3], Ian Smith[1], & James A. Landay[1, 3]

[1]Intel Research Seattle
Seattle, WA 98105 USA
[sunny.consolvo, beverly.harrison, anthony.lamarca, louis.l.legrand]
@intel.com, ttoscos@indiana.edu, mike@ludic-labs.com, iansmith@acm.org

[2] The Information School
DUB Group
University of Washington
Seattle, WA 98195 USA
[consolvo, dwmc, klasnja]
@u.washington.edu

[3] Computer Science & Engineering
DUB Group
University of Washington
Seattle, WA 98195 USA
[landay, jfroehli, libby]
@cs.washington.edu

ABSTRACT
Recent advances in small inexpensive sensors, low-power processing, and activity modeling have enabled applications that use on-body sensing and machine learning to infer people's activities throughout everyday life. To address the growing rate of sedentary lifestyles, we have developed a system, UbiFit Garden, which uses these technologies and a personal, mobile display to encourage physical activity. We conducted a 3-week field trial in which 12 participants used the system and report findings focusing on their experiences with the sensing and activity inference. We discuss key implications for systems that use on-body sensing and activity inference to encourage physical activity.

Author Keywords
persuasive technology, sensing, activity inference, mobile phone, ambient display, fitness, activity-based applications.

ACM Classification Keywords
H.5.2 User Interfaces, H.5.m Miscellaneous.

INTRODUCTION
Recent advances in small inexpensive sensors, low-power processing, and activity modeling have enabled new classes of technologies that use on-body sensing and machine learning to automatically infer people's activities throughout the day. These emerging technologies have seen success with participants in controlled and "living" lab settings [11] and with researchers in situ [18]. The next step is to conduct in situ studies with the target user population. Such studies expose important issues, for example, how the systems are used as part of everyday experiences, where the technology is brittle, and user reactions to activity inference and the presentation of those inferences.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2008, April 5–10, 2008, Florence, Italy.
Copyright 2008 ACM 978-1-60558-011-1/08/04…$5.00.

One application domain for on-body sensing and activity inference is addressing the growing rate of sedentary lifestyles. Regular physical activity is critical to everyone's physical and psychological health, regardless of their being normal weight, overweight, or obese [6,16]. Physical activity reduces risk of premature mortality, coronary heart disease, type II diabetes, colon cancer, and osteoporosis, and has also been shown to improve symptoms associated with mental health conditions such as depression and anxiety. Yet despite the importance of physical activity, many adults in the U.S. do not get enough exercise [1].

Technologies that apply on-body sensing and activity inference to the fitness domain are faced with a challenge regarding which physical activities should be detected. The American College of Sports Medicine (ACSM) recommends that physical activity be regular and include *cardiorespiratory training* (or "cardio") where large muscle groups are involved in dynamic activity such as running or cycling; *resistance training*, that is weight training that builds muscular strength and endurance; and *flexibility training* where muscles are slowly elongated to improve or maintain range of motion [22]. Technologies that attempt to encourage physical activity should support the *range of activities* that contribute to a physically active lifestyle, rather than focus on a single activity such as walking.

Our goal in this work is to investigate users' experiences with a system that we have developed, UbiFit Garden, which uses on-body sensing, activity inference, and a novel personal, mobile display to encourage physical activity. While our future work will focus on how the system affects awareness and sustained behavior change, at this stage, we are exploring how the system affects individuals' everyday lives, how they interpret and reflect on the data about their physical activities, and how they interact with that data. We conducted a three-week field trial (n=12) with participants who were representative of UbiFit Garden's target audience. In this paper, we discuss the types of physical activities participants performed, how those activities were recorded and manipulated, and participants' qualitative reactions to activity inference and manual journaling. We also discuss participants' general reactions to the system.

done by Fish'n'Steps, Houston, and Shakra), providing positive reinforcement rather than punishment (drawing from Houston's success and problems found by Fish'n'Steps), providing frequent opportunities for self-reflection (similar to Houston, Chick Clique, and Shakra), and integrating use into everyday life (all projects).

THE UBIFIT GARDEN SYSTEM

Mistress Mary, quite contrary, how does your garden grow? With silver bells, and cockle shells, and marigolds all in a row.
– Frances Hodgson Burnett

We have designed a healthy lifestyle technology, *UbiFit Garden*, which uses on-body sensing, real-time statistical modeling, and a personal, mobile display to encourage regular physical activity. UbiFit Garden is designed for individuals who have recognized the need to incorporate regular physical activity into their everyday lives but have not yet done so, at least not consistently[1].

The UbiFit Garden system consists of three components: (1) fitness device, (2) interactive application, and (3) glanceable display. The *fitness device* automatically infers and communicates information about several types of physical activities to the glanceable display and interactive application. The *interactive application* includes detailed information about the individual's physical activities and a journal where activities can be added, edited, and deleted. The *glanceable display* uses a non-literal, aesthetic representation of physical activities and goal attainment to motivate behavior (Fig 1). It resides on the background screen, or "wallpaper," of a mobile phone to provide a subtle reminder whenever and wherever the phone is used.

We are using an iterative process to design UbiFit Garden. In addition to drawing from prior work (including our own [5]), we conducted a survey (n=75) with respondents from 13 states across the U.S. that covered a range of attitudes and behaviors with mobile devices and physical activity. This survey tested assumptions about the glanceable display and elicited general feedback. Overall, respondents were very positive about the concept and confirmed that the display was understandable. A majority of respondents could imagine using UbiFit Garden. Common concerns had to do with assuming that all exercise data would have to be manually entered into the phone or that the phone would have to be carried during exercise (the fitness device and interactive application were not addressed in the survey).

UbiFit Garden's Fitness Device
UbiFit Garden is part of a larger research program to explore how sensing and activity inference technologies can be applied to real world problems, such as encouraging people

[1] UbiFit Garden targets the contemplation, preparation, and action stages of change of the *Transtheoretical Model*, which describes the stages through which individuals progress to intentionally modify addictive and other problematic behaviors [19].



Figure 1. UbiFit Garden's glanceable display. a) at the beginning of the week—small butterflies indicate recent goal attainments; the absence of flowers means no activity this week; b) a garden with workout variety; c) the display on a mobile phone—the large butterfly indicates this week's goal was met.

to engage in healthy activities. UbiFit Garden relies on the *Mobile Sensing Platform (MSP)* [4,9], a research platform for mobile sensing and inference applications.

The MSP is a pager-sized, battery powered computer with sensors chosen to facilitate a wide range of mobile sensing applications (Fig 2). The MSP's sensors include: 3-d accelerometer, barometer, humidity, visible and infrared light, temperature, microphone, and compass. It includes a 400MHz XScale microprocessor, 32MB RAM, 2GB of flash memory for storing programs and logging data, and a rechargeable lithium ion battery. The MSP's Bluetooth networking allows it to communicate with other Bluetooth-enabled devices such as mobile phones.

UbiFit Garden uses the MSP to automatically infer physical activities in real time. The MSP runs a set of boosted decision stump classifiers that have been trained to infer walking, running, cycling, using an elliptical trainer, and using a stair machine. These inferences are derived from two sensors: the 3-d accelerometer and barometer. The sensor data is processed and the activity is inferred on the MSP, then the inferences are communicated via Bluetooth to a mobile phone that runs the interactive application and glanceable display (Fig 3). The MSP communicates a list of activities and their predicted likelihoods to the phone four times per second. The phone application aggregates and "smoothes" these fine-grain, noisy data resulting in "human scale"



Figure 2. UbiFit Garden's fitness device–the MSP. At left is the MSP worn by a woman; at center is the MSP in its waist-mount case; and at right are the sensor boards inside the casing.

HEALTH & WELLNESS

UBIFIT





Consolvo *et al.*, CHI'08; Consolvo *et al.*, UbiComp'08

HEALTH & WELLNESS

UBIFIT INFLUENCE



Google Fit Activity Tracking



Apple Watch Activity Tracking

CHI'09

Froehlich, Dillahunt, *et al.*
UbiGreen: Investigating a Mobile Tool for Tracking & Supporting Green Transportation Habits



TAWANNA DILLAHUNT

UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits

Jon Froehlich[1], Tawanna Dillahunt[2], Predrag Klasnja[3,4], Jennifer Mankoff[2], Sunny Consolvo[4], Beverly Harrison[4], James A. Landay[1,4]

[1]CSE; [3]The Information School
DUB Institute, U. of Washington
Seattle, WA 98195 USA
{jfroehli, landay}@cs.washington.edu; klasnja@u.washington.edu

[2]HCI Institute
Carnegie Mellon University
Pittsburgh, PA 15213 USA
{tdillahu, jmankoff}@cs.cmu.edu

[4]Intel Research Seattle
Seattle, WA 98105 USA
{sunny.consolvo, beverly.harrison}@intel.com

ABSTRACT
The greatest contributor of CO_2 emissions in the average American household is personal transportation. Because transportation is inherently a mobile activity, mobile devices are well suited to sense and provide feedback about these activities. In this paper, we explore the use of personal ambient displays on mobile phones to give users feedback about sensed and self-reported transportation behaviors. We first present results from a set of formative studies exploring our respondents' existing transportation routines, willingness to engage in and maintain green transportation behavior, and reactions to early mobile phone "green" application design concepts. We then describe the results of a 3-week field study (N=13) of the UbiGreen Transportation Display prototype, a mobile phone application that semi-automatically senses and reveals information about transportation behavior. Our contributions include a working system for semi-automatically tracking transit activity, a visual design capable of engaging users in the goal of increasing green transportation, and the results of our studies, which have implications for the design of future green applications.

Author Keywords
Sustainability, transportation, ubicomp, ambient displays

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
In 2005, Americans consumed 100 quadrillion British thermal units (BTUs) of energy [32], almost six times the worldwide average per person [20]. This in turn caused the release of 2.2 billion metric tons of carbon dioxide (CO_2), a greenhouse gas assumed to be a major cause of adverse climate change. To reverse this trend, action will be required on many levels, including policy, infrastructure, and individual change. Given the growing prevalence of mobile phones with sensing capabilities, one compelling opportunity to potentially impact human behavior is to offer immediate feedback about how currently sensed behaviors affect the environment. In this paper, we explore the use of personal ambient displays on mobile phones to give users feedback about their sensed and self-reported transportation behaviors (Figure 1).





Figure 1 (left) The UbiGreen Transportation Display shows transit behavior as "wallpaper" on a phone's screen. Here the tree is nearly full of leaves, indicating that the user has completed several green trips for the week. (top) The MSP sensor worn near the waist and the phone's GSM cell tower data are used to semi-automatically infer transportation mode.

Researchers have identified three areas responsible for a majority of energy consumption in American households: home heating and cooling; shopping and eating (and the associated transportation of goods); and commuting, flying and other daily transportation activities [3,35]. In this paper, we focus on the latter (personal transportation), the greatest individual contributor of CO_2 emissions (26%) in the average American household [35].

There is extensive literature in the areas of environmental sociology, public policy, and more recently, conservation psychology that discuss the promotion of environmentally responsible behavior [1,2,26,33]. Past work has shown that motivators such as public commitment, frequent feedback, and personalization can positively impact environmentally responsible behavior [1]. Since the 1990s, information campaigns and other programs have attempted to engage individuals in voluntary greening of transportation behavior

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.
Copyright 2009 ACM 978-1-60558-246-7/09/04...$5.00.

trips (32%). Also similar to our survey results, for 73% of all car trips, greener transportation options existed.

This study also helped reveal the hidden complexities behind the perception and selection of a transportation mode. One participant noted in the exit interview that when biking for transportation, he did not think of it as exercise. Reframing short trips as an opportunity for exercise could potentially make a difference in selecting vehicular travel vs. healthier (and more environmentally friendly) options. If a participant indicated in an ESM survey that s/he drove, we asked if bicycling or walking were viable alternatives. In those cases when bicycling or walking were indeed viable options, our participants reported 52% of the time that they would have been more likely to select bicycling or walking had they thought of health benefits (e.g., caloric expenditure) when making the travel decision.

Finally, the participants were shown an early version of our design concept for a mobile phone application—an icon-based design representing green activity with a growing tree (similar to the tree design in Figure 1). All were able to understand the interface elements without prompting and were positive about having such a representation of their transportation activities on their mobile phones.

Design Implications of Online Survey and ESM Study
Our formative studies suggested that users could benefit from a mobile application that provides awareness of transportation routines and that they would be interested in such an application. Given the range of considerations that impact transportation choice, a design need not focus solely on emphasizing green behavior and may incorporate auxiliary benefits such as cost and health. Other factors such as stress, ability to do other things while traveling (*e.g.,* reading) may also be relevant. Prior work [17] also underscores the many factors that affect transportation choice—not all of which are environmentally related. We highlight these secondary benefits in our design.

Although initially we were interested in building a social-mobile application around green transportation behaviors, our participants' ambivalence about sharing information led us to focus on a single-user application. We plan to explore multi-user applications in future research.

As users did not express a preference between the iconic representations *vs.* numeric representations of transportation behaviors in our online survey data, we decided to use iconic representations for our mobile application. Prior literature enumerates a few of the advantages of iconic visual displays: (1) they may be more aesthetically pleasing in a peripheral viewing situation [31]; (2) once learned, they can easily and quickly convey glanceable information [24]; and (3) they may evoke other responses such as emotional attachment [13,14]. However, iconic representations often do not offer the same level of detail as their numeric counterparts.

UBIGREEN MOBILE APPLICATION
Based on the results of our formative work, we created the UbiGreen Transportation Display, a mobile phone-based application that provides personal awareness about green transportation behaviors through iconic feedback. Small graphical rewards are earned by taking "green" transportation such as riding the bus or train, walking, biking, or carpooling. Although each of these activities has different CO_2 emissions, we counted them equally, as each is preferable to driving alone. Once a green transit activity is sensed, the background (wallpaper) of the user's phone is updated accordingly. A phone's wallpaper represents a critical area of screen real estate as it is seen nearly every time the device is picked up and used. In this way, the wallpaper functions as a type of personal ambient display [9,24].

Our designs are partly based on a finding from social psychology that cognitive representations of different concepts become linked if those concepts are repeatedly encountered together [21]. We take advantage of this fact by jointly presenting a representation of eco-friendly transportation and representations of other goals—such as saving money, getting exercise, etc.—that the user may care about. The interface emphasizes these sub-goals automatically when green transportation is taken.

We were also influenced by research in conservation psychology that showed how caring for animals helps humans connect with nature [25]. Dillahunt *et al.* adapted this notion and explored how "virtual polar bears" could be used to motivate green behaviors [13]. They found that

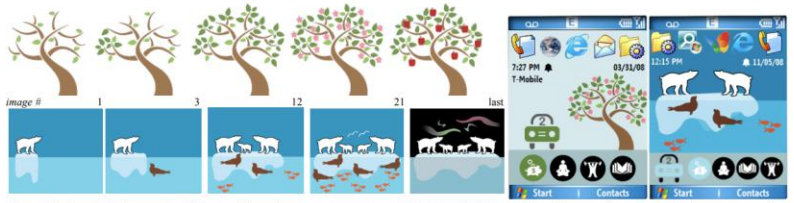


Figure 3: (top-left) A sample of images from the tree progression and **(bottom-left)** a sample of images from the polar bear progression. **(right)** Screenshots showing the graphics in context. In both examples, the user recently carpooled (as indicated by the car with the "2" in the windshield). Since carpooling saves money, the piggy bank is highlighted.

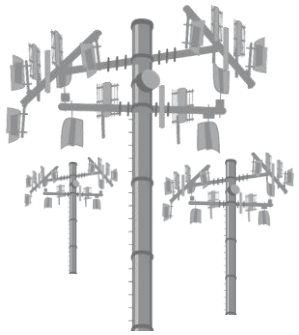
ENVIRONMENTAL SUSTAINABILITY

UBIGREEN TRANSIT SENSING

The MSP
Activity inference



Cell Towers
Transit inference



Manual Entry
Context-triggered mobile surveys





Minimum activity duration: **7 mins**

Froehlich *et al.,* CHI'09

ENVIRONMENTAL SUSTAINABILITY

UBIGREEN



tree
design:

wed thu fri day

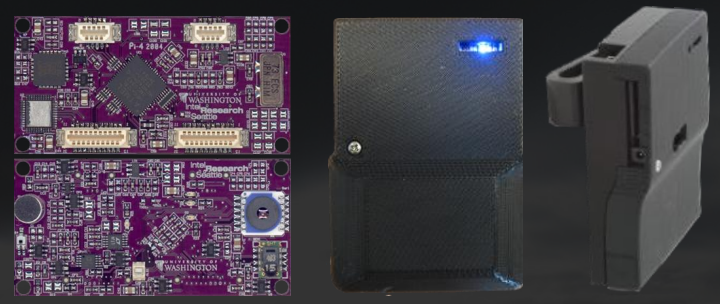


SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**

MSP AS UNIQUE ENABLER





LEVERAGE POINTS

MICROSOFT RESEARCH



ERIC HORVITZ

JOHN KRUMM

WHAT IS THIS DEVICE?



Garmin Geko 201

← GPS Data Logger

Launched Feb, 2003

HOW ABOUT THIS ONE?



← Original iPhone

Launched June 29, 2007
No dedicated GPS chip

STUDYING HUMAN MOVEMENT WAS HARD & EXPENSIVE



Garmin Geko 201



Original iPhone

LEVERAGE POINTS

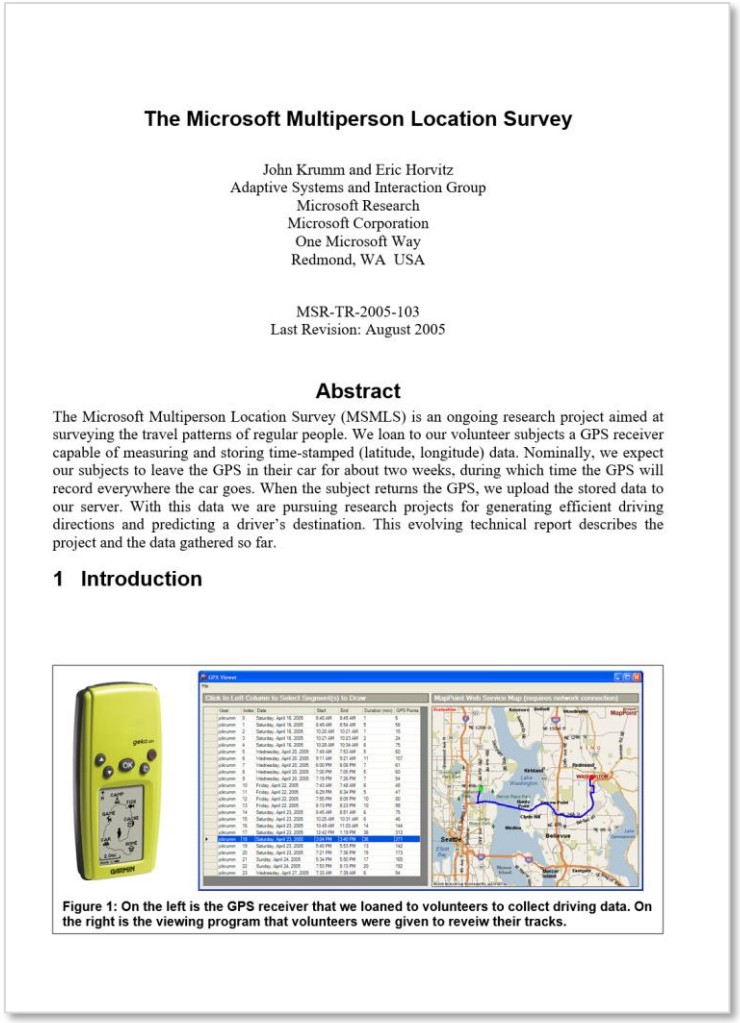
MICROSOFT MULTIPERSON LOCATION SURVEY: 2004-09



JOHN KRUMM



ERIC HORVITZ



The Microsoft Multiperson Location Survey

John Krumm and Eric Horvitz
Adaptive Systems and Interaction Group
Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA USA

MSR-TR-2005-103
Last Revision: August 2005

Abstract

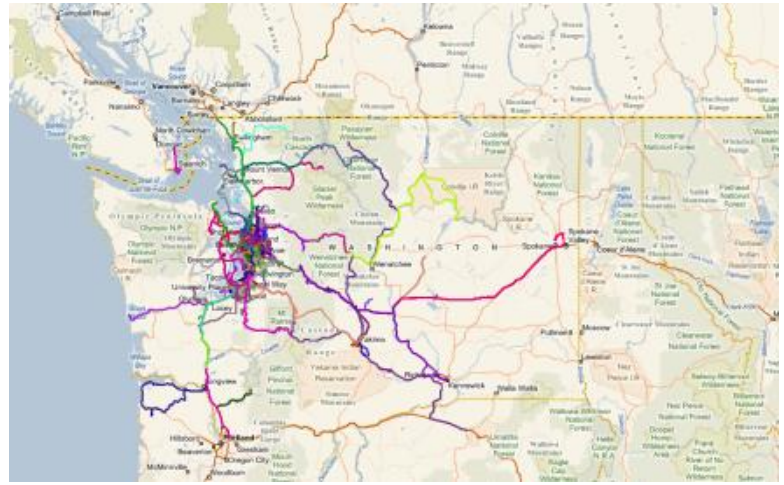
The Microsoft Multiperson Location Survey (MSMLS) is an ongoing research project aimed at surveying the travel patterns of regular people. We loan to our volunteer subjects a GPS receiver capable of measuring and storing time-stamped (latitude, longitude) data. Nominally, we expect our subjects to leave the GPS in their car for about two weeks, during which time the GPS will record everywhere the car goes. When the subject returns the GPS, we upload the stored data to our server. With this data we are pursuing research projects for generating efficient driving directions and predicting a driver's destination. This evolving technical report describes the project and the data gathered so far.

1 Introduction

Figure 1: On the left is the GPS receiver that we loaned to volunteers to collect driving data. On the right is the viewing program that volunteers were given to reveiw their tracks.



Garmin Geko 201



2.5M DATA POINTS

Installed GPS loggers in cars

Collected data from ~2004-2009

~250 participants

2005

2006

2007

2008

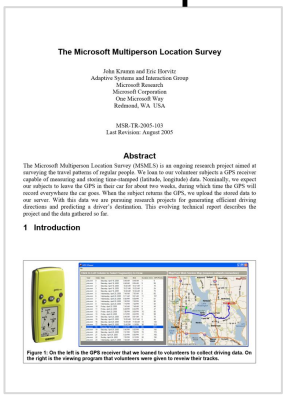
2009

2010

2011

2012

MSR Tech Report'05



The Microsoft Multiperson Location Survey

John Krumm and Eric Horvitz
Adaptive Systems and Interaction Group
Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA USA

MSR-TR-2005-103
Last Revision: August 2005

Abstract

The Microsoft Multiperson Location Survey (MMLS) is an ongoing research project aimed at surveying the travel patterns of regular people. We loan to our volunteer subjects a GPS receiver capable of measuring and storing time-stamped (latitude, longitude) data. Nominally, we expect our subjects to leave the GPS in their car for about two weeks, during which time the GPS will record everywhere the car goes. When the subject returns the GPS, we upload the stored data to our server. With this data we are pursuing research projects for generating efficient driving directions and predicting a driver's destination. This evolving technical report describes the project and the data gathered so far.

1 Introduction

Figure 1: On the left is the GPS receiver that we loaned to volunteers to collect driving data. On the right is the viewing program that volunteers were given to review their tracks.

AAAI'06



Trip Router with Individualized Preferences (TRIP): Incorporating Personalization into Route Planning

Julia Letchner — University of Washington — letchner@cs.washington.edu
John Krumm — Microsoft Research — jckrumm@microsoft.com
Eric Horvitz — Microsoft Research — horvitz@microsoft.com

Abstract

Introduction

SAE'06

Predestination: Inferring Destinations from Partial Trajectories

Real Time Destination Prediction Based On Efficient Routes

Paper Number 2006-01-0811

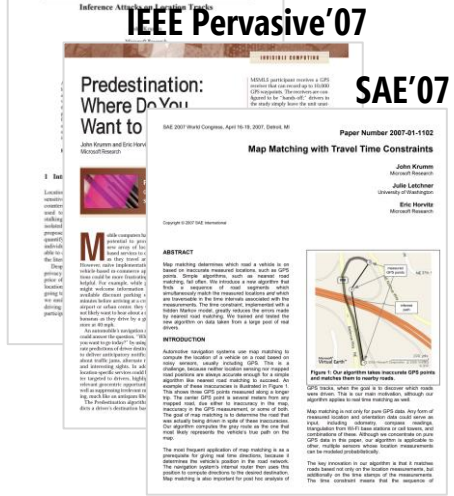
John Krumm
Microsoft Research

UbiComp'06

UbiComp 2006: The Eighth International Conference on Ubiquitous Computing, September 17-21, Orange County, CA, USA

Predestination: Inferring Destinations from Partial Trajectories

SAE'07



Paper Number 2007-01-1102

Map Matching with Travel Time Constraints

John Krumm
Microsoft Research

Julie Letchner
University of Washington

Eric Horvitz
Microsoft Research

ABSTRACT

INTRODUCTION

Pervasive'07

Fifth International Conference on Pervasive Computing (Pervasive 2007), May 13-16, Toronto, Ontario, Canada

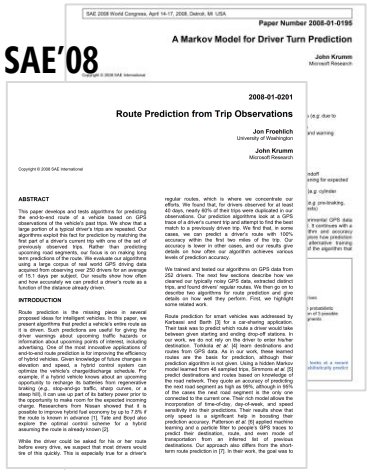
Inference Attacks on Location Tracks

IEEE Pervasive'07

Predestination: Where Do You Want to

John Krumm and Eric Horvitz
Microsoft Research

SAE'08



SAE 2008 World Congress, April 14-17, 2008, Detroit, MI USA

Paper Number 2008-01-0195

A Markov Model for Driver Turn Prediction

John Krumm
Microsoft Research

SAE'08

2008-01-0201

Route Prediction from Trip Observations

Jon Froehlich
University of Washington

Jon Krumm
Microsoft Research

ABSTRACT

INTRODUCTION

Pervasive'09



Realistic Driving Trips For Location Privacy

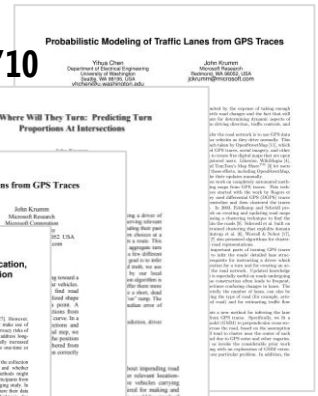
John Krumm

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052 USA
jckrumm@microsoft.com

Abstract

1 Trip Simulations For Privacy

SIGSPATIAL'10



Probabilistic Modeling of Traffic Lanes from GPS Traces

Yihua Chen
Department of Electrical Engineering
University of Washington
Seattle, WA 98195, USA
yhchen@ee.washington.edu

John Krumm
Microsoft Research
Redmond, WA 98052, USA
jckrumm@microsoft.com

PUC'10

Where Will They Turn: Predicting Turn Proportions At Intersections

GIS'10



Detecting Road Intersections from GPS Traces

Mitra Fathi
College of Computing
Georgia Institute of Technology

John Krumm
Microsoft Research
Microsoft Corporation

UbiComp'10

Exploring End User Preferences for Location Obfuscation, Location-Based Services, and the Value of Location

A.J. Berkshire Brush, John Krumm and James Scott
Microsoft Research
{ajbrush, john.krumm, jws}@microsoft.com

ABSTRACT

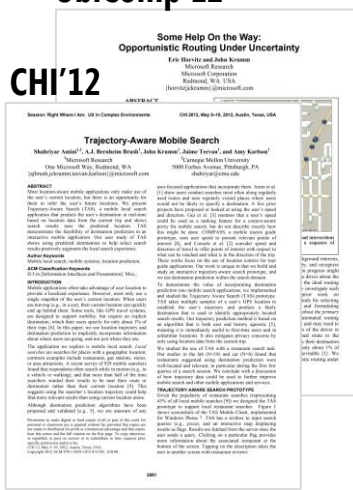
Author Keywords
Privacy, Location, Computational Location Privacy, Obfuscation, Anonymization

ACM Classification Keywords
K.4.1 COMPUTERS AND SOCIETY: Social Issues

General Terms
Human Factors

INTRODUCTION

UbiComp'12



Some Help On the Way: Opportunistic Routing Under Uncertainty

Eric Horvitz and John Krumm
Microsoft Research
Microsoft Corporation
Redmond, WA USA
{horvitz|jckrumm}@microsoft.com

ABSTRACT

CHI'12

Session: Right Where I Am: UX in Complex Environments — CHI 2012, May 5-10, 2012, Austin, Texas, USA

Trajectory-Aware Mobile Search

Shahriyar Amini[1,2], A.J. Bernheim Brush[1], John Krumm[1], Jaime Teevan[1], and Amy Karlson[1]
[1]Microsoft Research — [2]Carnegie Mellon University
One Microsoft Way, Redmond, WA — 5000 Forbes Avenue, Pittsburgh, PA
{ajbrush,jckrumm,teevan,karlson}@microsoft.com — shahriyar@cmu.edu

ABSTRACT

INTRODUCTION

CST'11



Chapter 7
Trajectory Analysis for Driving

John Krumm

Abstract

7.1 Introduction

SAE'08

Froehlich & Krumm

Route Prediction from Trip Observations



JOHN KRUMM

2008-01-0201

Route Prediction from Trip Observations

Jon Froehlich
University of Washington

John Krumm
Microsoft Research

Copyright © 2008 SAE International

ABSTRACT

This paper develops and tests algorithms for predicting the end-to-end route of a vehicle based on GPS observations of the vehicle's past trips. We show that a large portion of a typical driver's trips are repeated. Our algorithms exploit this fact for prediction by matching the first part of a driver's current trip with one of the set of previously observed trips. Rather than predicting upcoming road segments, our focus is on making long term predictions of the route. We evaluate our algorithms using a large corpus of real world GPS driving data acquired from observing over 250 drivers for an average of 15.1 days per subject. Our results show how often and how accurately we can predict a driver's route as a function of the distance already driven.

INTRODUCTION

Route prediction is the missing piece in several proposed ideas for intelligent vehicles. In this paper, we present algorithms that predict a vehicle's entire route as it is driven. Such predictions are useful for giving the driver warnings about upcoming traffic hazards or information about upcoming points of interest, including advertising. One of the most innovative applications of end-to-end route prediction is for improving the efficiency of hybrid vehicles. Given knowledge of future changes in elevation and speed, a hybrid control system can optimize the vehicle's charge/discharge schedule. For example, if a hybrid vehicle knows about an upcoming opportunity to recharge its batteries from regenerative braking (*e.g.*, stop-and-go traffic, sharp curves, or a steep hill), it can use up part of its battery power prior to the opportunity to make room for the expected incoming charge. Researchers from Nissan showed that it is possible to improve hybrid fuel economy by up to 7.8% if the route is known in advance [1]. Tate and Boyd also explore the optimal control scheme for a hybrid assuming the route is already known [2].

While the driver could be asked for his or her route before every drive, we suspect that most drivers would tire of this quickly. This is especially true for a driver's regular routes, which is where we concentrate our efforts. We found that, for drivers observed for at least 40 days, nearly 60% of their trips were duplicated in our observations. Our prediction algorithms look at a GPS trace of a driver's current trip and attempt to find the best match to a previously driven trip. We find that, in some cases, we can predict a driver's route with 100% accuracy within the first two miles of the trip. Our accuracy is lower in other cases, and our results give details on how often our algorithm achieves various levels of prediction accuracy.

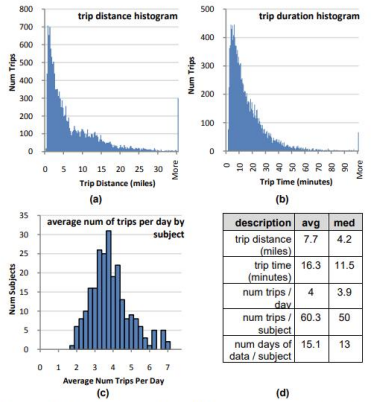
We trained and tested our algorithms on GPS data from 252 drivers. The next few sections describe how we cleaned our typically noisy GPS data, extracted distinct trips, and found drivers' regular routes. We then go on to describe two algorithms for route prediction and give details on how well they perform. First, we highlight some related work.

Route prediction for smart vehicles was addressed by Karbassi and Barth [3] for a car-sharing application. Their task was to predict which route a driver would take between given starting and ending drop-off stations. In our work, we do not rely on the driver to enter his/her destination. Torkkola *et al.* [4] learn destinations and routes from GPS data. As in our work, these learned routes are the basis for prediction, although their prediction algorithm is not given. Using a hidden Markov model learned from 46 sampled trips, Simmons *et al.* [5] predict destinations and routes based on knowledge of the road network. They quote an accuracy of predicting the next road segment as high as 99%, although in 95% of the cases the next road segment is the only one connected to the current one. Their rich model allows the incorporation of time-of-day, day-of-week, and speed sensitivity into their predictions. Their results show that only speed is a significant help in boosting their prediction accuracy. Patterson *et al.* [6] applied machine learning and a particle filter to people's GPS traces to predict their destination, route, and even mode of transportation from an inferred list of previous destinations. Our approach also differs from the short-term route prediction in [7]. In their work, the goal was to

After filtering, our corpus of trips appeared much more reasonable, although some invalid trips remained, and some valid trips were removed. Acknowledging that our filtering process was not perfect, we tried to err on the side of removing too many valid trips with the goal of minimizing invalid trips that would distort our prediction accuracy experiments.

THE TRIP DATA

Trip filtering reduced the number of trips from 27,479 to 14,523 (a 47.2% reduction). For our analysis, we also required that subjects have more than ten total trips; this excluded twelve subjects. We felt that ten or fewer trips were not enough to properly establish a driver's regular patterns. Thus, our final dataset included 14,468 trips from 240 subjects. On average, our subjects took 4 trips per day. The average trip length was 7.7 miles and trip time 16.3 minutes. For comparison, according to the 2001 National Household Travel Survey, an average driver in a metropolitan area the size of Seattle makes 4.07 trips per day with an average length of 9.87 miles [10]. Figure 6 reveals some high level statistics and distributions of our trip data and Figure 7 displays the trip data overlaid on two maps, one of Washington State and the other of Seattle.



description	avg	med
trip distance (miles)	7.7	4.2
trip time (minutes)	16.3	11.5
num trips / day	4	3.9
num trips / subject	60.3	50
num days of data / subject	15.1	13

Figure 6. A histogram of trip distances is shown in (a) and trip durations in (b). Most subjects drove between 3-5 trips per day on average as can be observed in (c). Finally, (d) shows a small table of trip statistics.

FROM TRIPS TO REGULAR ROUTES

When making a route prediction, we try to pick which previously seen route a driver is on. We make a distinction between trips, which we extract from the GPS data, and routes, which we build from repeated trips. A trip describes a driver's path through time and space using time stamped GPS data. A route is simply an abstraction of a trip (or trips) without the temporal component. That is, a route is a collection of latitude, longitude pairs that define a directed path. A *regular* route is a path that a driver drives often. This section describes how we extract routes from trip data.

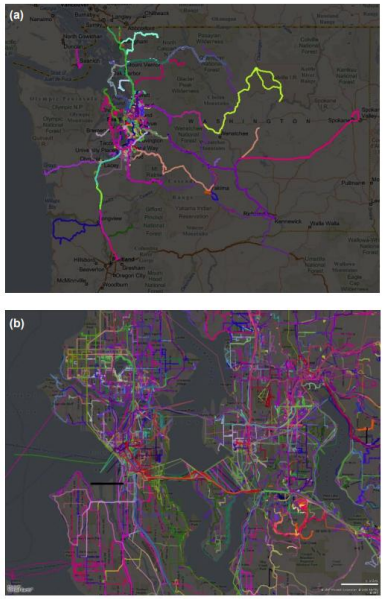
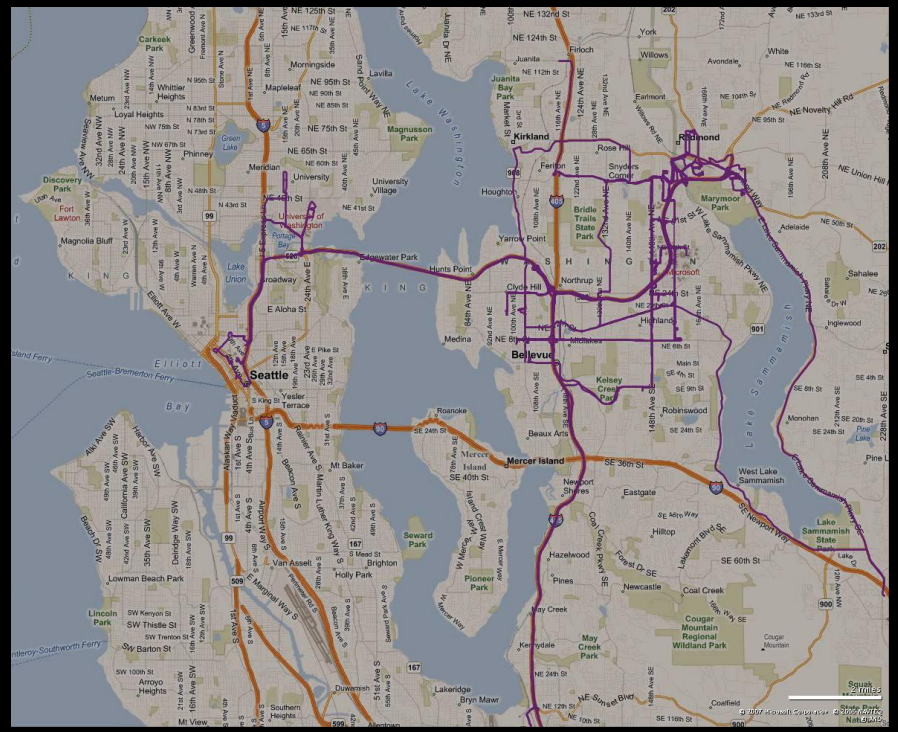


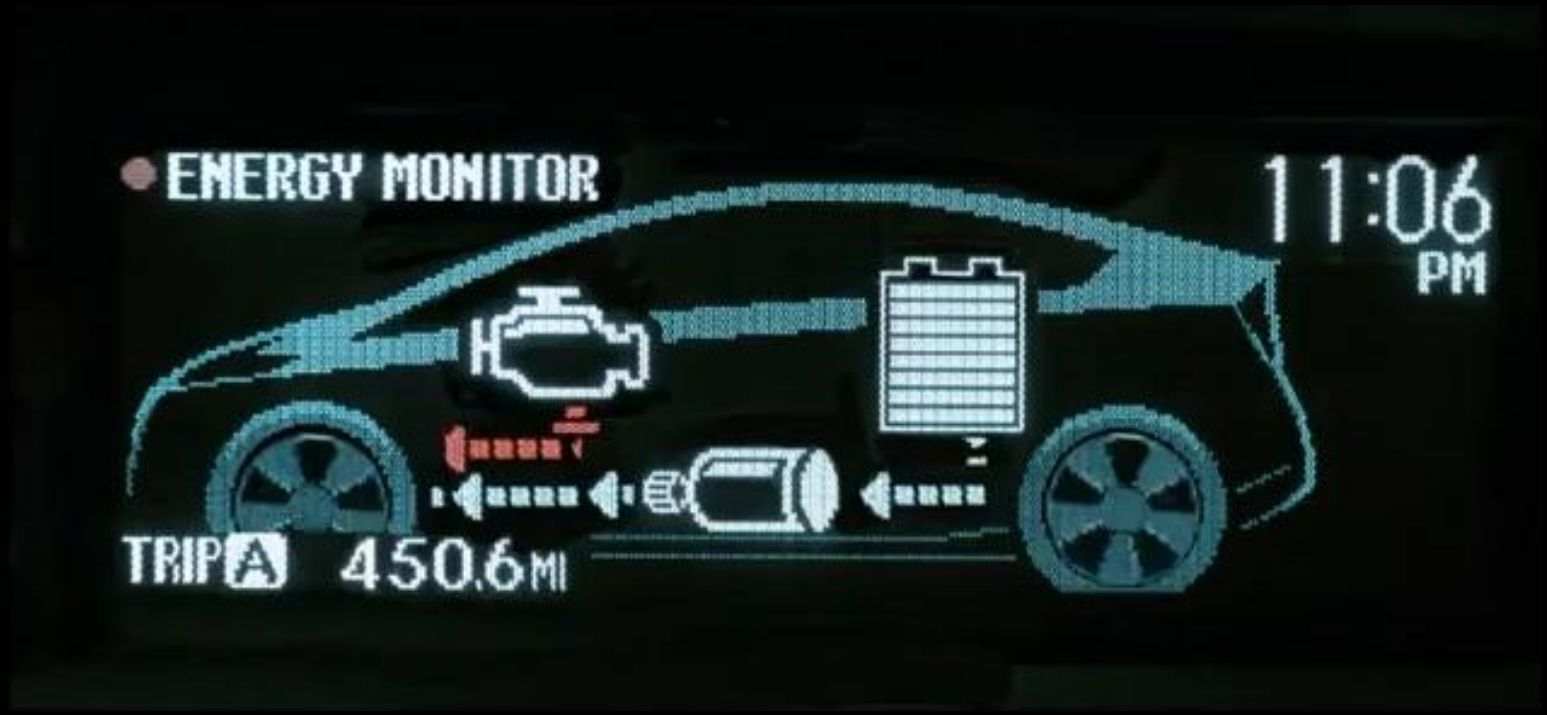
Figure 7. A map of Washington (a) and the Seattle metropolitan area (b) with overlaid trip data from the MSMLS dataset. The trips are colored by subject. Figure (b) is the same map view as Figure 3 but with cleaned and filtered trips. Note that the maps have been darkened to emphasize the trip data.

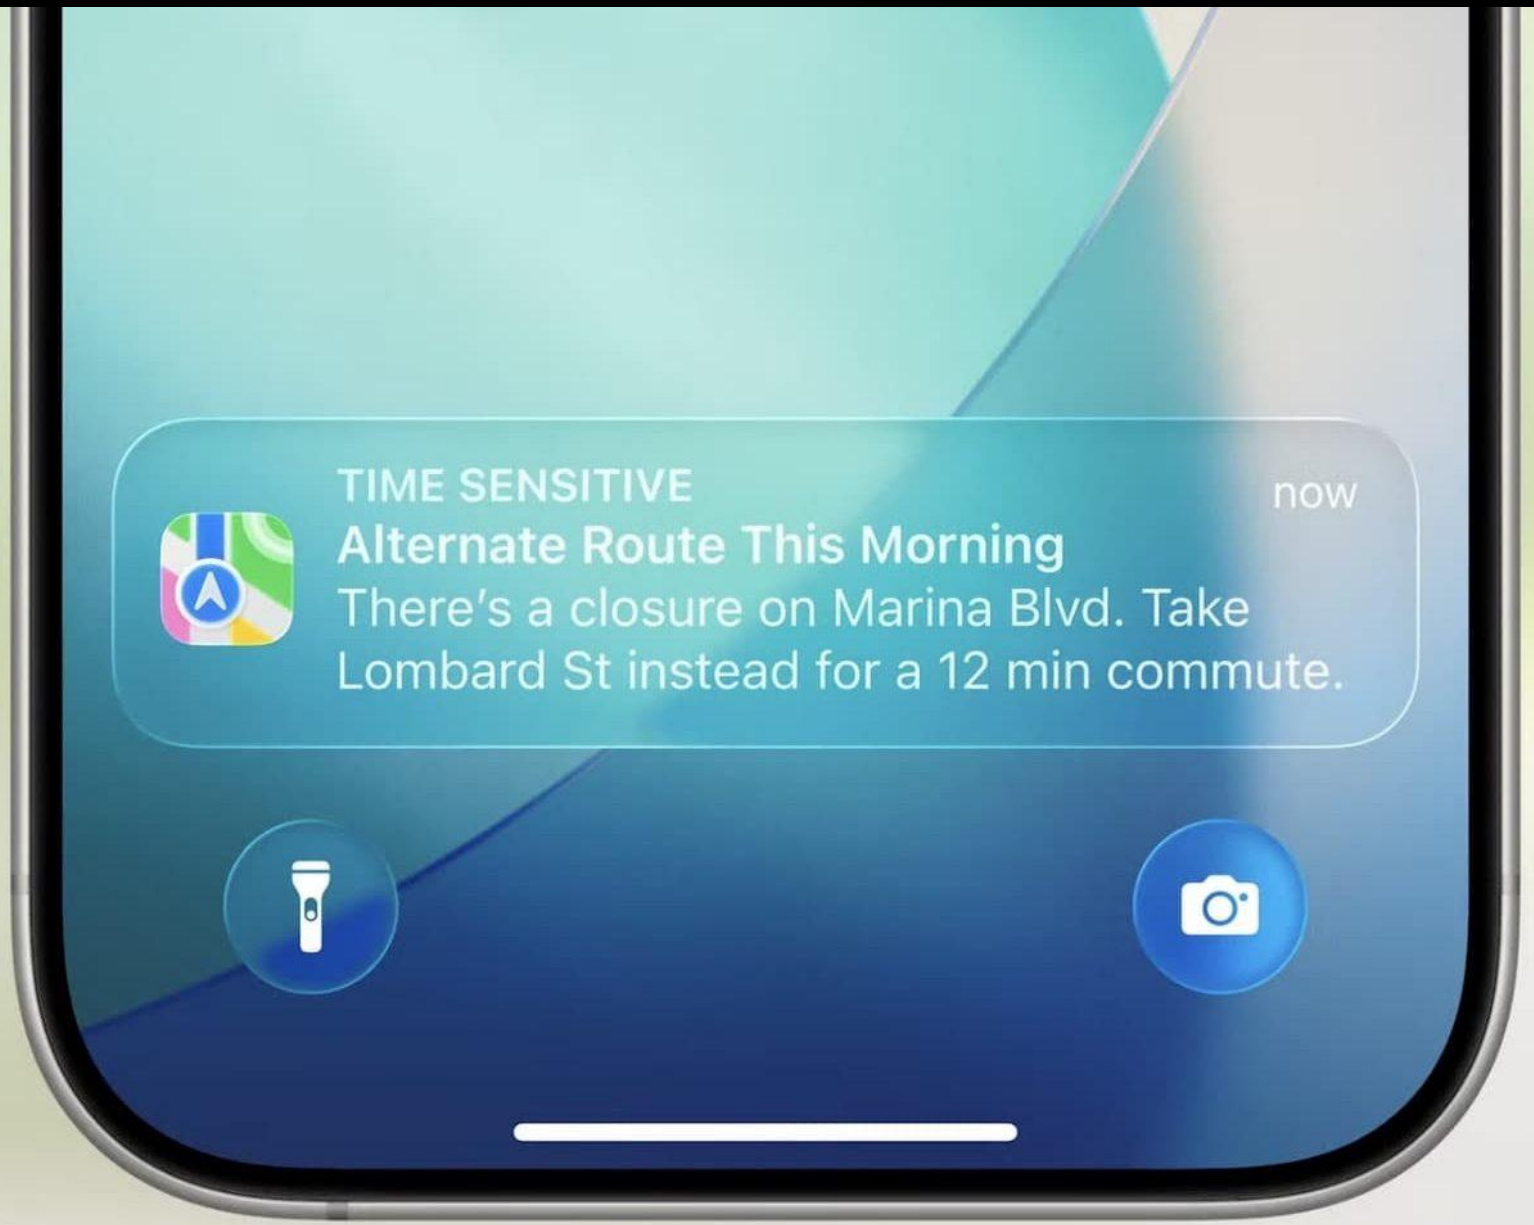
Trip Similarity

Intuitively, trips that a driver takes along the same path should be considered the same route. To do this programmatically, we must formally define some notion of trip similarity. If two or more trips are similar, then we combine them into a route. This section shows how we assess the similarity of two trips. The algorithm relies solely on the (latitude, longitude) data from our filtered trips. It does not require an extra step of inferring which roads were traversed.

OPTIMIZE BATTERY USAGE BASED ON ROUTE PREDICTION






TIME SENSITIVE
now
Alternate Route This Morning
There's a closure on Marina Blvd. Take Lombard St instead for a 12 min commute.

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**

GPS DATASET AS UNIQUE ENABLER



2009

I'm **feeling good**. My research has **momentum**.


HCIC 2009
Snow Mountain Ranch
Fraser, Colorado
Photo by Jon E. Froehlich


HCIC'09
BANQUET DINNER


“But that’s **not** HCI…”
-HCI LUMINARY

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your work

Sensing and Feedback of Everyday Activities to Promote Environmental Behaviors

Jon E. Froehlich

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2011

Program Authorized to Offer Degree:
Department of Computer Science and Engineering

Sensing and Feedback to Promote Environmentally Sustainable Behaviors

@jonfroehlich



dub design: use: build:

UNIVERSITY of WASHINGTON

CHI'10

Froehlich, Findlater, Landay

The Design of Eco-Feedback Technology



The Design of Eco-Feedback Technology

Jon Froehlich[1], Leah Findlater[2], James Landay[1]
[1]Computer Science and Engineering, [2]The Information School
DUB Institute, University of Washington
{ jfroehli, leahkf, landay}@uw.edu

ABSTRACT
Eco-feedback technology provides feedback on individual or group behaviors with a goal of reducing environmental impact. The history of eco-feedback extends back more than 40 years to the origins of environmental psychology. Despite its stated purpose, few HCI eco-feedback studies have attempted to measure behavior change. This leads to two overarching questions: (1) what can HCI learn from environmental psychology and (2) what role should HCI have in designing and evaluating eco-feedback technology? To help answer these questions, this paper conducts a comparative survey of eco-feedback technology, including 89 papers from environmental psychology and 44 papers from the HCI and UbiComp literature. We also provide an overview of predominant models of proenvironmental behaviors and a summary of key motivation techniques to promote this behavior.

Author Keywords
Eco-feedback, Environmental HCI, Reflective HCI, Survey

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI)

General Terms
Design, Human Factors

INTRODUCTION
As environmental issues such as climate change, air pollution, and water scarcity become more salient in the global consciousness, so too have they become more active targets of research within HCI and Ubiquitous Computing [6, 19, 57]. One particularly popular form of environmental HCI research is the design and study of *eco-feedback technology*, which we define as technology that provides feedback on individual or group behaviors with a goal of reducing environmental impact (adapted from [39] and [28], see Figure 1 for examples). Despite this goal, few HCI eco-feedback studies have even attempted to measure behavior change. Although eco-feedback may be seen as an extension of research in *persuasive technology* [17], it actually extends back much further to over 40 years of research in environmental psychology. This leads to two

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
Copyright 2010 ACM 978-1-60558-929-9/10/04....$10.00.



Figure 1. Examples of eco-feedback technology. (left-to-right) The Infotropism display uses sensors and living plants to provide feedback about recycling and waste disposal [29]. WaterBot provides ambient feedback information about water usage [3]. The UbiGreen Transportation Display semi-automatically senses and feeds back information about transportation to encourage green transit [19].

interrelated questions: (1) What can HCI learn from environmental psychology and (2) what should be the role of the HCI community in contributing to eco-feedback research? To explore these questions in detail, we present a review of the related environmental psychology literature as well as a comparative survey of eco-feedback studies in both HCI and environmental psychology.

Eco-feedback technology is based on the working hypothesis that most people lack awareness and understanding about how their everyday behaviors such as driving to work or showering affect the environment; technology may bridge this "environmental literacy gap" by automatically sensing these activities and feeding related information back through computerized means (e.g., mobile phones, ambient displays, or online visualizations). HCI and UbiComp researchers have built eco-feedback technologies for a variety of domains including energy consumption [28], water usage [3], transportation [19], and waste disposal practices [29].

Contributing to this growing interest in eco-feedback technology is the parallel advancement and availability of sensing systems for environmentally related activities (e.g., human activity inference [35]) *and* interactive displays to feedback this data (e.g., iPods and mobile phones). Such advances provide a rich space of opportunities for new types of eco-feedback that could not be considered in the past. Moreover, the next generation of resource measurement systems (often referred to as "smart meters") will soon provide real-time (or near real-time) data on electricity, gas, and water usage in homes and businesses. This will produce tremendous amounts of data that can be

PC, with some amount of historical data available for self-comparison. Almost all (10/12) of the devices used were semi-interactive, but interactions were often limited, for example, to pressing a button that would cycle through statistics like the current day's electricity rate or the amount of the last month's bill (e.g., [48]).

In contrast, the eco-feedback designs in the HCI papers were much more diverse and fully explained. Of the 27 HCI papers that provide some sort of study of their eco-feedback technology, only four papers do not disclose a screenshot of the interface. In addition, the studies employed a range of presentation mediums for their feedback including: ambient displays (e.g., [3, 25, 42, 44]), mobile phone applications (e.g., [19, 45]), desktop games (e.g., [4]), and social websites [38]—see Figure 3.

Unfortunately, many of the eco-feedback designs in HCI do not link back to work in environmental or behavioral psychology. In our survey, less than half of the HCI eco-feedback papers referenced behavioral psychology literature and 58% referenced environmental psychology literature. Even more dramatically, no study in environmental psychology referred back to HCI. This represents a profound gap between disciplines. Interestingly, one author McCalley (e.g., [40]) has published in both fields—having published in both the Persuasive conference and journals in psychology and energy. Perhaps a future goal for HCI should be to initiate collaborations with environmental psychologists.

Discussion of Treatment
The primary motivation of eco-feedback technologies in both disciplines is to promote proenvironmental behaviors. Despite the relatively simple interfaces and lack of focus on design, the environmental psychology studies have achieved impressive results, a finding which should be cause for reflection by eco-feedback researchers in HCI. HCI researchers/practioners should ground their designs in the basic principles uncovered by environmental psychology. They can then apply the unique methodologies and approaches found in HCI (e.g., user centered design) to further the design of eco-feedback technology.

Although the environmental psychology studies show that eco-feedback can reduce consumption, they do not clarify the extent to which this impact is based on specific design elements. Considering only the designs that appeared in the environmental psychology studies, we can see questions that HCI researchers are well-suited to study: How

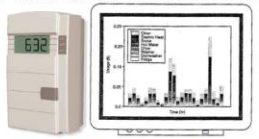


Figure 2. Two of the most commonly reported eco-feedback designs in environmental psychology: (left) a simple LCD display [32]; (right) a bar chart showing current and historical consumption data [13].

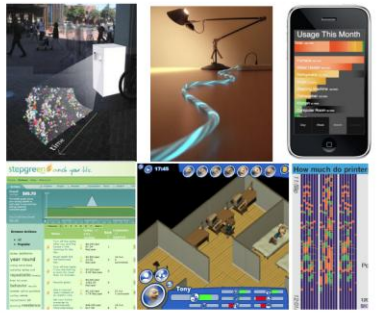


Figure 3. (clockwise): JetSam: ambient display for trash [42]; Power-Aware-Cord: ambient energy usage display [25]; WattBot: mobile phone home energy feedback [45]; Stepgreen: proenvironmental behavior tracking social website [38]; ThePowerHouse: resource eco-feedback in a virtual game environment [4]; Imprint: environmental impact of printing ambient public display [44].

important is it that eco-feedback be even minimally interactive? What types of information and presentation mediums are most effective (e.g., graphs vs. abstract ambient representations)? To what degree does the physical placement and access to the device impact its overall effectiveness? Answering these questions should allow us to identify how environmental psychologists may improve on the advances they have already made.

Consumption Targets of Eco-Feedback Technology
Eco-feedback technologies have been developed to target many types of consumption. The most common target is residential electricity usage: 41% of the papers in HCI and 92% of the papers in environmental psychology. This emphasis is both a reflection of the impact that electricity usage behaviors have on the environment as well as the ease with which energy usage can be automatically sensed.

As a field partially composed of computer scientists and designers, HCI researchers often have the resources to construct both their own novel sensing systems as well as their own feedback interfaces. HCI has also developed techniques to test and iterate on prototypes for exploring interactions and interfaces independent of the current state of technology (e.g., Wizard-of-Oz evaluations). As such, HCI has explored eco-feedback technologies for a larger set of behaviors than have been studied in environmental psychology. We found 20 HCI papers on eco-feedback for electricity, 4 on water, 4 on transportation, 4 on carbon tracking, 3 on garbage and recycling behaviors, 3 on the environmental impact of product purchases, 2 on paper usage and 1 on eco-feedback for a virtual game world.

Discussion of Consumption Targets
The HCI studies have shown that people are open to new types of eco-feedback for behaviors outside of energy usage. One role for HCI is then to challenge the limitations

Sensing and Feedback of Everyday Activities to Promote Environmental Behaviors

Jon E. Froehlich

A dissertation submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

University of Washington

2011

Program Authorized to Offer Degree:
Department of Computer Science and Engineering

UW Outstanding Dissertation Award

HCIC'09

Promoting Energy Efficient Behaviors in the Home through Feedback: The Role of Human-Computer Interaction

Jon Froehlich
DUB Group
CSE, University of Washington, Seattle, WA 98195
jfroehli@cs.washington.edu



Figure 1. The first two examples illustrate highly localized feedback. The other two examples provide aggregate energy usage information that can be placed anywhere in the home. (from left-to-right). The Kill-A-Watt provides energy usage information for appliances plugged into the proxy outlet. The MIT HeatSink (Arroyo, 2005) illuminates water according to temperature directly at the point of consumption. The Energy Detective (TED) presents overall electricity consumption information numerically via a monochrome display. The EnergyMate, a conceptual design, is an LCD flat panel display meant to be positioned in a visible place in the home and offers real time feedback about energy usage.

Abstract

The consumption of energy is unlike most consumable goods. It is abstract, invisible, and untouchable. Without a tangible manifestation, home energy usage often goes unnoticed. Advances in resource monitoring systems will soon provide real-time data on electricity, gas, and water usage in the home. This will produce a tremendous amount of data that can be analyzed and fed back to the user—creating a rich space of opportunities for HCI research. This paper outlines common misconceptions of energy usage in the home, establishes the potential of feedback to change energy consumption behavior, and introduces ten design dimensions of feedback technology with which to build and evaluate future systems.

Introduction

The United States consumes one quarter of the world's energy resources, despite accounting for less than five percent of the world's population (US Department of Energy, 2002). The residential sector accounts for 21% of the nation's energy use and the average American household spends nearly $2,000 on energy bills per year (US Department of Energy, 2006). Home energy and personal transport are the top two contributors of the average American's CO_2 emissions into the environment (Weber and Matthews, 2007), accounting for over 50% of their total carbon footprint. To date, the primary methods applied to improving energy efficiency and/or reducing energy usage have been technological and economic (Armel, 2008). For example, the production of hybrid or hydrogen vehicles has been emphasized as a major solution to CO_2 reduction and oil dependence. However, there is growing evidence that a human-centered, *behavioral* approach should also be pursued to educate, inform, and motivate energy efficient human behaviors.

In a study evaluating the energy consumption of 10 identical Habitat for Humanity all electric homes outfitted with the same appliances and equipment, homes were found to exhibit a large range in energy consumption, with the most energy intensive home consuming 2.6 times more energy than the least (Parker et al., 2008). Indeed, it has

CHI'09 Workshop

This paper was presented at the CHI09 workshop: "Defining the Role of HCI in the Challenges of Sustainability."
It is also published as a UW CSE Technical Report, CSE-08-10-01

Sensing Opportunities for Personalized Feedback Technology to Reduce Consumption

Jon Froehlich, Kate Everitt, James Fogarty, Shwetak Patel, James Landay
DUB Institute
Department of Computer Science and Engineering
University of Washington, Seattle, WA 98195, USA
{jfroehli, everitt, jfogarty, shwetak, landay}@cs.washington.edu

ABSTRACT
Most people are unaware of how their daily activities affect the environment. Previous studies have shown that feedback technology is one of the most effective strategies in reducing electricity usage in the home. In this position paper, we expand the notion of feedback systems to a broad range of human behaviors that have an impact on the environment. In particular, we enumerate five areas of consumption: electricity, water, personal transportation, product purchases, and garbage disposal. For each, we outline their effect on the environment and review and propose methods for automatically sensing them to enable new types of feedback systems.

Author Keywords
Sensing, Feedback Technology, Sustainability, UbiComp

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Everyday human behaviors (e.g., home energy use and personal travel) are directly responsible for 28% of US energy consumption and 41% of US CO_2 emissions (Shui and Dowlatabadi, 2005). Most people, however, are unaware of how their daily activities affect the environment nor how often they engage in those activities. Feedback has been shown to be one of the most effective strategies in reducing electricity usage in the home (Geller, 1982). Because of limitations in sensing, however, there has been a dearth of feedback technologies in other domains, such as personal transportation and product usage. With the advent of low-cost sensing technologies, fast computation, and advances in machine learning, we now have the potential to provide personal, relevant feedback in real time for a variety of consumption activities.

In this paper, we highlight the state-of-the-art in sensing electricity, water, personal transportation, product purchases, and garbage disposal. In particular, we focus on resource consumption that is within an individual's primary locus of control. For each of the five consumption areas, we briefly outline their impact on the environment and review methods for automatically sensing them. For some areas, like electricity, a large number of sensing approaches already exist to measure energy consumption (driven in large part by utility companies' need to charge based on utilization as well as the ease in which this data can be sensed)—for others, like transportation or product purchases, very few sensing methods exist. In some cases, we also propose our own methods or suggest ways in which existing sensing systems could be reappropriated for the purposes of feedback technology to reduce consumption.

The goal here is twofold: (1) to inform the reader about the ways in which environmentally impactful human behaviors can be sensed; (2) to inspire thinking about ways in which these new types of sensor data may be aggregated, analyzed, and fed back to the individual in order to increase awareness about environmentally impactful activities and motivate sustainable behaviors. The scope of this paper prevents us from discussing particular feedback designs or visualizations (for a review see Pierce et al., 2008 and Froehlich, 2009) and instead focuses primarily on the sensing aspects.

Some open questions worth considering while reading through the rest of the paper:

- How can we automatically sense and measure the ways in which human activities impact the environment? Is this something that even should be done?
- How do we present this information to the user? Will the effective feedback strategies used to reduce home energy consumption (Fischer, 2008) translate to other environmentally impactful activities such as transit and trash disposal habits?
- Certain activities, such as plane rides, dwarf the impact of other activities. What are the implications of such contrasts for motivating sustainable behaviors?
- Assuming environmentally impactful activities can be sensed, how can we reduce the cognitive burden of consuming all of this information, particularly given that "the environment" is rarely a first-order concern of individuals when making decisions?
- How do we reconcile the paradox of creating new sensing/feedback technologies to reduce energy consumption when the manufacturing, distribution and powering of such technologies has an environmental impact itself?

CHI'09

UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits

Jon Froehlich[1], Tawanna Dillahunt[2], Predrag Klasnja[3,4], Jennifer Mankoff[2], Sunny Consolvo[4], Beverly Harrison[4], James A. Landay[1,4]
[1]CSE, [2]The Information School
DUB Institute, U. of Washington
Seattle, WA 98195 USA
{jfroehli, landay}@cs.washington.edu; klasnja@u.washington.edu
[3]HCI Institute
Carnegie Mellon University
Pittsburgh, PA 15213 USA
{dillahu, jmankoff}@cs.cmu.edu
[4]Intel Research Seattle
Seattle, WA 98105 USA
{sunny.consolvo, beverly.harrison}@intel.com

ABSTRACT
The greatest contributor of CO_2 emissions in the average American household is personal transportation. Because transportation is inherently a mobile activity, mobile devices are well suited to sense and provide feedback about these activities. In this paper, we explore the use of personal ambient displays on mobile phones to give users feedback about sensed and self-reported transportation behaviors. We first present results from a set of formative studies exploring our respondents' existing transportation routines, willingness to engage in and maintain green transportation behavior, and reactions to early mobile phone "green" application design concepts. We then describe the results of a 3-week field study (N=13) of the UbiGreen Transportation Display prototype, a mobile phone application that semi-automatically senses and reveals information about transportation behavior. Our contributions include a working system for semi-automatically tracking transit activity, a visual design capable of engaging users in the goal of increasing green transportation, and the results of our studies, which have implications for the design of future green applications.

Author Keywords
Sustainability, transportation, ubicomp, ambient displays

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
In 2005, Americans consumed 100 quadrillion British thermal units (BTUs) of energy [32], almost six times the worldwide average per person [20]. This in turn caused the release of 2.2 billion metric tons of carbon dioxide (CO_2), a greenhouse gas assumed to be a major cause of adverse climate change. To reverse this trend, action will be required on many levels, including policy, infrastructure,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2009, April 4-9, 2009, Boston, Massachusetts, USA.
Copyright 2009 ACM 978-1-60558-246-7/09/04...$5.00.





Figure 1 (left) The UbiGreen Transportation Display shows transit behavior as "wallpaper" on a phone's screen. Here the tree is nearly full of leaves, indicating that the user has completed several green trips for the week. (top) The MSP sensor worn near the waist and the phone's GSM cell tower data are used to semi-automatically infer transportation mode.

and individual change. Given the growing prevalence of mobile phones with sensing capabilities, one compelling opportunity to potentially impact human behavior is to offer immediate feedback about how currently sensed behaviors affect the environment. In this paper, we explore the use of personal ambient displays on mobile phones to give users feedback about their sensed and self-reported transportation behaviors (Figure 1).

Researchers have identified three areas responsible for a majority of energy consumption in American households: home heating and cooling; shopping and eating (and the associated transportation of goods); and commuting, flying and other daily transportation activities [3,35]. In this paper, we focus on the latter (personal transportation), the greatest individual contributor of CO_2 emissions (26%) in the average American household [35].

There is extensive literature in the areas of environmental sociology, public policy, and more recently, conservation psychology that discuss the promotion of environmentally responsible behavior [1,2,26,33]. Past work has shown that motivators such as public commitment, frequent feedback, and personalization can positively impact environmentally responsible behavior [1]. Since the 1990s, information campaigns and other programs have attempted to engage individuals in voluntary greening of transportation behavior

IJCAI'09

Sensing and Predicting the Pulse of the City through Shared Bicycling

Jon Froehlich[1], Joachim Neumann[2], Nuria Oliver[2]
[1]Computer Science and Engineering
University of Washington
jfroehli@cs.washington.edu
[2]Telefonica Research, Barcelona, Spain
{joachim, nuriao}@tid.es

Abstract
City-wide urban infrastructures are increasingly reliant on network technology to improve and expand their services. As a side effect of this digitalization, large amounts of data can be sensed and analyzed to uncover patterns of human behavior. In this paper, we focus on the *digital footprints* from one type of emerging urban infrastructure: shared bicycling systems. We provide a spatiotemporal analysis of 13 weeks of bicycle station usage from Barcelona's shared bicycling system, called Bicing. We apply clustering techniques to identify shared behaviors across stations and skew how these behaviors relate to location, neighborhood, and time of day. We then compare experimental results from four predictive models of near-term station usage. Finally, we analyze the impact of factors such as time of day and station activity in the prediction capabilities of the algorithms.



Figure 1. (top) A Bicing station; a close-up of a bicycle and parking slot; and a user at a station kiosk using RFID to check-out a bicycle; (bottom) The 390 Bicing stations distributed across the city of Barcelona, Spain

1 Introduction

Observing and modeling human movement in urban environments is central to traffic forecasting, understanding the spread of biological viruses, designing location-based services, and improving urban infrastructure. However, little has changed since Whyte (1980) observed in his "Street Life Project" that the *actual* usage of New York's streets and squares clashed with the original ideas of architects and city planners. A key difficulty faced by urban planners, virologists, and social scientists is that obtaining large, real-world observational data of human movement is challenging and costly (Brockman et al., 2006).

As websites have evolved to offer geo-located services, new sources of real-world behavioral data have begun to emerge. For example, Battiestuny *et al.* (2007) and Girardin *et al.* (2008) used geo-tagging patterns of photographs in Flickr to automatically detect interesting real-world events and draw conclusions about the flow of tourists in a city. In addition, as city-wide urban infrastructures such as buses, subways, public utilities, and roads become digitized, other sources of real-world datasets that can be implicitly sensed are emerging. Ratti *et al.* (2006) and González *et al.* (2008) used cellular network data to study city dynamics and human mobility. McNamara *et al.* (2008) used data collected from an RFID-enabled subway system to predict co-location patterns amongst the transit users. Such sources of data are ever-expanding and offer large, under-explored datasets of physically-based interactions with the real world.

In this paper, we introduce a novel source of real-world human behavioral data from a new type of urban infrastructure: shared bicycling systems. We show how station usage data from Barcelona's Bicing system (Figure 1) can be used to infer cultural and geographic aspects of the city and predict future bicycling station usage behavior, which corresponds to human movement in the city.

In particular, the main contributions of this paper are: (1) demonstrating the potential of using shared bicycling as a data source to gain insights into city dynamics and aggregated human behavior; (2) exploring the relationship between spatiotemporal patterns of bicycle usage and underlying city behavior and geography; and (3) studying patterns in bicycle station usage, including the prediction of usage patterns and an analysis of how factors such as the time of the day affect this prediction. In our analysis, we emphasize not just what the bicycling station usage data

UbiComp'09

HydroSense: Infrastructure-Mediated Single-Point Sensing of Whole-Home Water Activity

Jon Froehlich[1], Eric Larson[2], Tim Campbell[2], Conor Haggerty[4], James Fogarty[1], Shwetak N. Patel[1,2]
[1]Computer Science & Engineering, [2]Electrical Engineering,
[3]Mechanical Engineering, [4]Community, Environment, and Planning
DUB Institute, University of Washington
Seattle, WA 98195
{ jfroehli@cs, eclarson@u, tcampbell@u, conorh@u, jfogarty@cs, shwetak@cs } .washington.edu

ABSTRACT
Recent work has examined infrastructure-mediated sensing as a practical, low-cost, and unobtrusive approach to sensing human activity in the physical world. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed by their manifestations in an environment's existing infrastructures (e.g., a home's water, electrical, and HVAC infrastructures). This paper presents HydroSense, a low-cost and easily-installed single-point sensor of pressure within a home's water infrastructure. HydroSense supports both identification of activity at individual water fixtures within a home (e.g., a particular toilet, a kitchen sink, a particular shower) as well as estimation of the amount of water being used at each fixture. We evaluate our approach using data collected in ten homes. Our algorithms successfully identify fixture events with 97.9% aggregate accuracy and can estimate water usage with error rates that are comparable to empirical studies of traditional utility-supplied water meters. Our results both validate our approach and provide a basis for future improvements.

Author Keywords
Infrastructure-mediated sensing, activity sensing.

ACM Classification Keywords
H5.2. Information Interfaces and Presentation: User Interfaces; H1.2. Models and Principles: User/Machine Systems.

General Terms
Algorithms, Experimentation, Measurement

INTRODUCTION AND MOTIVATION
Effective feedback for sensing and modifying human activity in the physical world are a cornerstone of ubiquitous computing research and practice. Many complementary approaches have been developed, including recent interest in *infrastructure-mediated sensing* [7, 14, 15, 16, 17, 20]. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed via their manifestations in an environment's existing infrastructures (a home's water [7], electrical [15, 16, 17, 20], and HVAC [14] infrastructures). Because of its practical, low-cost, and unobtrusive nature, infrastructure-mediated sensing offers significant promise as a general method.

This work focuses on infrastructure-mediated sensing of home water activity. Water is essential to many home activities (e.g., washing, cleaning, cooking, drinking, gardening) which are in turn central to important potential ubiquitous computing applications (e.g., helping elders live more independently, helping people monitor their own water usage to reduce waste). Previous work monitoring home water usage [6, 7, 9] required multiple sensing points, exposed piping, could not infer *both* fixture and flow, and received limited or no validation in actual homes.

This paper presents HydroSense, a low-cost, *single-point* solution for activity sensing mediated by a home's existing water infrastructure. HydroSense is based on continuous analysis of *pressure* within a home's water infrastructure. Specifically, we *identify individual water fixtures* (e.g., a particular toilet, a kitchen sink, a particular shower) within a home according to the unique *pressure waves* that propagate to the sensor when valves are opened or closed. We also estimate the *amount of water being used* at a fixture based on the magnitude of the resulting *pressure drop* within the water infrastructure. Our work represents a significant advance over prior research in several regards:

First, HydroSense can be *easily installed* at any accessible location within a home's existing water infrastructure. Typical installations will be at an exterior hose bib, utility sink spigot, or water heater drain valve. If unavailable or not easily accessed (e.g., in an apartment unit), HydroSense can also be installed at the water connection point for a dishwasher, clothes washer, or toilet. All of these are simple screw-on installation points, with no need for a plumber.

Second, HydroSense's analysis of *pressure* provides the unique capability of sensing both the *individual fixture* at

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
UbiComp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.
Copyright 2009 ACM 978-1-60558-431-7/09/09...$10.00.

CHI'10

The Design of Eco-Feedback Technology

Jon Froehlich[1], Leah Findlater[2], James Landay[1]
[1]Computer Science and Engineering, [2]The Information School
DUB Institute, University of Washington
{ jfroehli, leahkf, landay}@uw.edu

ABSTRACT
Eco-feedback technology provides feedback on individual or group behaviors with a goal of reducing environmental impact. The history of eco-feedback extends back more than 40 years to the origins of environmental psychology. Despite its stated purpose, few HCI eco-feedback studies have attempted to measure behavior change. This leads to two overarching questions: (1) what can HCI learn from environmental psychology and (2) what role should HCI have in designing and evaluating eco-feedback technology? To help answer these questions, this paper conducts a comparative survey of eco-feedback technology, including 89 papers from environmental psychology and 44 papers from the HCI and UbiComp literature. We also provide an overview of predominant models of pro-environmental behaviors and a summary of key motivation techniques to promote this behavior.

Author Keywords
Eco-feedback, Environmental HCI, Reflective HCI, Survey

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI)

General Terms
Design, Human Factors

INTRODUCTION
As environmental issues such as climate change, air pollution, and water scarcity become more salient in the global consciousness, so too have they become more active targets of research within HCI and Ubiquitous Computing [6, 19, 57]. One particularly popular form of environmental HCI research is the design and study of *eco-feedback technology*, which we define as technology that provides feedback on individual or group behaviors with a goal of reducing environmental impact (adapted from [39] and [28], see Figure 1 for examples). Despite this goal, few HCI eco-feedback studies have even attempted to measure behavior change. Although eco-feedback may be seen as an extension of research in *persuasive technology* [17], it actually extends back much further to over 40 years of research in environmental psychology. This leads to two



Figure 1. Examples of eco-feedback technology. (left-to-right) The Ishiyopin display uses sensors and living plants to provide feedback about recycling and waste disposal [28]. WaterSpot provides ambient feedback information about water usage [3]. The UbiGreen Transportation Display semi-automatically senses and feeds back information about transportation to encourage green transit [19].

interrelated questions: (1) What can HCI learn from environmental psychology and (2) what should be the role of the HCI community in contributing to eco-feedback research? To explore these questions in detail, we present a review of the related environmental psychology literature as well as a comparative survey of eco-feedback studies in both HCI and environmental psychology.

Eco-feedback technology is based on the working hypothesis that most people lack awareness and understanding about how their everyday behaviors such as driving to work or showering affect the environment; technology may bridge this "environmental literacy gap" by automatically sensing these activities and feeding related information back through computerized means (e.g., mobile phones, ambient displays, or online visualizations). HCI and UbiComp researchers have built eco-feedback technologies for a variety of domains including energy consumption [28], water usage [3], transportation [19], and waste disposal practices [29].

Contributing to this growing interest in eco-feedback technology is the parallel advancement and availability of sensing systems for environmentally related activities (e.g., human activity inference [35]) and interactive displays to feedback this data (e.g., iPods and mobile phones). Such advances provide a rich space of opportunities for new types of eco-feedback that could not be considered in the past. Moreover, the next generation of resource measurement systems (often referred to as "smart meters") will soon provide real-time (or near real-time) data on electricity, gas, and water usage in homes and businesses. This will produce tremendous amounts of data that can be

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
Copyright 2010 ACM 978-1-60558-929-9/10/04...$10.00.

Pervasive'11

A Longitudinal Study of Pressure Sensing to Infer Real-World Water Usage Events in the Home

Jon Froehlich[1,*], Eric Larson[2,*], Elliot Saba[2], Tim Campbell[3], Les Atlas[2], James Fogarty[1], and Shwetak Patel[1,2]

[1] Computer Science and Engineering
[2] Electrical Engineering,
[3] Mechanical Engineering
University of Washington, Seattle
{jfroehli,eclarson,sabae,tcampbell,lesatlas, jfogarty,shwetak}@uw.edu

Abstract. We present the first longitudinal study of pressure sensing to infer *real-world* water usage events in the home (e.g., dishwasher, upstairs bathroom sink, downstairs toilet). In order to study the pressure-based approach *out in the wild*, we deployed a ground truth sensor network for five weeks in three homes and two apartments that *directly monitored* valve-level water usage by *fixtures* and *appliances*. We use this data to, first, demonstrate the practical challenges in constructing water usage activity inference algorithms and, second, to inform the design of a new probabilistic-based classification approach. Inspired by algorithms in speech recognition, our novel Bayesian approach incorporates template matching, a language model, grammar, and prior probabilities. We show that with a single pressure sensor, our probabilistic algorithm can classify real-world water usage at the fixture level with 90% accuracy and at the fixture-category level with 96% accuracy. With two pressure sensors, these accuracies increase to 94% and 98%. Finally, we show how our new approach can be trained with fewer examples than a strict template-matching approach alone.

Keywords: Water sensing, activity inference, sustainability, field deployments.

1 Introduction

Low-cost and easy-to-install methods to sense and model human activity in the home have long been a focus of UbiComp research. Because water is fundamental to many activities of human life (e.g., bathing, cooking), sensing disaggregated water usage has emerged as a particularly promising area for human activity inference in the home [6, 8, 19]. In addition, these sensing systems can play a vital role in collecting highly granular consumption information for enabling eco-feedback and sustainability applications (e.g., [7]). In previous work, we introduced HydroSense [8], a pressure-based sensing solution that disaggregates water usage at the fixture level from a single installation point. HydroSense identifies the unique pressure waves generated when

CHI'12

The Design and Evaluation of Prototype Eco-Feedback Displays for Fixture-Level Water Usage Data

Jon Froehlich[1,*], Leah Findlater[2,#], Marilyn Ostergren[3], Solai Ramanathan[4], Josh Peterson[3], Inness Wragg[4], Eric Larson[5], Fabia Fu[4], Mazhengmin Bai[5], Shwetak N. Patel[1,2], James A. Landay[1]
Computer Science and Engineering[1], Electrical Engineering[2],
Pre-Engineering[3], Interaction Design[4], DxArts[5], Information School[#]
University of Washington, Seattle
{ostergm, eclarson, shwetak, landay}@uw.edu
Computer Science[*]
College of Information Studies[#]
University of Maryland, College Park
{jonf, leahkf}@umd.edu

ABSTRACT
Few means currently exist for home occupants to learn about their water consumption: *e.g.*, where water use occurs, whether such use is excessive and what steps can be taken to conserve. Emerging water sensing systems, however, can provide detailed usage data at the level of individual water fixtures (*i.e., disaggregated* usage data). In this paper, we perform formative evaluations of two sets of novel eco-feedback displays that take advantage of this disaggregated data. The first display set isolates and examines specific elements of an eco-feedback design space such as *data and time granularity*. Displays in the second set act as *design probes* to elicit reactions about competition, privacy, and integration into domestic space. The displays were evaluated via an online survey of 651 North American respondents and in-home, semi-structured interviews with 10 families (20 adults). Our findings are relevant not only to the design of future water eco-feedback systems but also for other types of consumption (e.g., electricity and gas).

Author Keywords
Eco-feedback, water, sustainability, iterative design

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION
Cities across the world are facing an escalating demand for potable water due to growing populations, higher population densities and warmer climates [12,13]. As new sources of water become more environmentally and economically costly to extract, water suppliers and governments are shifting their focus from finding new supplies to using existing supplies more efficiently [13,17,18,20]. One challenge in improving residential efficiency, however, is the lack of awareness that occupants have about their in-home water consumption habits. This disconnect makes it difficult, even for motivated individuals, to make informed decisions about what steps can be taken to conserve [7].



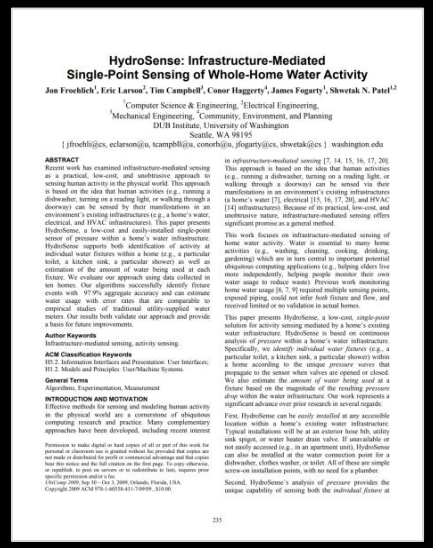
Figure 1: In our in-home interviews, participants selected preferred locations in their home to place our prototype water usage display.

Eco-feedback has been offered as one strategy to encourage conservation and help build the connection between home activities and resource use (see [4,6,9] for a review). However, most past work has focused on energy, with water-based eco-feedback largely limited to sensing and feedback at the *point-of-consumption* and to simple ambient and/or LED-based displays [2,19,21,22,30]. Although this type of feedback can potentially reduce usage at the installed fixture [30], it is limited in its ability to convey broader patterns of use or to compare across fixtures. These systems have also disproportionately focused on faucet and shower usage, which account for only 22% of water use in the average North American home [29].

In this paper, we explore a range of eco-feedback designs enabled by *disaggregated* (*i.e.*, fixture-level) water usage data. Our work is inspired by emerging technologies that can sense water usage at *individual fixtures* with only one or a few sensors [5,11]. Such detailed data presents new opportunities for water-based eco-feedback systems to visualize not only how much water is being consumed but also *where* and *when* it is occurring (e.g., upstairs bathroom toilet, front lawn sprinkler). The key question then becomes how to most effectively visualize this information? Moreover, what aspect(s) of the disaggregated data, if any, are people interested in, and what sort of reactions do these visualizations provoke?

To address these questions, we designed two sets of novel eco-feedback displays. The first set is designed to isolate and examine a subset of eco-feedback *design dimensions* [7,26] within the context of water usage (e.g., data and

UbiComp'09



HydroSense: Infrastructure-Mediated Single-Point Sensing of Whole-Home Water Activity

Jon Froehlich[1], Eric Larson[2], Tim Campbell[2], Conor Haggerty[3], James Fogarty[1], Shwetak N. Patel[1,2]

[1]Computer Science & Engineering, [2]Electrical Engineering,
[3]Mechanical Engineering, [4]Community, Environment, and Planning
DUB Institute, University of Washington
Seattle, WA 98195
{ jfroehli@cs, eclarson@u, tcampbll@u, conorh@u, jfogarty@cs, shwetak@cs } washington.edu

ABSTRACT

Recent work has examined infrastructure-mediated sensing as a practical, low-cost, and unobtrusive approach to sensing human activity in the physical world. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed by their manifestations in an environment's existing infrastructures (e.g., a home's water, electrical, and HVAC infrastructures). This paper presents HydroSense, a low-cost and easily-installed single-point sensor of pressure within a home's water infrastructure. HydroSense supports both identification of activity at individual water fixtures within a home (e.g., a particular toilet, a kitchen sink, a particular shower) as well as estimation of the amount of water being used at each fixture. We evaluate our approach using data collected in ten homes. Our algorithms successfully identify fixture events with 97.9% aggregate accuracy and can estimate water usage with error rates that are comparable to empirical studies of traditional utility-supplied water meters. Our results both validate our approach and provide a basis for future improvements.

Author Keywords

Infrastructure-mediated sensing, activity sensing.

ACM Classification Keywords

H5.2. Information Interfaces and Presentation: User Interfaces; H1.2. Models and Principles: User/Machine Systems.

General Terms

Algorithms, Experimentation, Measurement

INTRODUCTION AND MOTIVATION

Effective methods for sensing and modeling human activity in the physical world are a cornerstone of ubiquitous computing research and practice. Many complementary approaches have been developed, including recent interest in *infrastructure-mediated sensing* [7, 14, 15, 16, 17, 20]. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed via their manifestations in an environment's existing infrastructures (a home's water [7], electrical [15, 16, 17, 20], and HVAC [14] infrastructures). Because of its practical, low-cost, and unobtrusive nature, infrastructure-mediated sensing offers significant promise as a general method.

This work focuses on infrastructure-mediated sensing of home water activity. Water is essential to many home activities (e.g., washing, cleaning, cooking, drinking, gardening) which are in turn central to important potential ubiquitous computing applications (e.g., helping elders live more independently, helping people monitor their own water usage to reduce waste). Previous work monitoring home water usage [6, 7, 9] required multiple sensing points, exposed piping, could not infer *both* fixture and flow, and received limited or no validation in actual homes.

This paper presents HydroSense, a low-cost, *single-point* solution for activity sensing mediated by a home's existing water infrastructure. HydroSense is based on continuous analysis of *pressure* within a home's water infrastructure. Specifically, we *identify individual water fixtures* (e.g., a particular toilet, a kitchen sink, a particular shower) within a home according to the unique *pressure waves* that propagate to the sensor when valves are opened or closed. We also estimate the *amount of water being used* at a fixture based on the magnitude of the resulting *pressure drop* within the water infrastructure. Our work represents a significant advance over prior research in several regards:

First, HydroSense can be *easily installed* at any accessible location within a home's existing water infrastructure. Typical installations will be at an exterior hose bib, utility sink spigot, or water heater drain valve. If unavailable or not easily accessed (e.g., in an apartment unit), HydroSense can also be installed at the water connection point for a dishwasher, clothes washer, or toilet. All of these are simple screw-on installation points, with no need for a plumber.

Second, HydroSense's analysis of *pressure* provides the unique capability of sensing both the *individual fixture* at

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
UbiComp 2009, Sept 30 – Oct 3, 2009, Orlando, Florida, USA.
Copyright 2009 ACM 978-1-60558-431-7/09/09...$10.00.

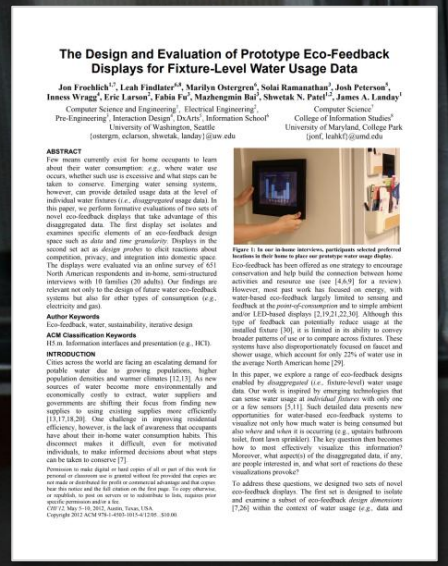


ERIC LARSON



Froehlich *et al.*, UbiComp'09

CHI'12



The Design and Evaluation of Prototype Eco-Feedback Displays for Fixture-Level Water Usage Data

Jon Froehlich[1,7], Leah Findlater[4,8], Marilyn Ostergren[4], Solai Ramanathan[3], Josh Peterson[3], Inness Wragg[4], Eric Larson[2], Fabia Fu[5], Mazhengmin Bai[5], Shwetak N. Patel[1,2], James A. Landay[1]

Computer Science and Engineering[1], Electrical Engineering[2], Pre-Engineering[3], Interaction Design[4], DxArts[5], Information School[6]
University of Washington, Seattle
{ostergm, eclarson, shwetak, landay}@uw.edu

Computer Science[7], College of Information Studies[8]
University of Maryland, College Park
{jonf, leahkf}@umd.edu

ABSTRACT

Few means currently exist for home occupants to learn about their water consumption: *e.g.*, where water use occurs, whether such use is excessive and what steps can be taken to conserve. Emerging water sensing systems, however, can provide detailed usage data at the level of individual water fixtures (*i.e., disaggregated* usage data). In this paper, we perform formative evaluations of two sets of novel eco-feedback displays that take advantage of this disaggregated data. The first display set isolates and examines specific elements of an eco-feedback design space such as *data* and *time granularity*. Displays in the second set act as *design probes* to elicit reactions about competition, privacy, and integration into domestic space. The displays were evaluated via an online survey of 651 North American respondents and in-home, semi-structured interviews with 10 families (20 adults). Our findings are relevant not only to the design of future water eco-feedback systems but also for other types of consumption (*e.g.*, electricity and gas).

Author Keywords

Eco-feedback, water, sustainability, iterative design

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION

Cities across the world are facing an escalating demand for potable water due to growing populations, higher population densities and warmer climates [12,13]. As new sources of water become more environmentally and economically costly to extract, water suppliers and governments are shifting their focus from finding new supplies to using existing supplies more efficiently [13,17,18,20]. One challenge in improving residential efficiency, however, is the lack of awareness that occupants have about their in-home water consumption habits. This disconnect makes it difficult, even for motivated individuals, to make informed decisions about what steps can be taken to conserve [7].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI'12, May 5-10, 2012, Austin, Texas, USA.
Copyright 2012 ACM 978-1-4503-1015-4/12/05...$10.00.

Figure 1: In our in-home interviews, participants selected preferred locations in their home to place our prototype water usage display.

Eco-feedback has been offered as one strategy to encourage conservation and help build the connection between home activities and resource use (see [4,6,9] for a review). However, most past work has focused on energy, with water-based eco-feedback largely limited to sensing and feedback at the *point-of-consumption* and to simple ambient and/or LED-based displays [2,19,21,22,30]. Although this type of feedback can potentially reduce usage at the installed fixture [30], it is limited in its ability to convey broader patterns of use or to compare across fixtures. These systems have also disproportionately focused on faucet and shower usage, which account for only 22% of water use in the average North American home [29].

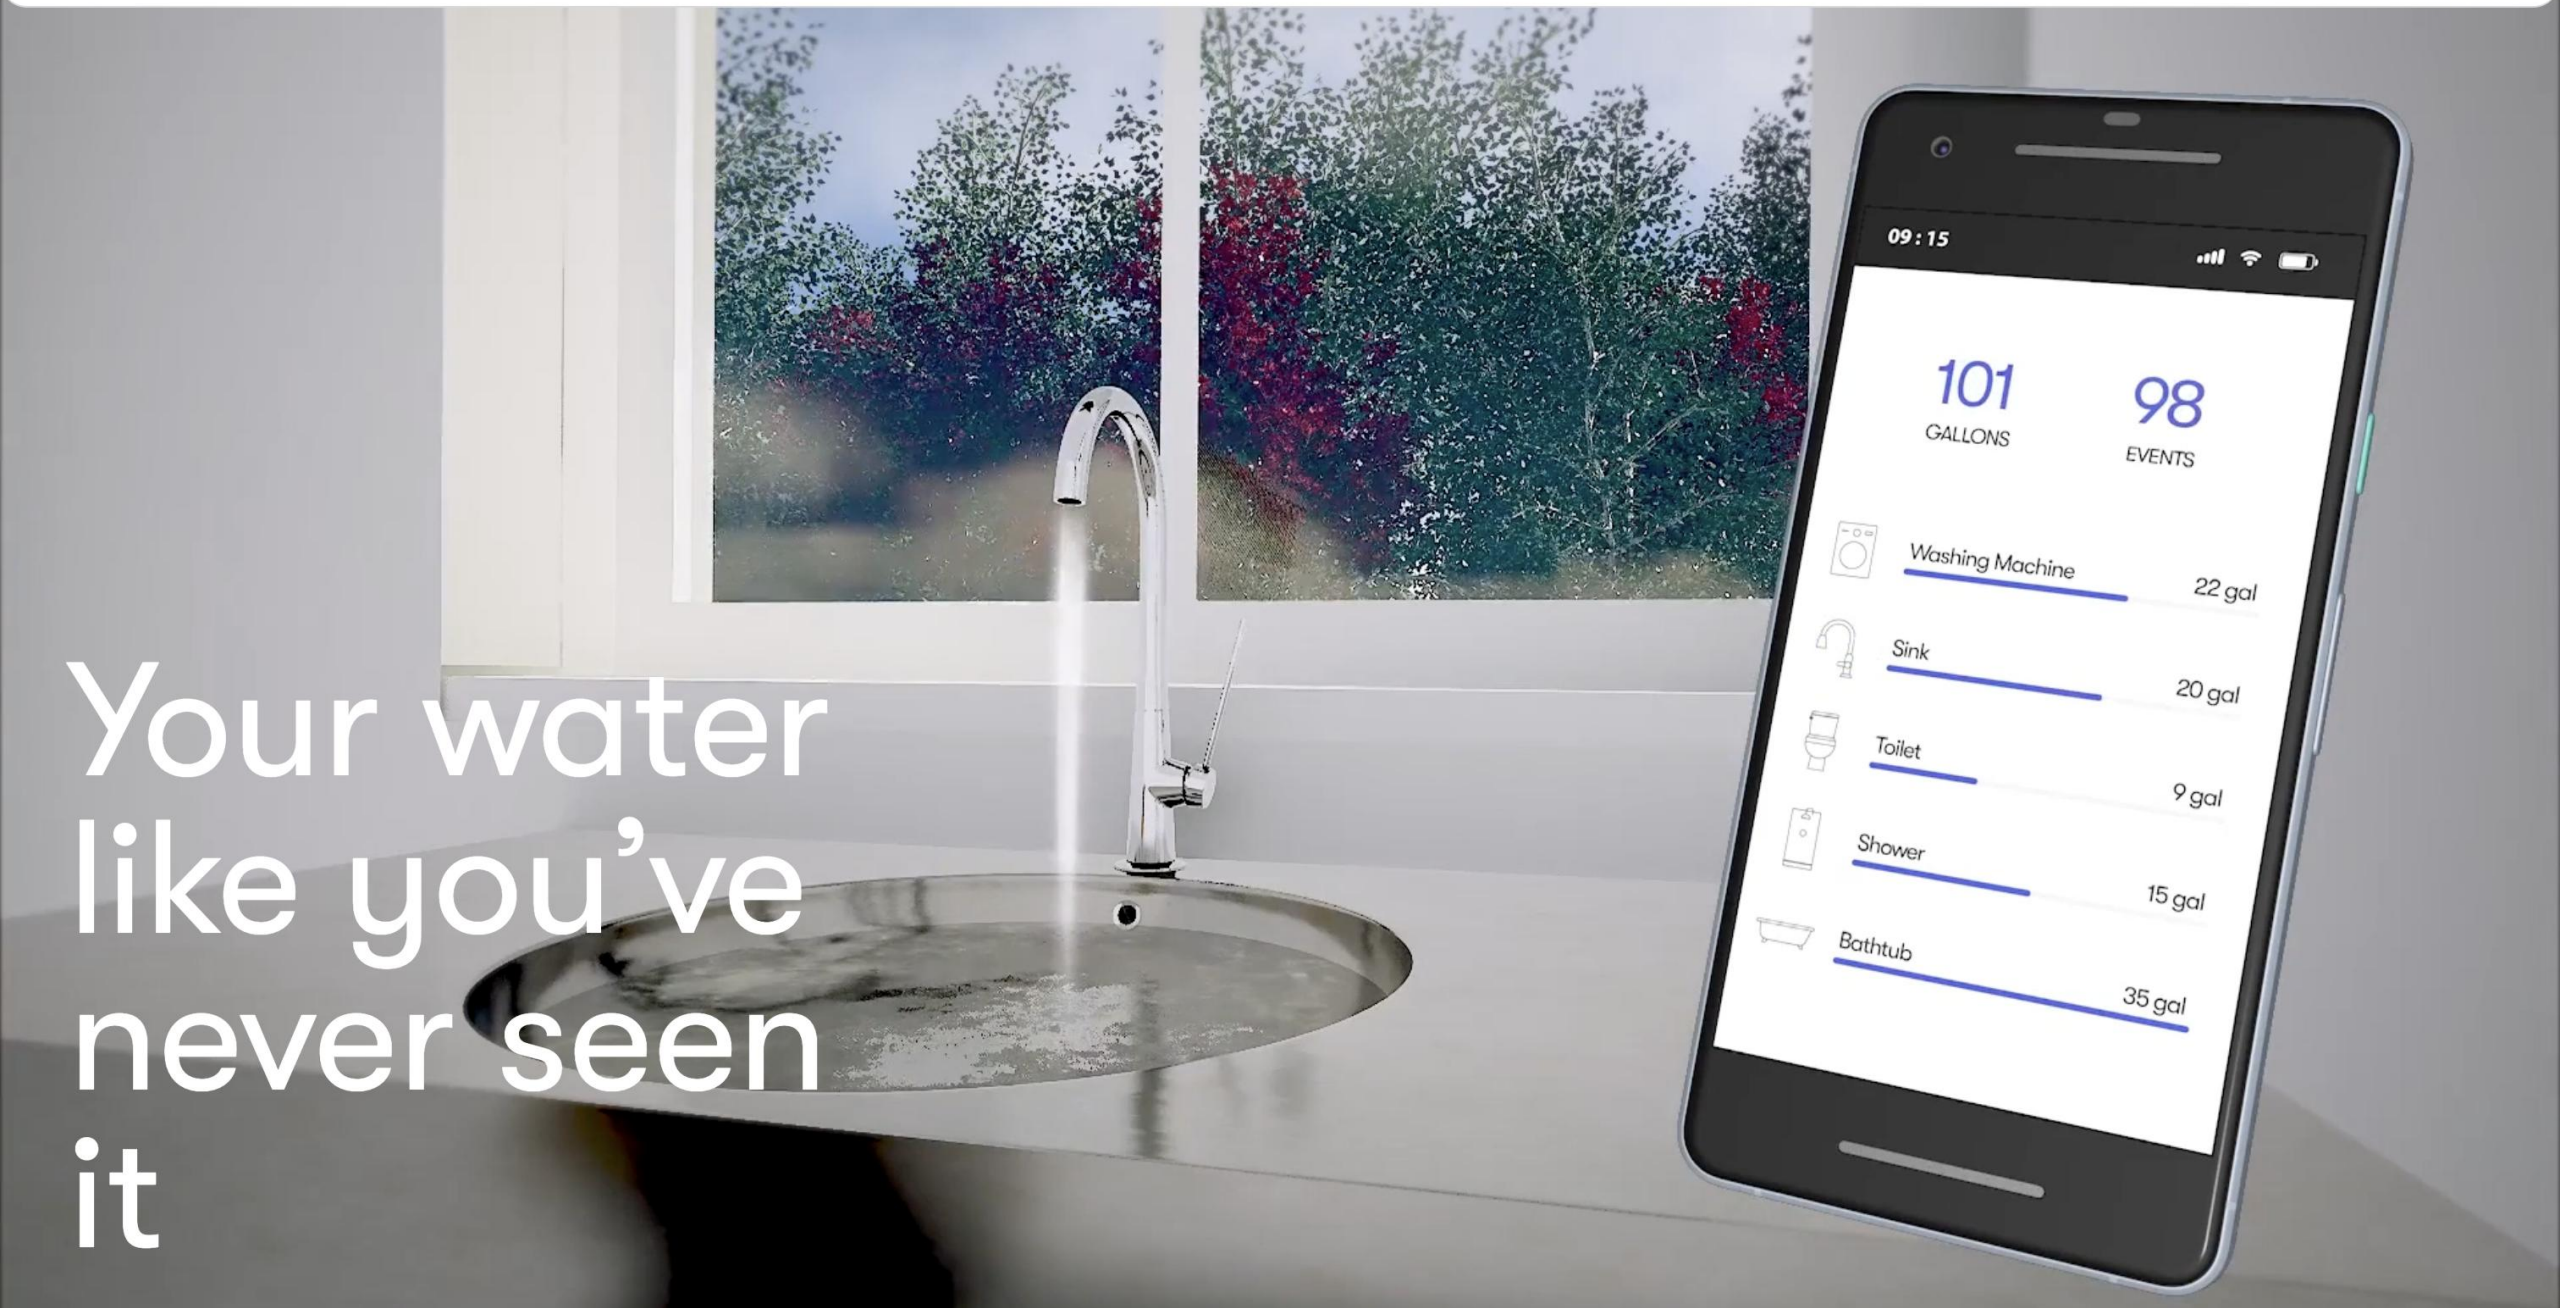
In this paper, we explore a range of eco-feedback designs enabled by *disaggregated* (*i.e.*, fixture-level) water usage data. Our work is inspired by emerging technologies that can sense water usage at *individual fixtures* with only one or a few sensors [5,11]. Such detailed data presents new opportunities for water-based eco-feedback systems to visualize not only how much water is being consumed but also *where* and *when* it is occurring (e.g., upstairs bathroom toilet, front lawn sprinkler). The key question then becomes how to most effectively visualize this information? Moreover, what aspect(s) of the disaggregated data, if any, are people interested in, and what sort of reactions do these visualizations provoke?

To address these questions, we designed two sets of novel eco-feedback displays. The first set is designed to isolate and examine a subset of eco-feedback *design dimensions* [7,26] within the context of water usage (e.g., data and

WATER ECO-FEEDBACK DISPLAY



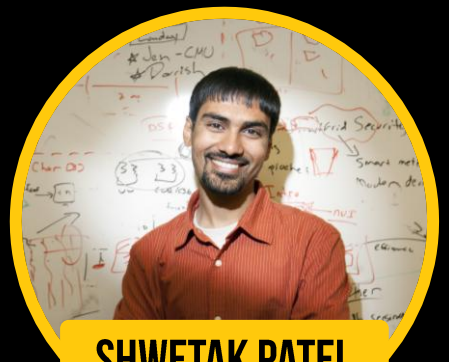
Froehlich *et al.*, CHI'12



Your water like you've never seen it



JAMES LANDAY
PhD Co-Advisor



SHWETAK PATEL
PhD Co-Advisor

GRADUATED 2011

Photo by Bruce Hemingway


MAKEABILITY LAB
University of Maryland


2013
MAKEABILITY LAB
HCIL
MAKEABILITY LAB

2024

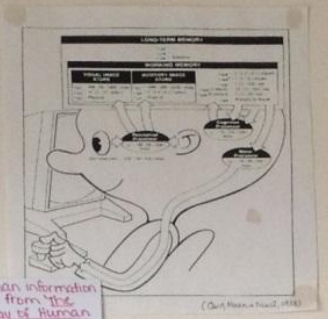
MAKEABILITY LAB



Our Mission

DESIGN, BUILD, & STUDY INTERACTIVE TOOLS & TECHNIQUES TO ADDRESS PRESSING SOCIETAL CHALLENGES

HCIL Begins



The human information processor from "The Psychology of Human Computer Interaction"

1983

Ben Shneiderman
Founding Director

HCIL
Human-Computer Interaction Lab



NOOBIE 1986
(A. DRUIN'S MASTER'S THESIS AT MIT)



Lab in Comp. & Space Sciences Bldg pre-1988



Ben Shneiderman Trading Card



SIGCHI SOCIAL IMPACT AWARD 2010

HCIL

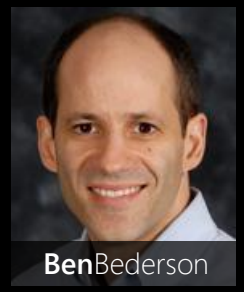


Photo by Ben Shneiderman


LeahFindlater


JenGolbeck


BenShneiderman

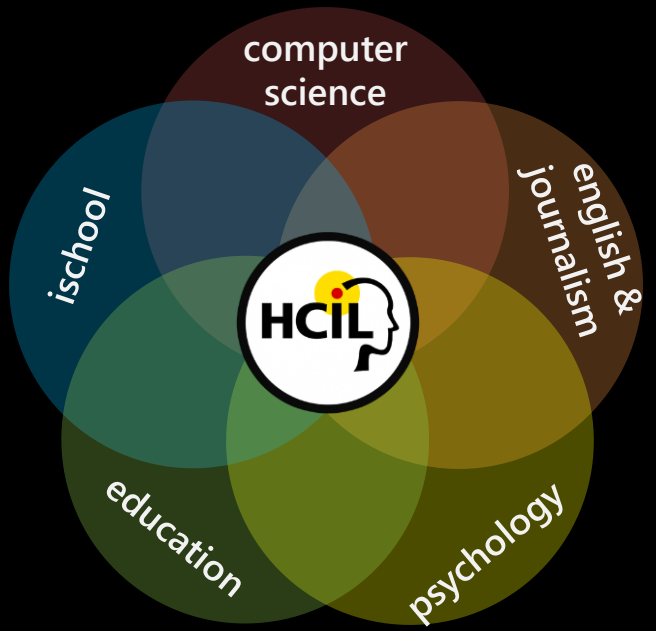

BenBederson


JonFroehlich


AnneRose

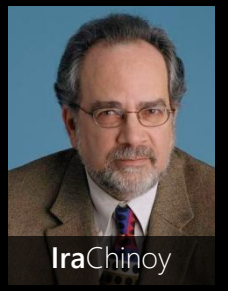

CatherinePlaisant


MarshiniChetty


computer science
ischool
english & journalism
education
psychology
HCIL

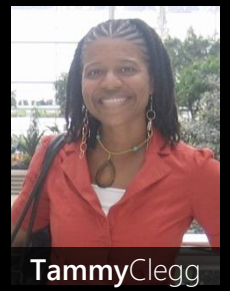

KariKraus

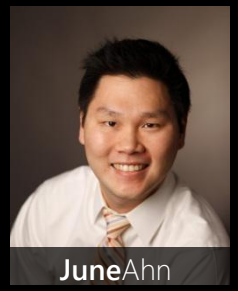

JennyPreece

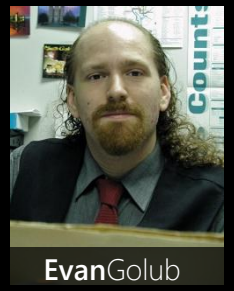

IraChinoy

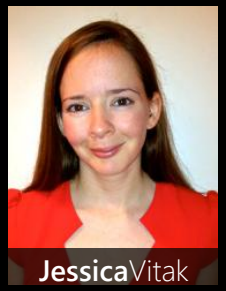

AllisonDruin


MonaLeighGuha


TammyClegg


JuneAhn


EvanGolub


JessicaVitak


KentNorman

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your work
6. The **people matter** the most

PARTICIPATORY DESIGN

Cooperative Inquiry











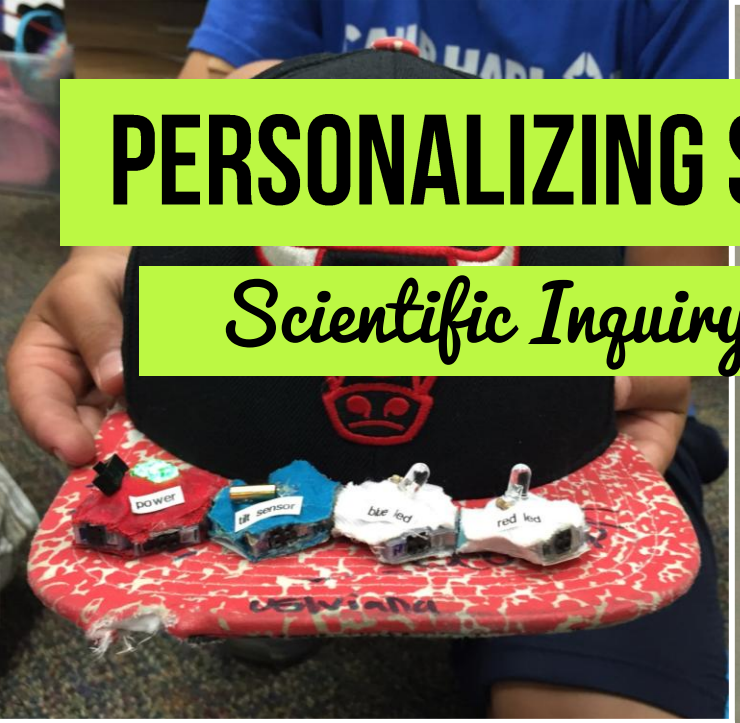
ALLISON DRUIN

Cooperative Inquiry: Druin, CHI'99; Guha, Druin, & Fails, 2013

Photos by Jon E. Froehlich

PERSONALIZING SCIENCE

Scientific Inquiry





Clegg, IDC'2012



TAMARA CLEGG

Photos by Jon E. Froehlich

SIGCHI SOCIETAL IMPACT TALK

SOCIAL IMPACT AREAS



ACCESSIBILITY



HEALTH
& WELLNESS



ENVIRONMENTAL
SUSTAINABILITY



STEM
EDUCATION



Norooz *et al.*, CHI'15

Photo by Jon E. Froehlich

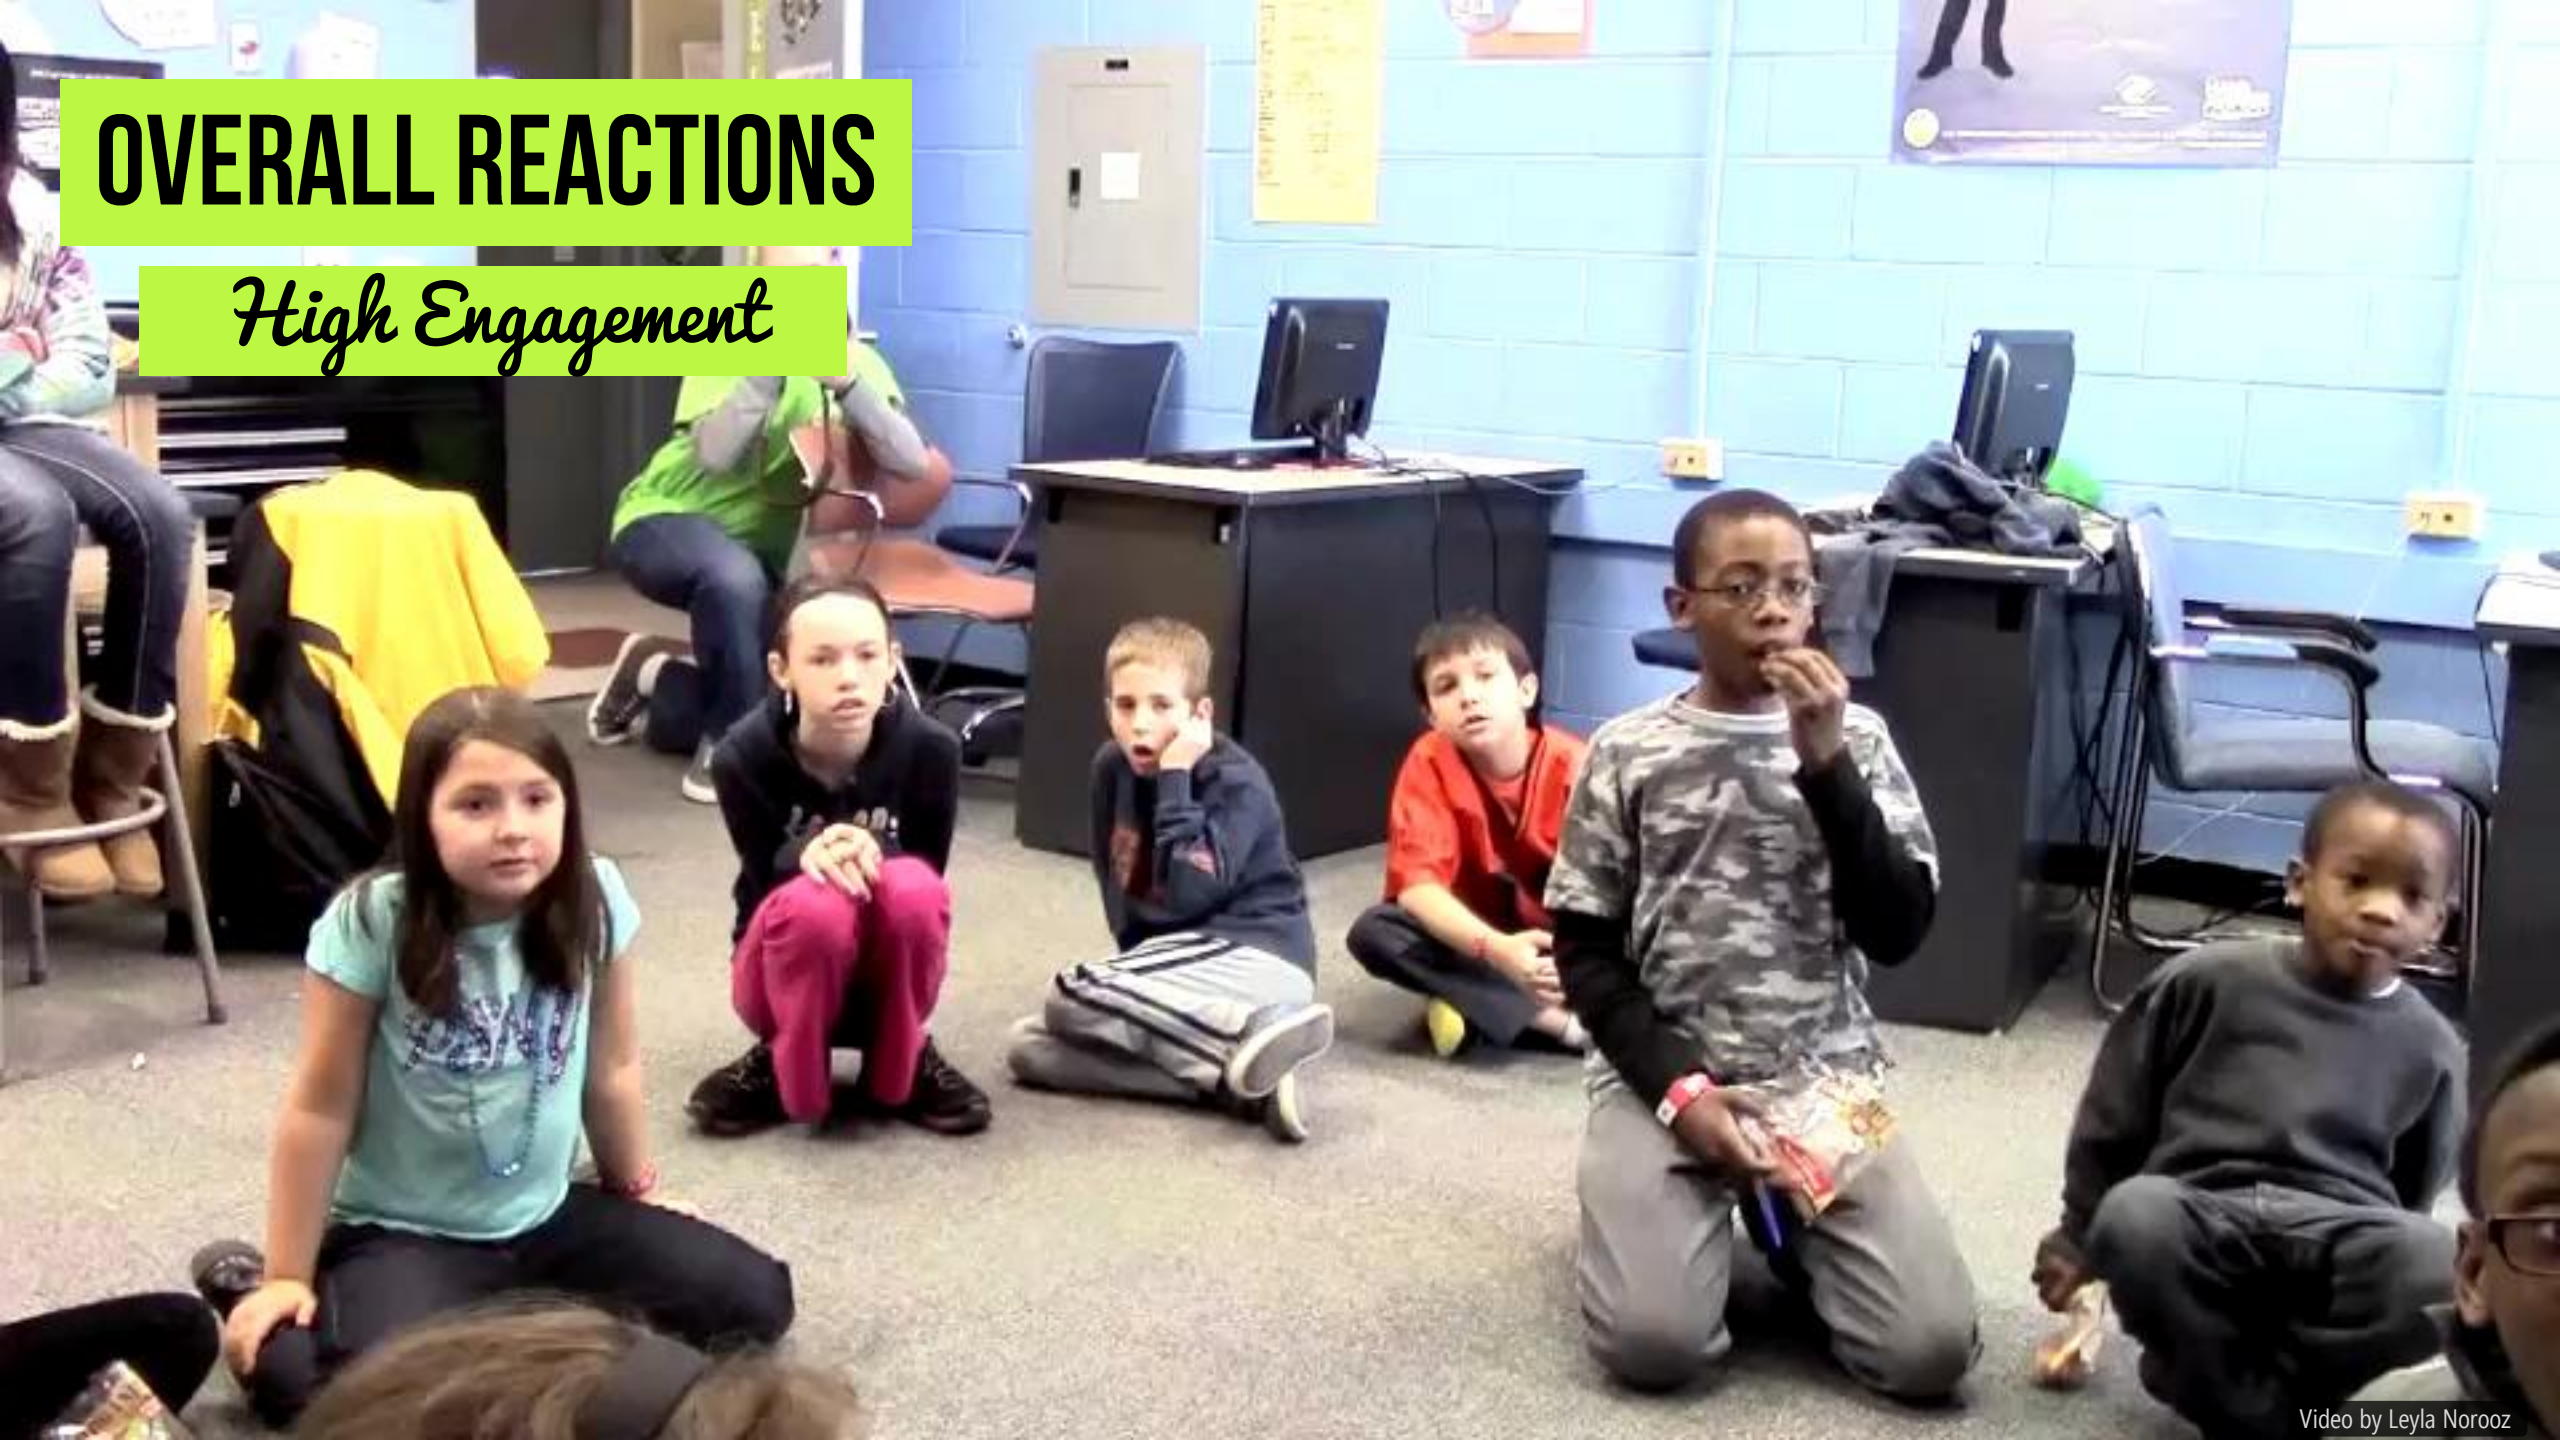


The heart and lungs visualize wearers' live heart and breathing rate.

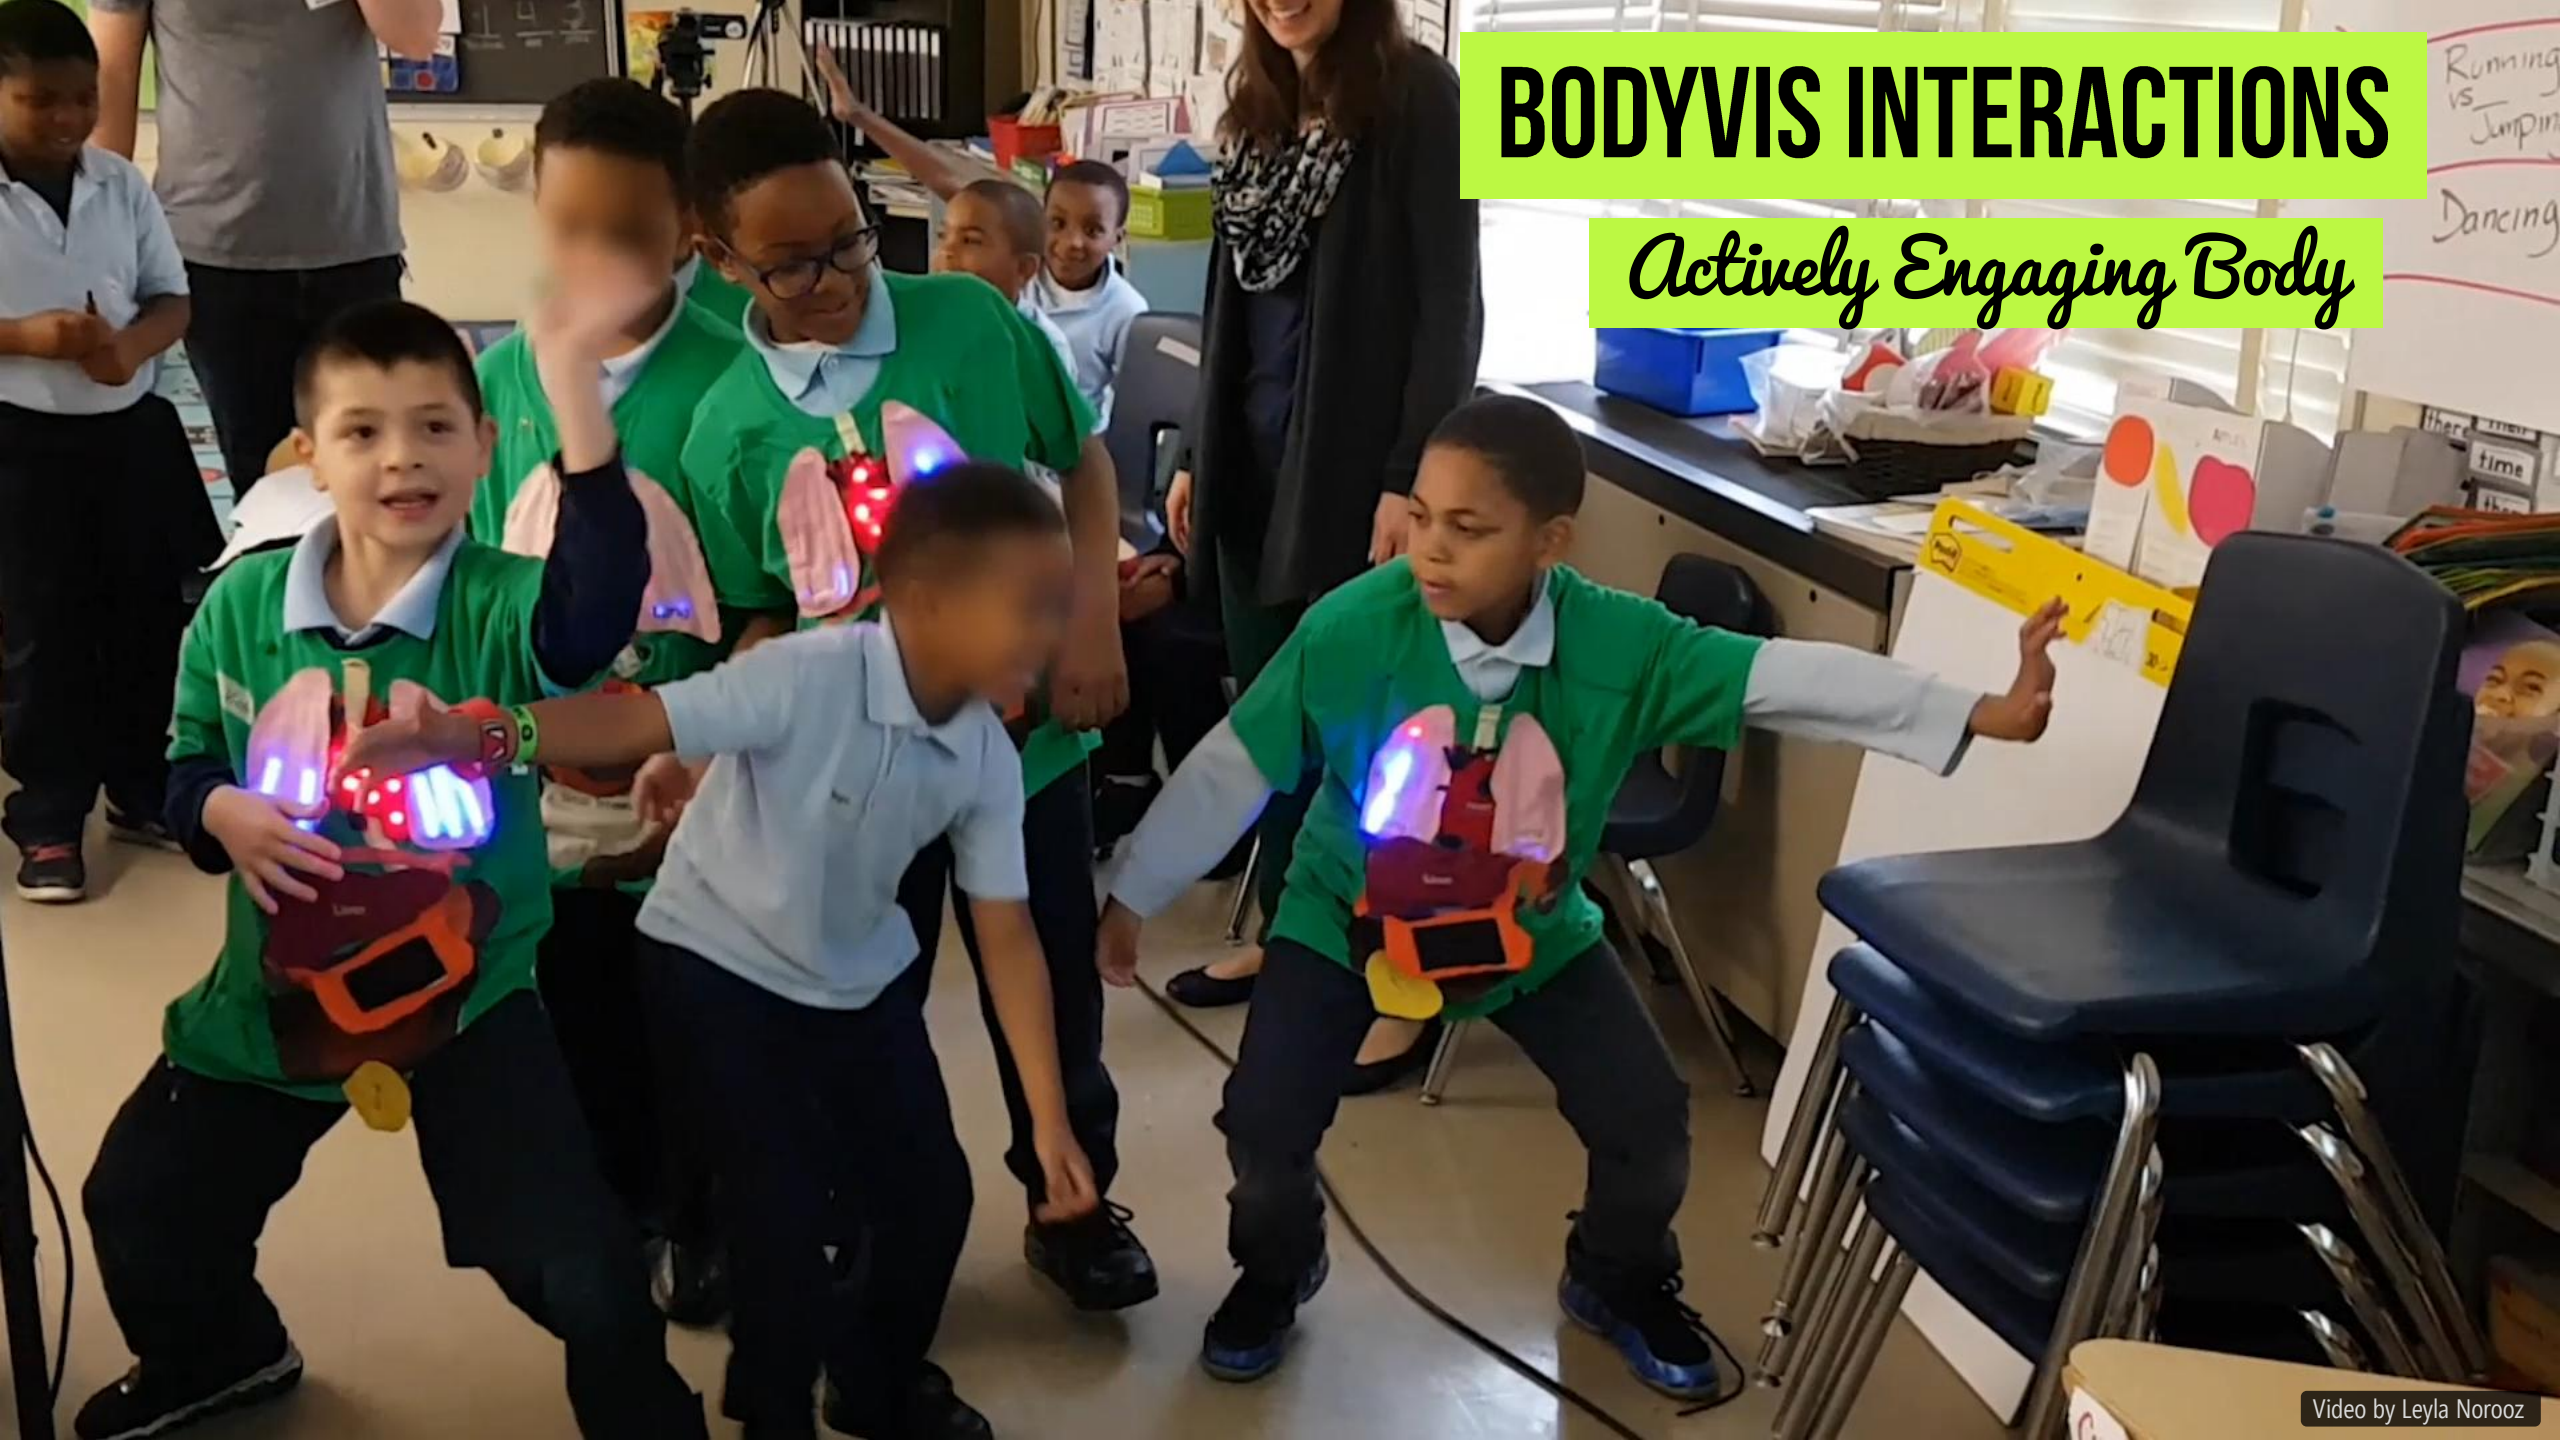
Video by Leyla Norooz

OVERALL REACTIONS

High Engagement

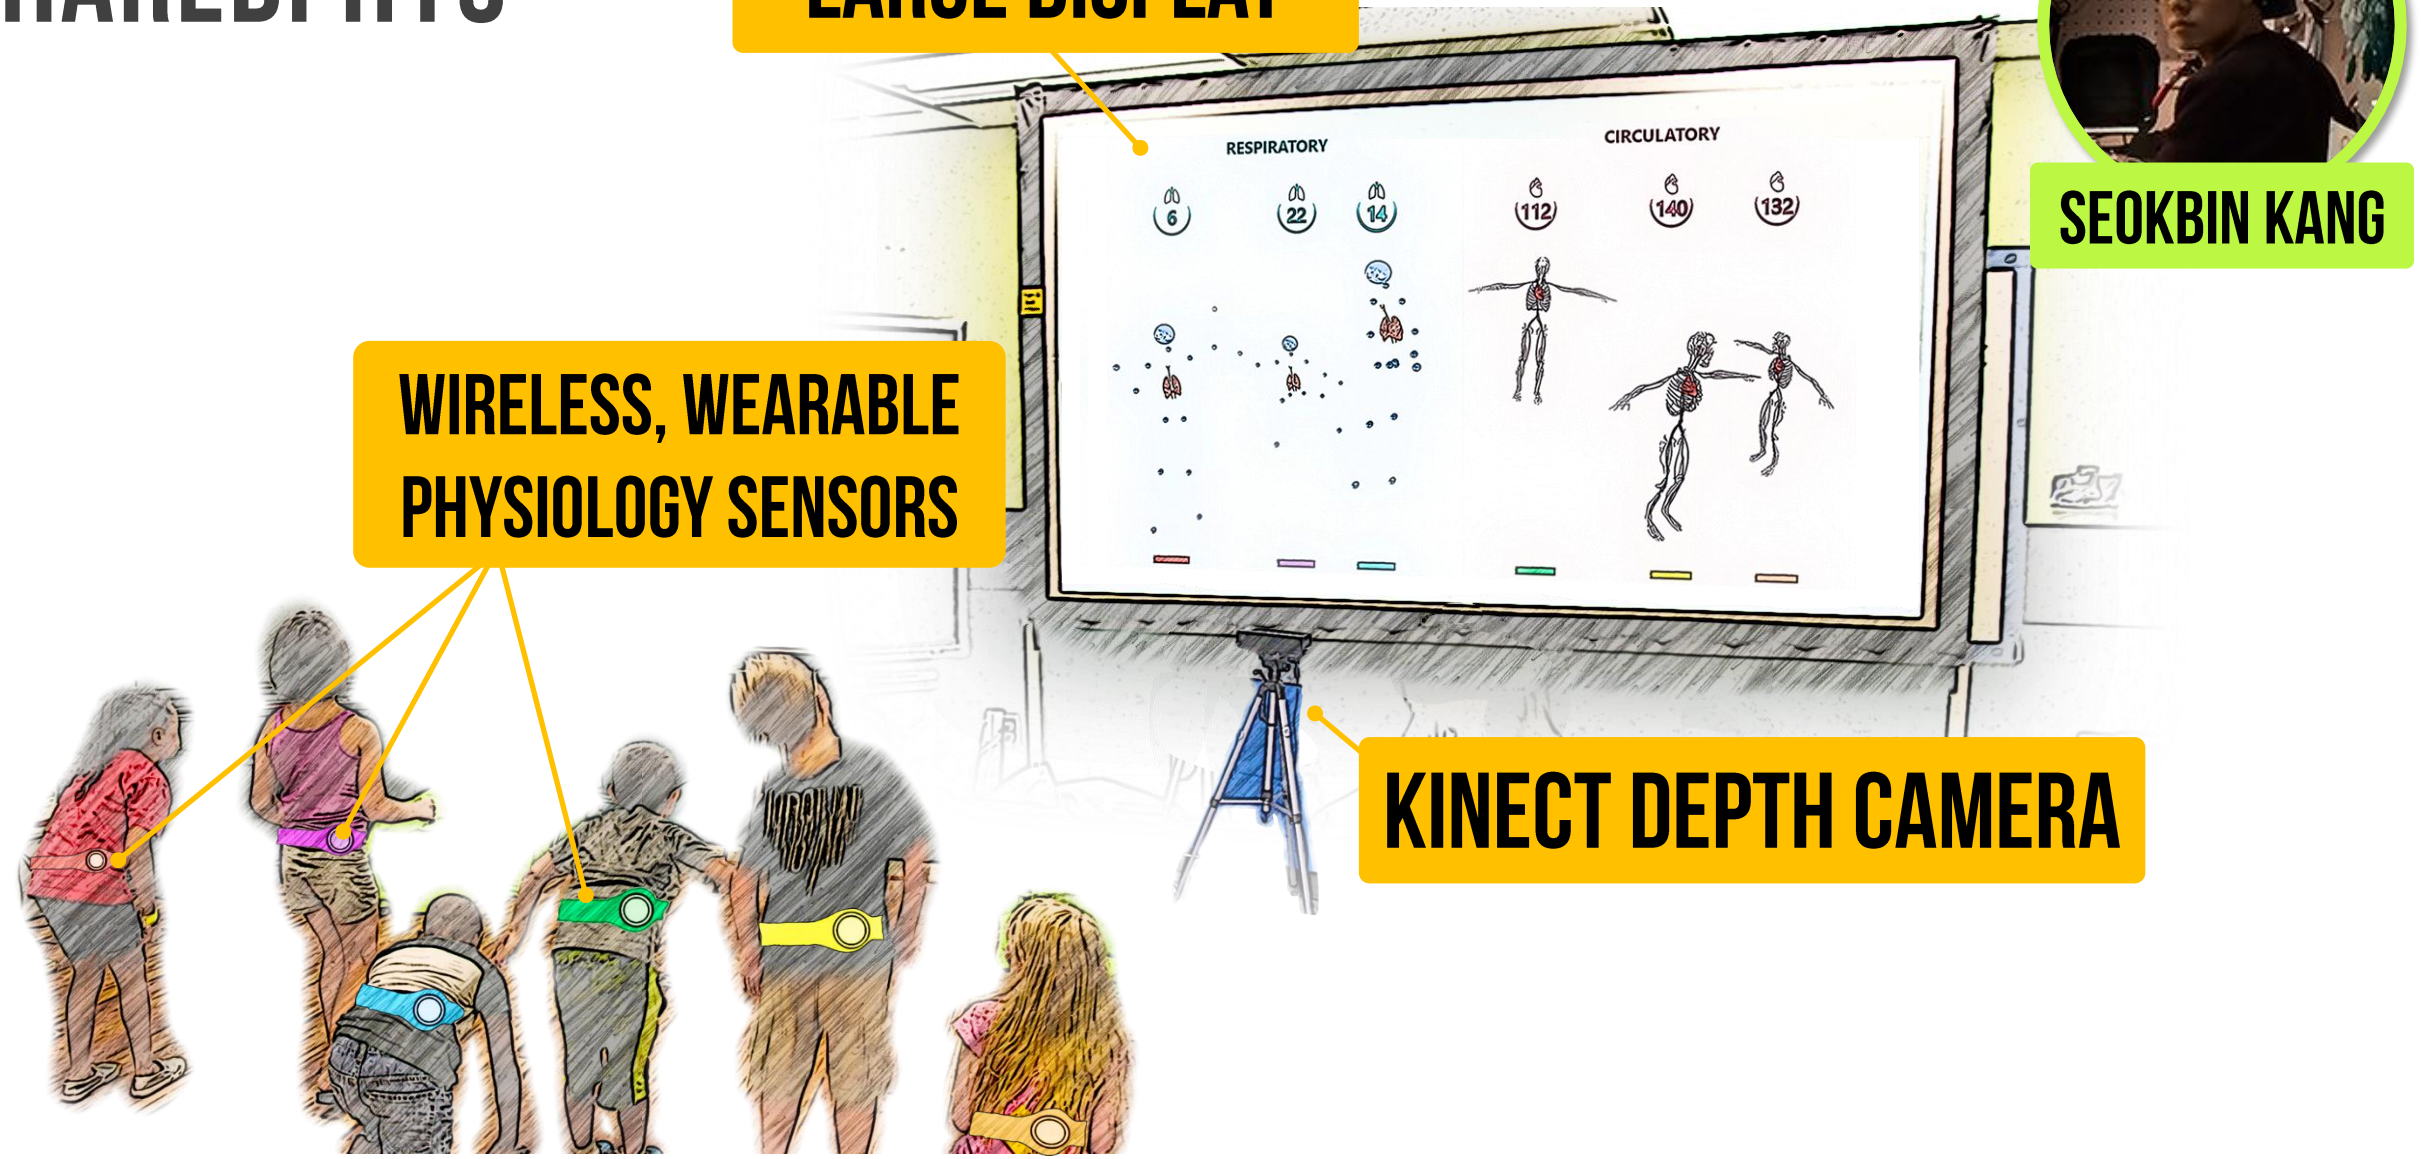


Video by Leyla Norooz


BODYVIS INTERACTIONS
Actively Engaging Body
Video by Leyla Norooz

STEM EDUCATION

SHAREDPHYS



Kang *et al.,* IDC'17; Clegg, CHI'17; Byrne, ICLS'18

PROTOTYPE #1

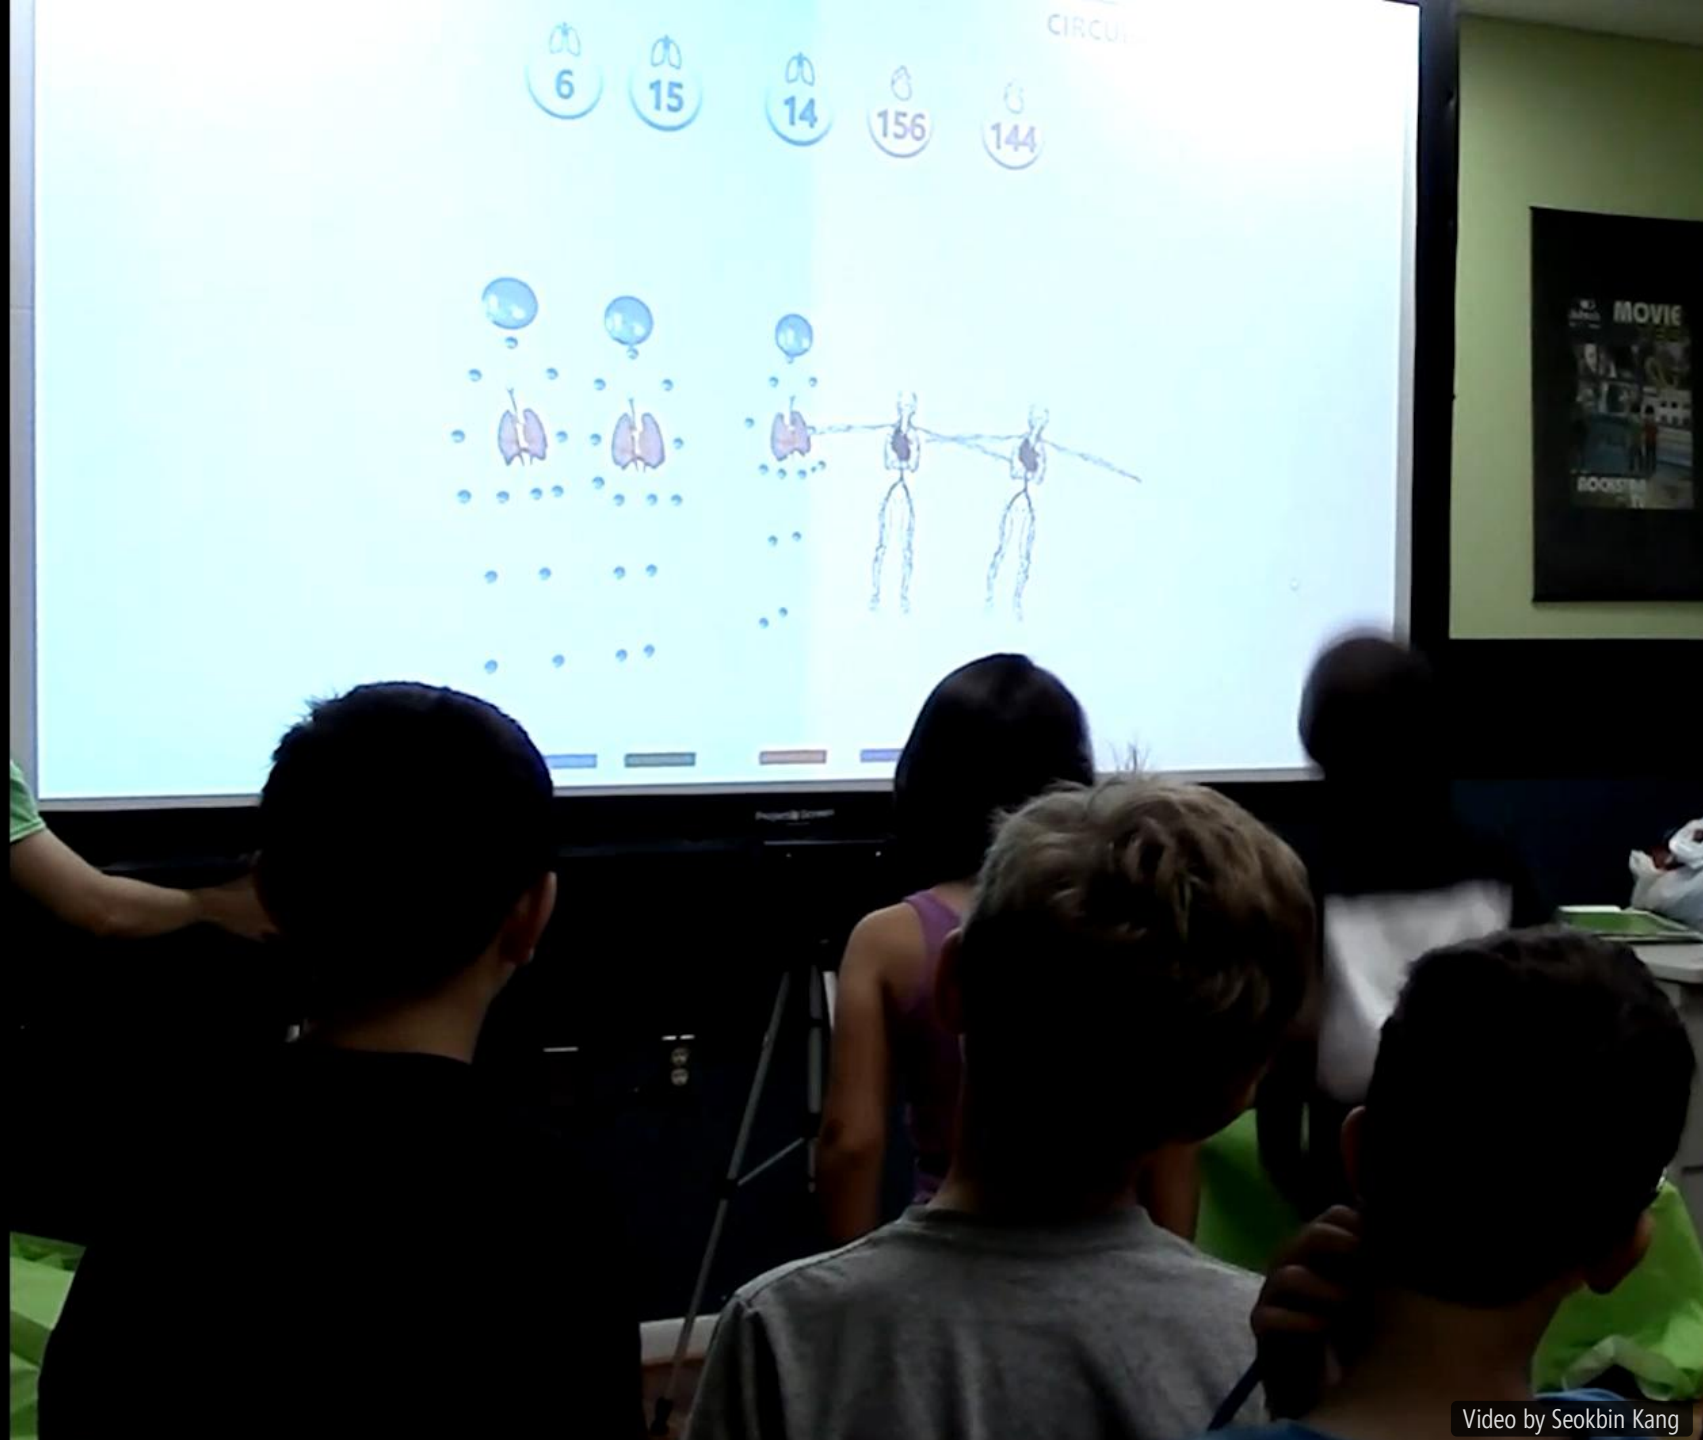
Magic Mirror

Mirror metaphor

Peer inside live body

Whole-body interactio

Kang *et al.,* IDC'17; Clegg, CHI'17; Byrne, ICLS'18



Video by Seokbin Kang

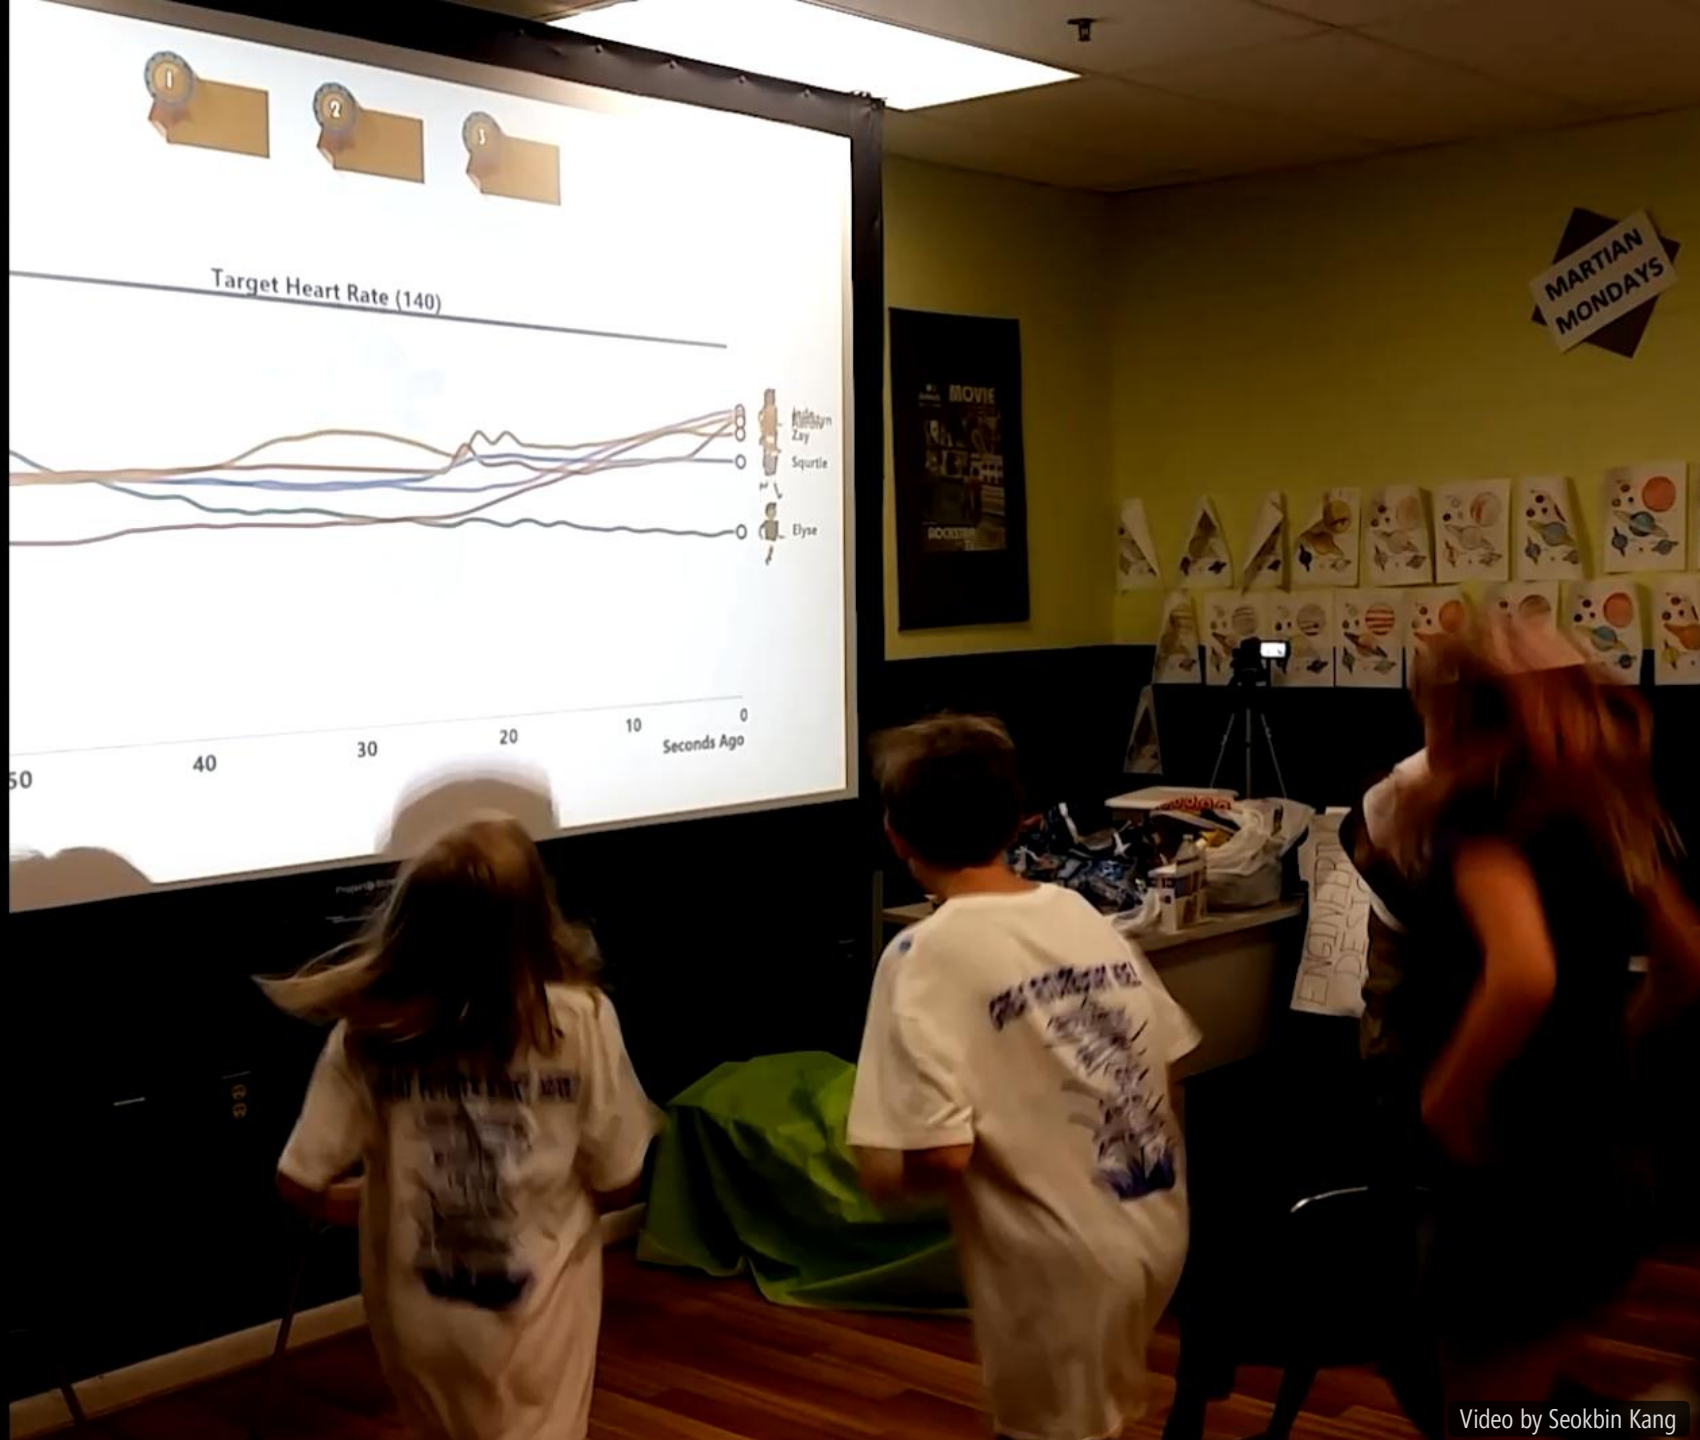
Prototype #2

Moving Graphs

Social comparison

Hypothesis testing

Basic statistics



Kang *et al.,* IDC'17; Clegg, CHI'17; Byrne, ICLS'18

Video by Seokbin Kang

UNEXPECTED FINDING

How Does It Work?







STEM EDUCATION

MAKERWEAR



MAJEED KAZEMITABAAR



LIANG HE



IDC'15, CHI'16 Best Poster, CHI'17 Best Paper

How can we...

enable young children to build their own interactive wearables?



Photo by Majeed Kazemitabaar


1
MAKERWEAR SYSTEM
TANGIBLE MODULES
Photo by Majeed Kazemitabaar

2 MAKERWEAR SYSTEM
MAGNETIC SOCKET MESH

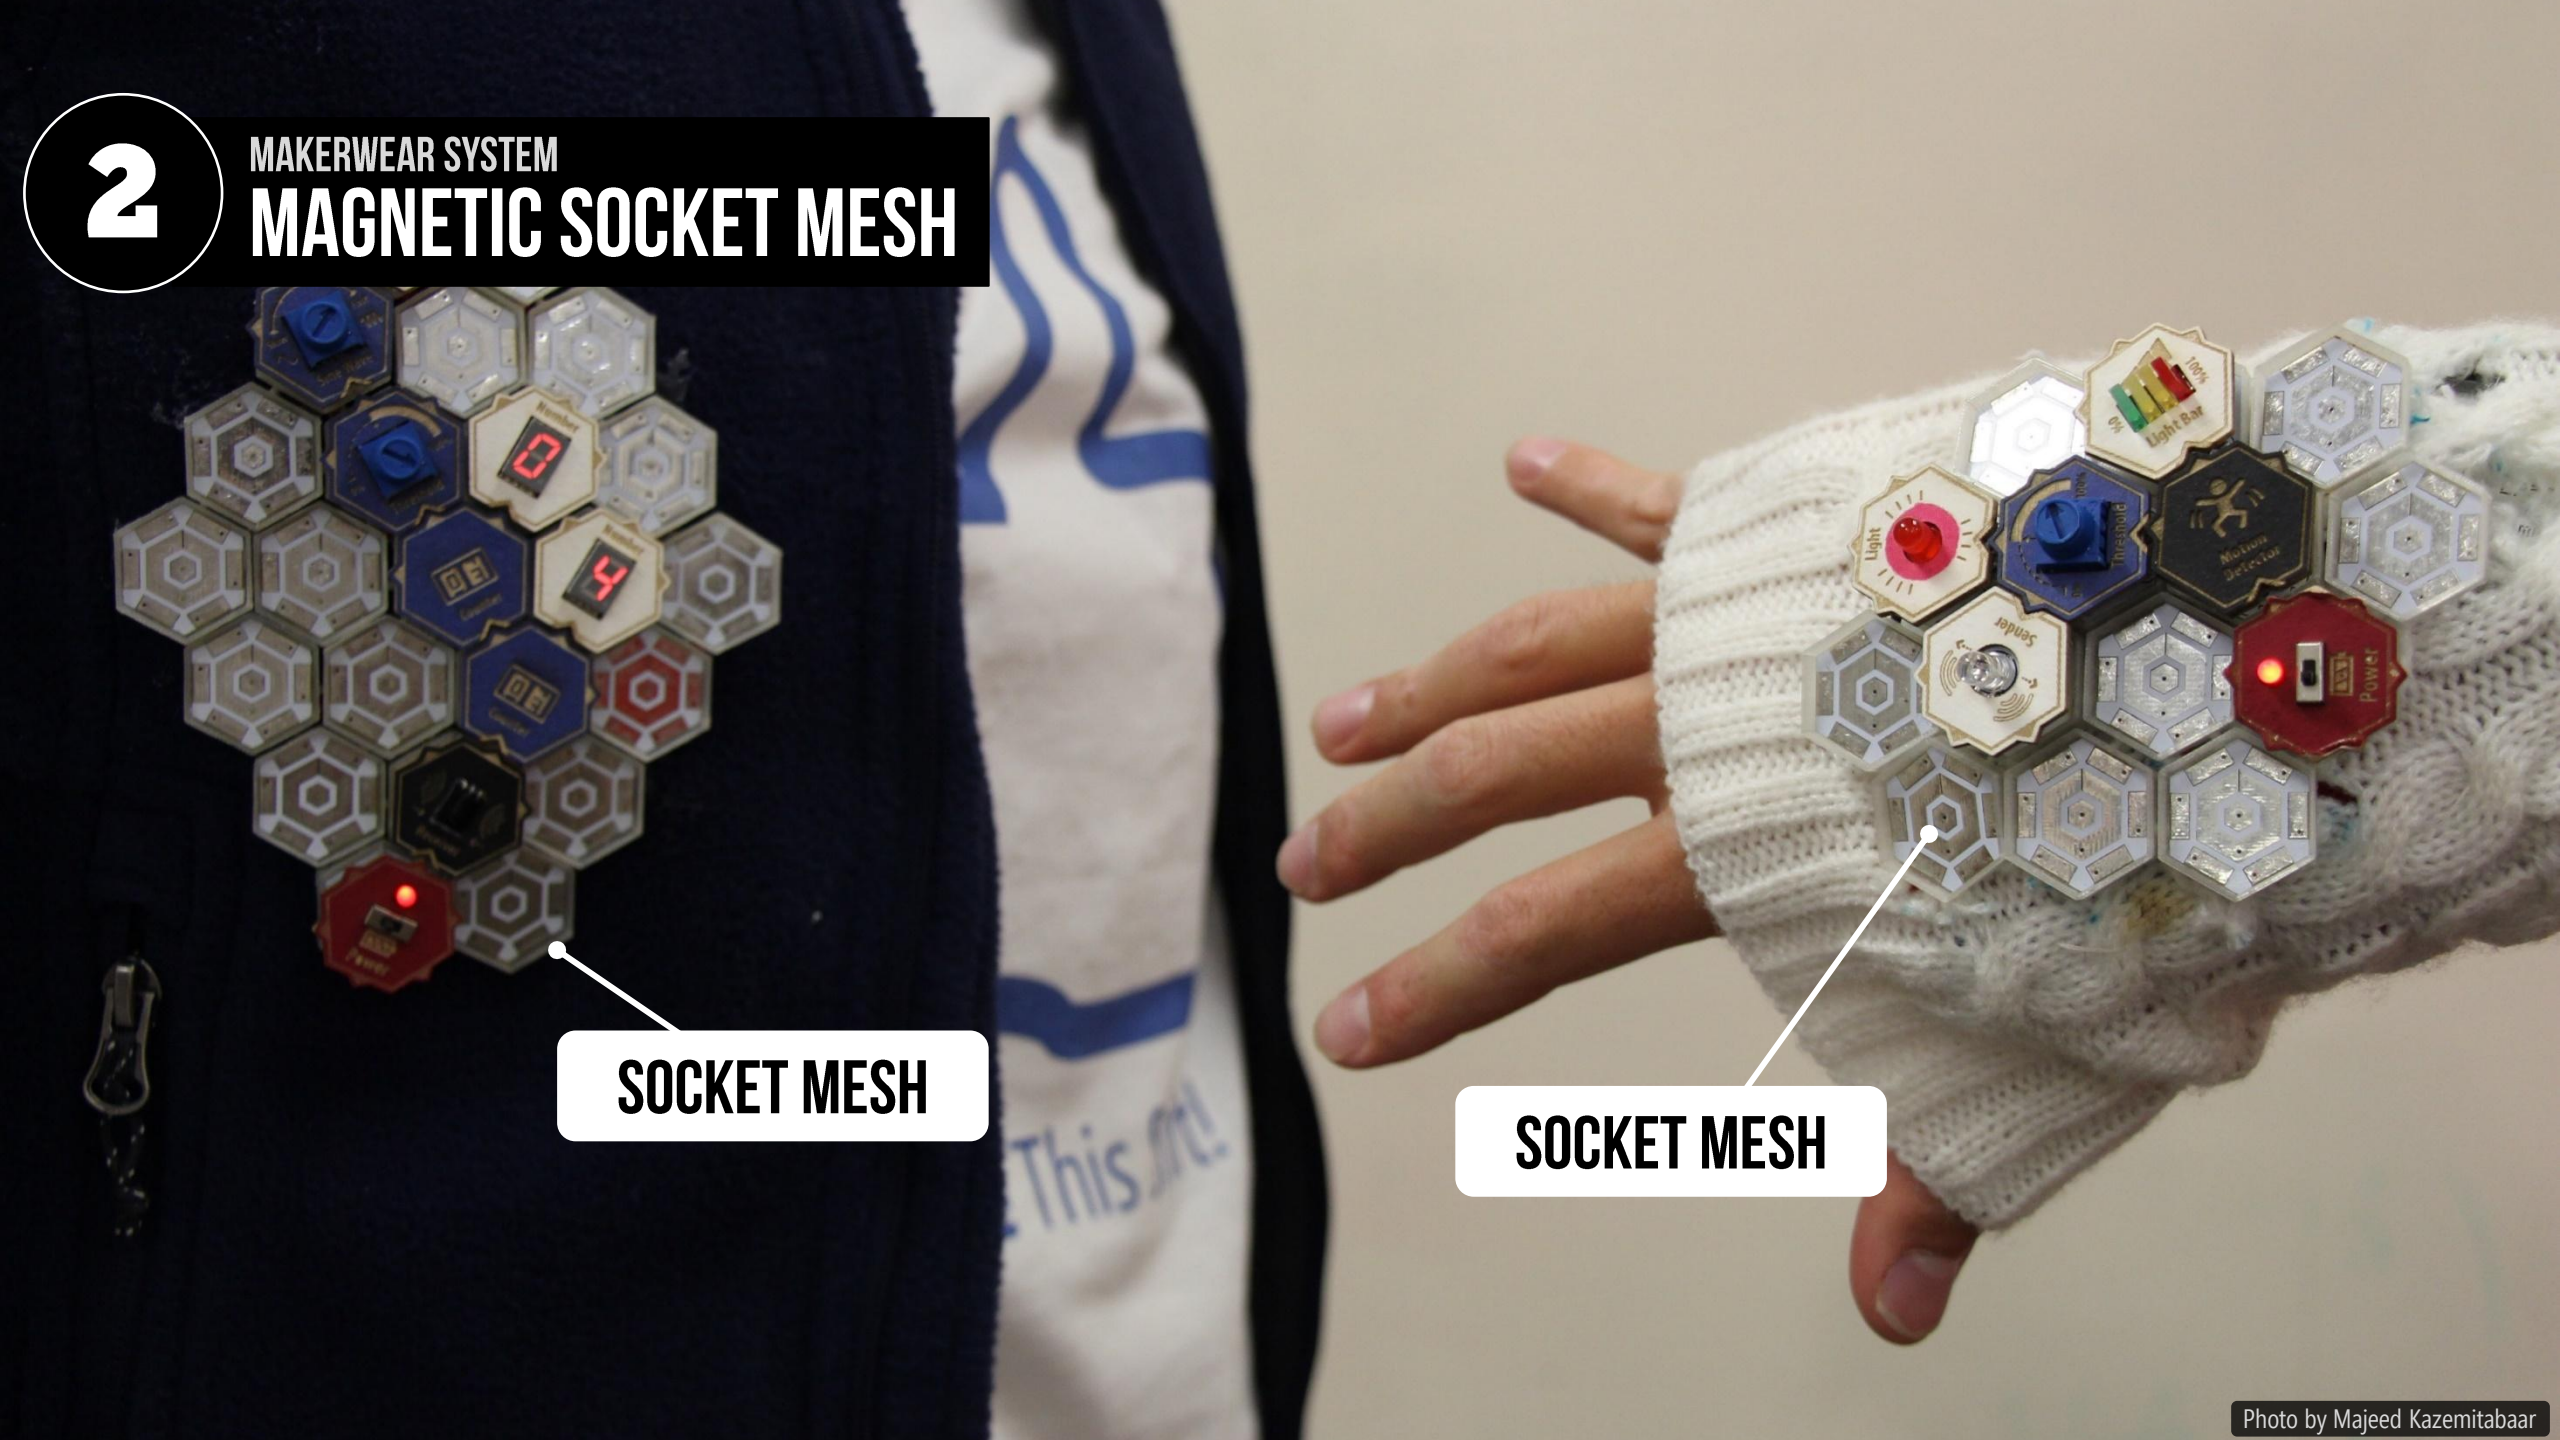
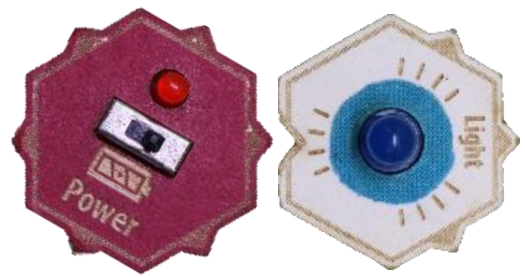


Photo by Majeed Kazemitabaar


Power


Power
Light


Power
Motion Detector
Light

MAKERWEAR EXAMPLES

“MOTION-REACTIVE CLOTHES”





Motion-reactive clothes!

MAKERWEAR EXAMPLES

“MOTION-REACTIVE CLOTHES”





Now with fade effect

MAKERWEAR

“LASER TAG ARMBAND”



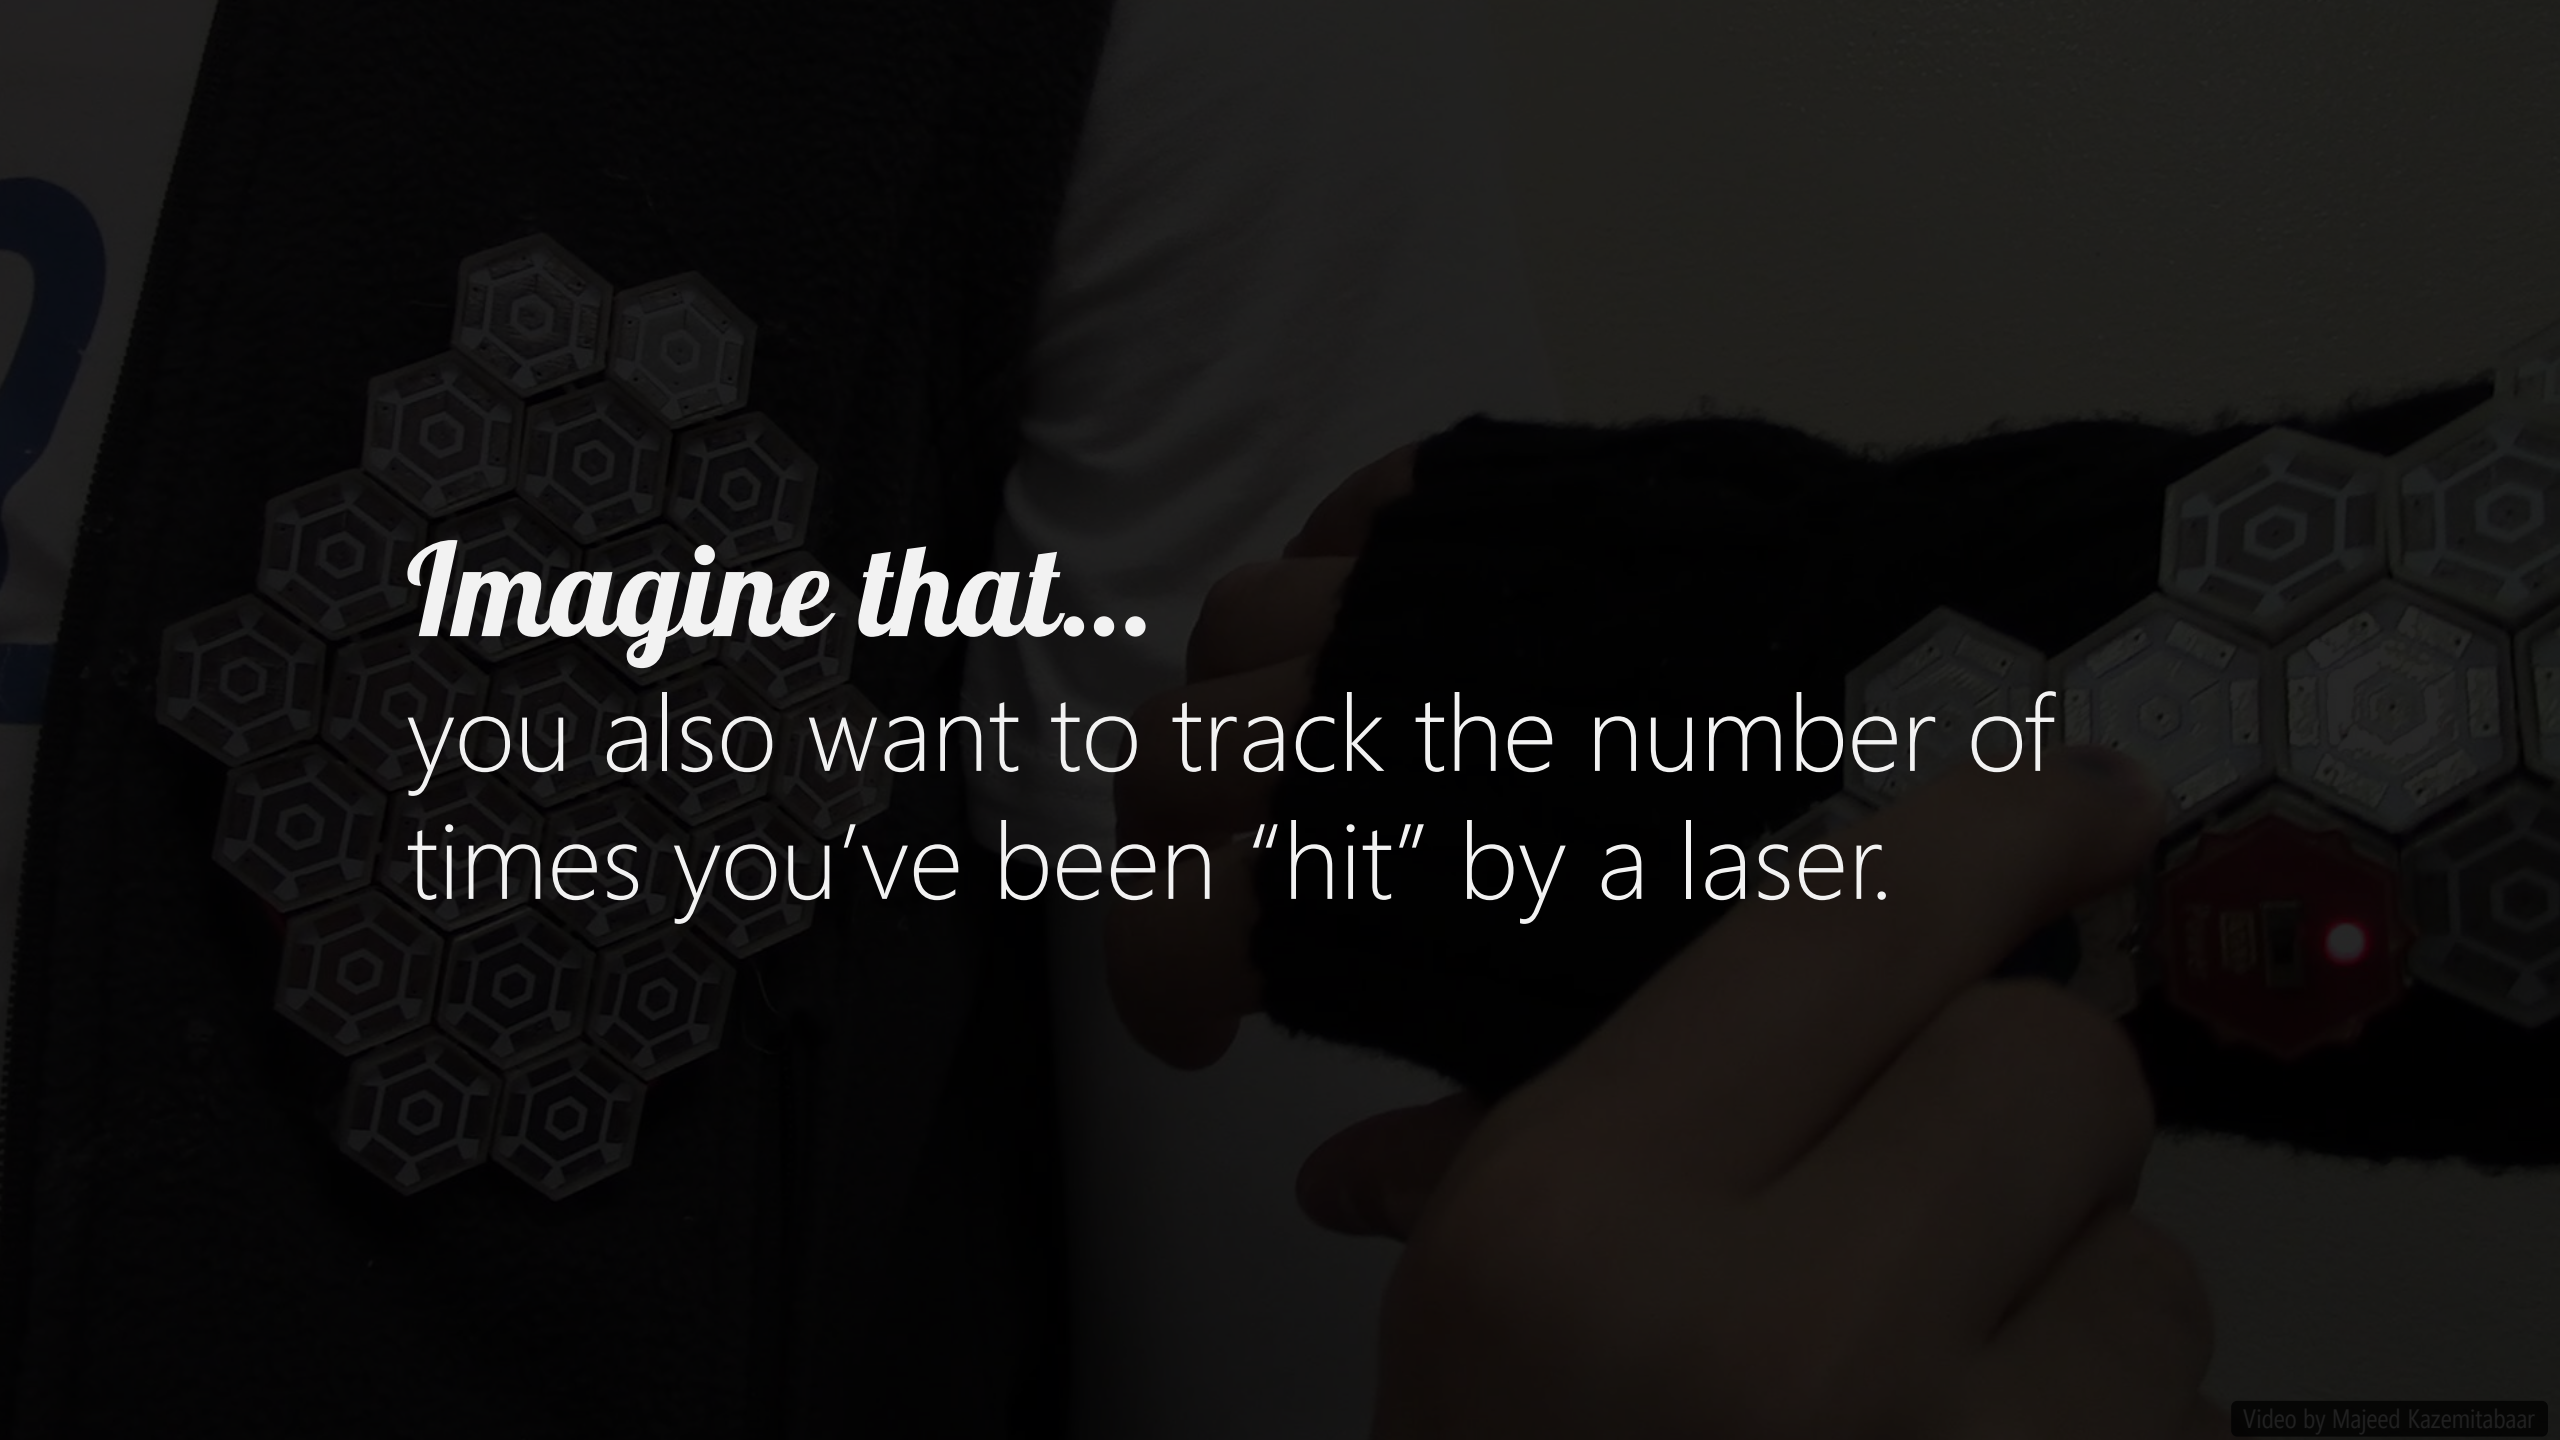
Video by Majeed Kazemitabaar

MAKERWEAR EXAMPLES

“LASER TAG ARMBAND”

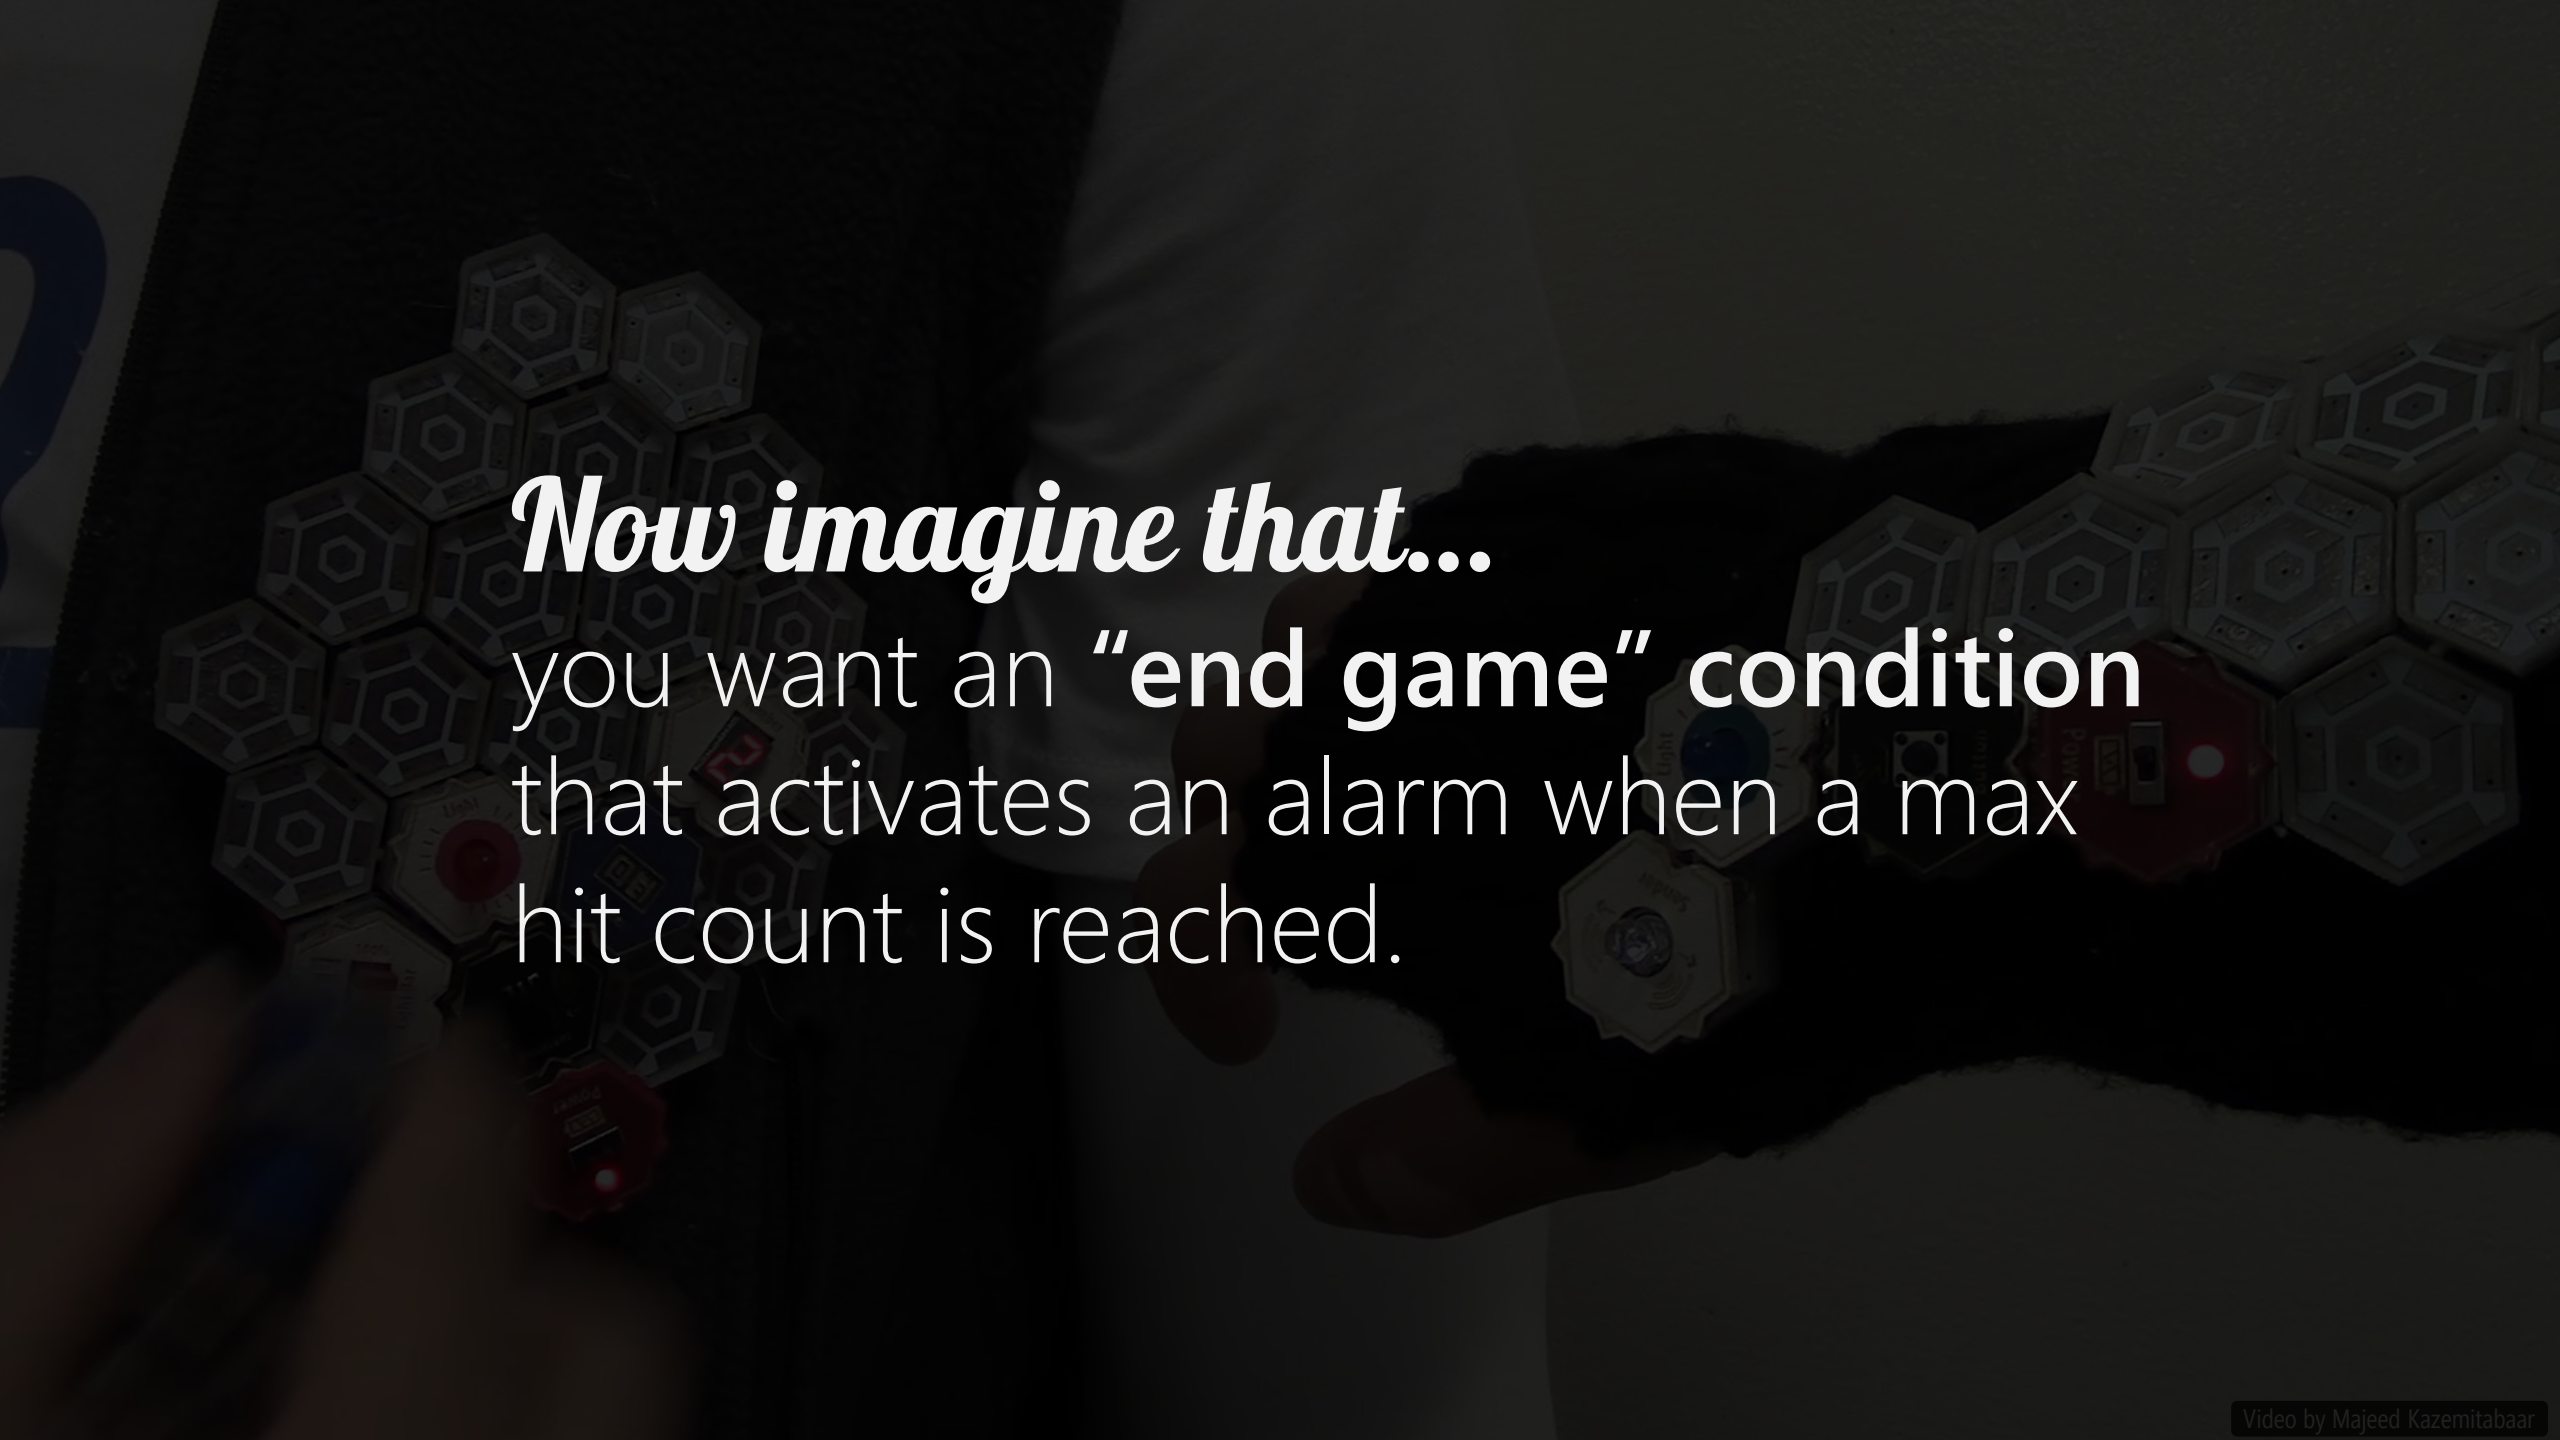


Video by Majeed Kazemitabaar


Imagine that...
you also want to track the number of times you've been "hit" by a laser.
Video by Majeed Kazemitabaar



Video by Majeed Kazemitabaar



Now imagine that...

you want an **"end game" condition** that activates an alarm when a max hit count is reached.

Video by Majeed Kazemitabaar



Video by Majeed Kazemitabaar



MAKERWEAR

OVERALL

Highly engaged in making

Wide variety of designs

Applied computational thinking

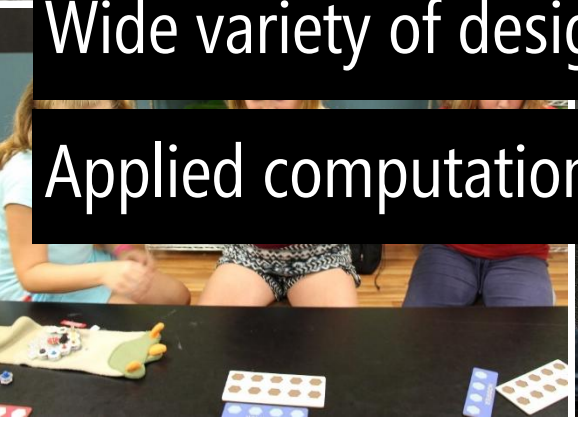


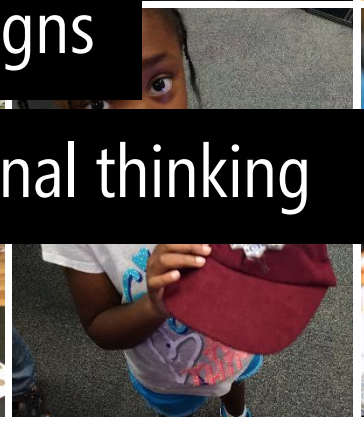



























MAKERWEAR
“MAGIC POKÉMON”
Age 9
Photo by Jon E. Froehlich

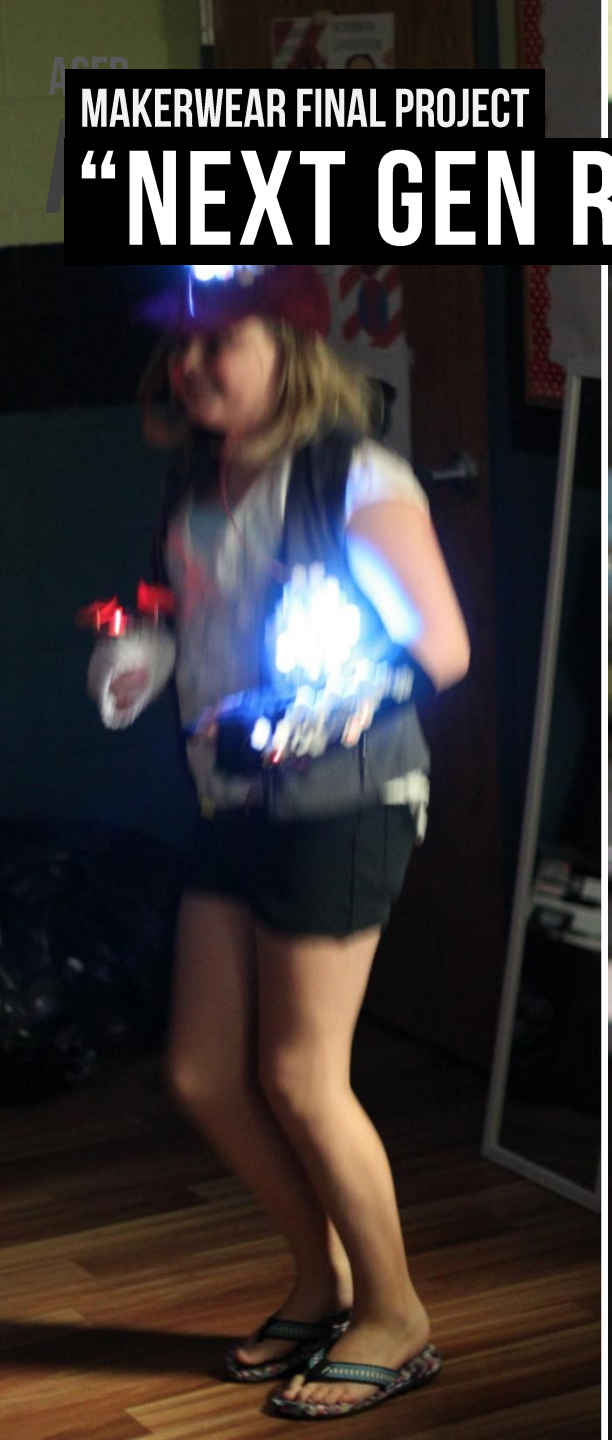

Video by Jon E. Froehlich


MAKERWEAR
“SMART LACROSSE STICK”
KEEP
GOING
GOING
GOING
Age 9


BAD
AD's
SUPERMAN
LOVES STEM
KEEP
GOING

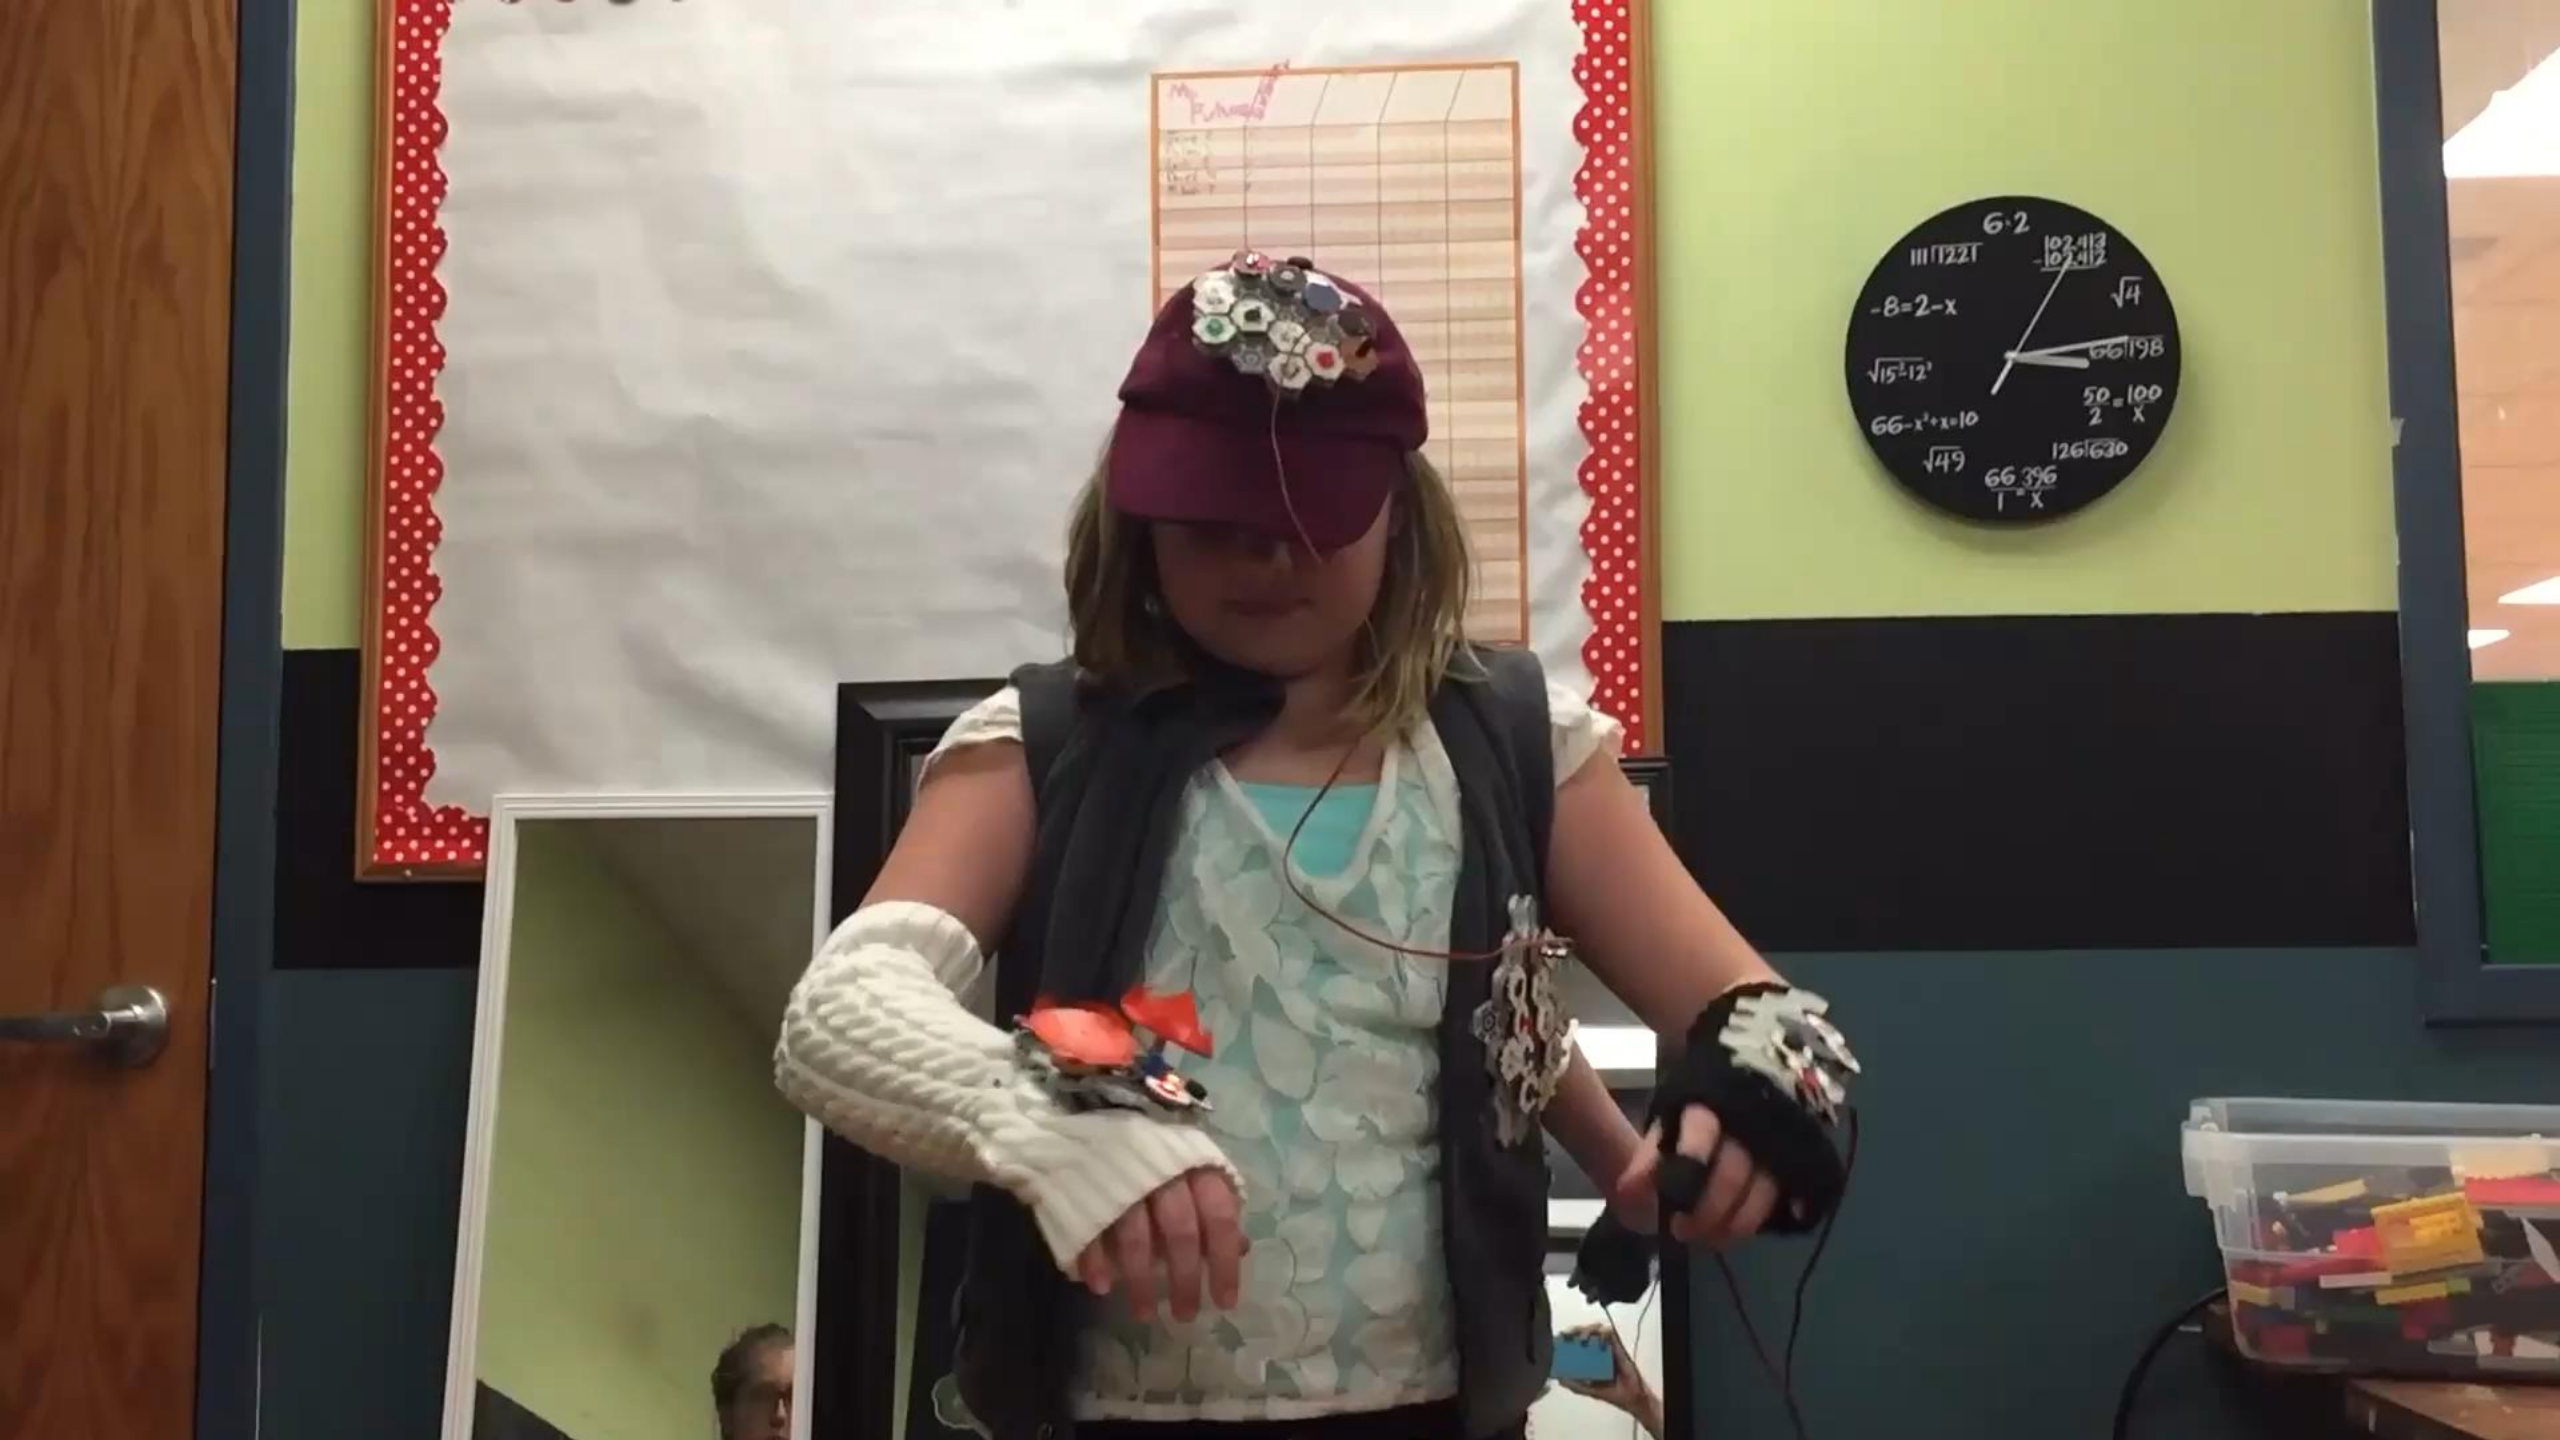
MAKERWEAR FINAL PROJECT

“NEXT GEN RUNNING CLOTHES”



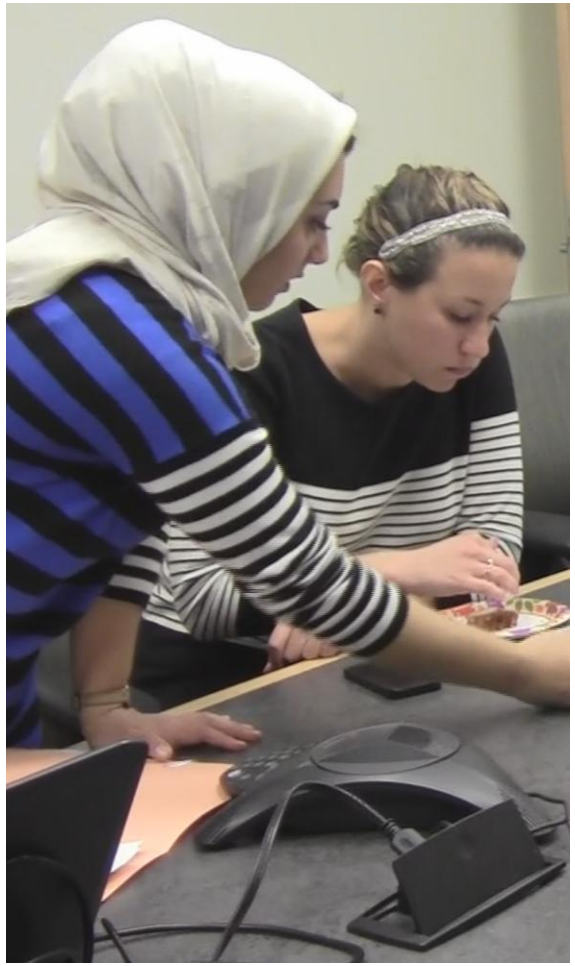


“Amelia”, Age 10


6·2
$\sqrt{4}$
$\sqrt{49}$

STEM EDUCATION

EVALUATIONS (N=240)



TEACHER INTERVIEWS



AFTER-SCHOOL PROGRAMS



CAMPS / WORKSHOPS

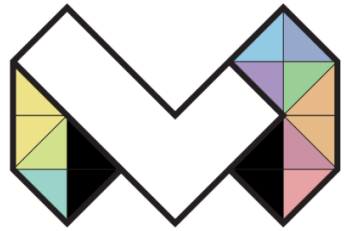


ELEMENTARY SCHOOLS

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your work
6. The **people matter** the most
7. Engage and serve **community**



Makeability Lab Handbook

By Professor Jon Froehlich and lab members
Allen School of Computer Science and Engineering
University of Washington

Welcome to the Makeability Lab! 👋 We are so happy that you're here. Let's do great things together!

This handbook is intended to establish lab principles, values, and expectations. When you join our group (woohoo!), you are expected to read it. If you're just considering the Makeability Lab, awesome! This document should help you understand *who* we are and *whether* our lab is a good match for your personality and interests.

The handbook is a working document. We value your comments, suggested edits, and constructive criticisms. As with the lab itself, this handbook is ours—we get to define who we are and what we want to be. This is both a collective opportunity and a collective responsibility. Being part of a lab, pursuing a PhD, and having a healthy advisor-advisee relationship requires significant work and investment.

If something is unclear, it's likely confusing to others too—so, please ask questions and/or suggest edits so we can collectively improve the handbook.

This handbook is a start and not an end—and hopefully will provoke self-reflection and opportunities for dialogue between us. Please also see our list of Resources, the Allen School Graduate Student Handbook, and the PhD Process webpage.

Let's do it!

Core Values

The following Makeability values represent who we are and continually strive to be. They help define and unify us.

Mutual respect and support. The Makeability Lab is founded on mutual respect, trust, and support. We respect and support one another through our words and actions. If someone makes you uncomfortable, please let me know immediately. If I am the source of discomfort, I would hope that you could approach me—either in the moment or afterwards. If you don't feel comfortable reaching out to me, you can involve our grad advising staff. Alternatively, you could fill out this anonymous Makeability Lab feedback form or a similar feedback form run by the Allen School.

Trust. Trust allows us to take risks, be vulnerable, and cooperate more efficiently[4]. It is foundational to any relationship—personal or professional. We must trust each other and our labmates. A challenge for all of us, however, is that trust is *earned* through interaction, experience, and observation. Let us invest in our relationships, be patient, and earn this trust.

Elevate others. How can you make others better as a result of your presence? In the Makeability Lab, we try to elevate others, not diminish them. Your presence in the lab should lift us higher. For example, show interest in other's work and well being, be constructive and kind in your critiques, volunteer to help review paper drafts, talks, and scholarship applications.

Community engagement. Our research is driven by the needs of people, communities, and society. Our intention should always be to *work with* and *contribute* to these communities.

Togetherness. Research and graduate school can feel insular. But we are in this together. Seek help. Offer help. Be there for one another. My most fulfilling and exciting collaborations in graduate school were with fellow students—working closely together, learning from one another, and challenging each other to do great work and be better.

Psychological safety. We engender and nurture a psychologically safe environment. All members are free to speak up, to try and fail, to be who they are. This does not mean that we must all agree or provide unequivocal praise but that people feel supported in taking risks and speaking up.

Culture of excellence. You are part of one of the top HCI research institutions in the world. In all parts of our work, we strive for excellence and hold ourselves to the highest standards, while acknowledging that, in some cases, perfection is the enemy of done.

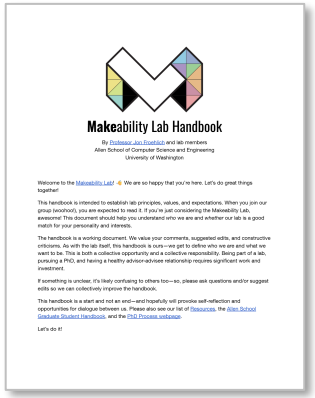
Passionate. We are passionate about what we do and are grateful for the opportunity.

Our best effort. We push ourselves and each other to do the very best that we can.

Accountability. We hold ourselves accountable. Make commitments to yourself and to your team and meet those commitments. If you struggle, seek help. Be dependable, reliable, and responsible.

---

[4] See Krot & Lewicka, IJEBM 2021



Makeability Lab Handbook

By Professor Jon Froehlich and lab members
Allen School of Computer Science and Engineering
University of Washington

Welcome to the Makeability Lab! 👋 We are so happy that you're here. Let's do great things together!

This handbook is intended to establish lab principles, values, and expectations. When you join our group (woohoo!), you are expected to read it. If you're just considering the Makeability Lab, awesome! This document should help you understand who we are and whether our lab is a good match for your personality and interests.

The handbook is a working document. We value your comments, suggested edits, and constructive criticisms. As with the lab itself, this handbook is ours—we get to define who we are and what we want to be. This is both a collective opportunity and a collective responsibility. Being part of a lab, pursuing a PhD, and having a healthy advisor-advisee relationship requires significant work and investment.

If something is unclear, it's likely confusing to others too—so, please ask questions and/or suggest edits so we can collectively improve the handbook.

This handbook is a start and not an end—and hopefully will provoke self-reflection and opportunities for dialogue between us. Please also see our list of Resources, the Allen School Graduate Student Handbook, and the PhD Process webpage.

Let's do it!

Core Values

The following Makeability values represent who we are and continually strive to be. They help define and unify us.

Mutual respect and support. The Makeability Lab is founded on mutual respect, trust, and support. We respect and support one another through our words and actions. If someone makes you uncomfortable, please let me know immediately. If I am the source of discomfort, I would hope that you could approach me—either in the moment or afterwards. If you don't feel comfortable reaching out to me, you can involve our grad advising staff. Alternatively, you could fill out this anonymous Makeability Lab feedback form or a similar feedback form run by the Allen School.

Trust. Trust allows us to take risks, be vulnerable, and cooperate more efficiently[4]. It is foundational to any relationship—personal or professional. We must trust each other and our labmates. A challenge for all of us, however, is that trust is *earned* through interaction, experience, and observation. Let us invest in our relationships, be patient, and earn this trust.

Elevate others. How can you make others better as a result of your presence? In the Makeability Lab, we try to elevate others, not diminish them. Your presence in the lab should lift us higher. For example, show interest in other's work and well being, be constructive and kind in your critiques, volunteer to help review paper drafts, talks, and scholarship applications.

Community engagement. Our research is driven by the needs of people, communities, and society. Our intention should always be to *work with* and *contribute* to these communities.

Togetherness. Research and graduate school can feel insular. But we are in this together. Seek help. Offer help. Be there for one another. My most fulfilling and exciting collaborations in graduate school were with fellow students—working closely together, learning from one another, and challenging each other to do great work and be better.

Psychological safety. We engender and nurture a psychologically safe environment. All members are free to speak up, to try and fail, to be who they are. This does not mean that we must all agree or provide unequivocal praise but that people feel supported in taking risks and speaking up.

Culture of excellence. You are part of one of the top HCI research institutions in the world. In all parts of our work, we strive for excellence and hold ourselves to the highest standards, while acknowledging that, in some cases, perfection is the enemy of done.

Passionate. We are passionate about what we do and are grateful for the opportunity.

Our best effort. We push ourselves and each other to do the very best that we can.

Accountability. We hold ourselves accountable. Make commitments to yourself and to your team and meet those commitments. If you struggle, seek help. Be dependable, reliable, and responsible.

[4] See Krot & Lewicka, IJEBM 2021

"**Community engagement.** Our research is driven by the needs of people, communities, and society. Our intention should always be to *work with* and *contribute* to these communities. "

STEM EDUCATION

PARTNERSHIPS

KID MUSEUM

BOYS & GIRLS CLUBS



Center for Young Children

futuremakers

KIDS MAKE THINGS BETTER.COM

PGCPS

Great By Choice

Prince George's County
Public School System

WE ARE CASA

CASA

SOMOS CASA



University of Maryland Kidsteam

UNIVERSITY OF MARYLAND 1856

COLLEGE OF EDUCATION

STEM Masters in Education Program











SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, in your vision
6. The **people matter** the most
7. Engage and serve **community**



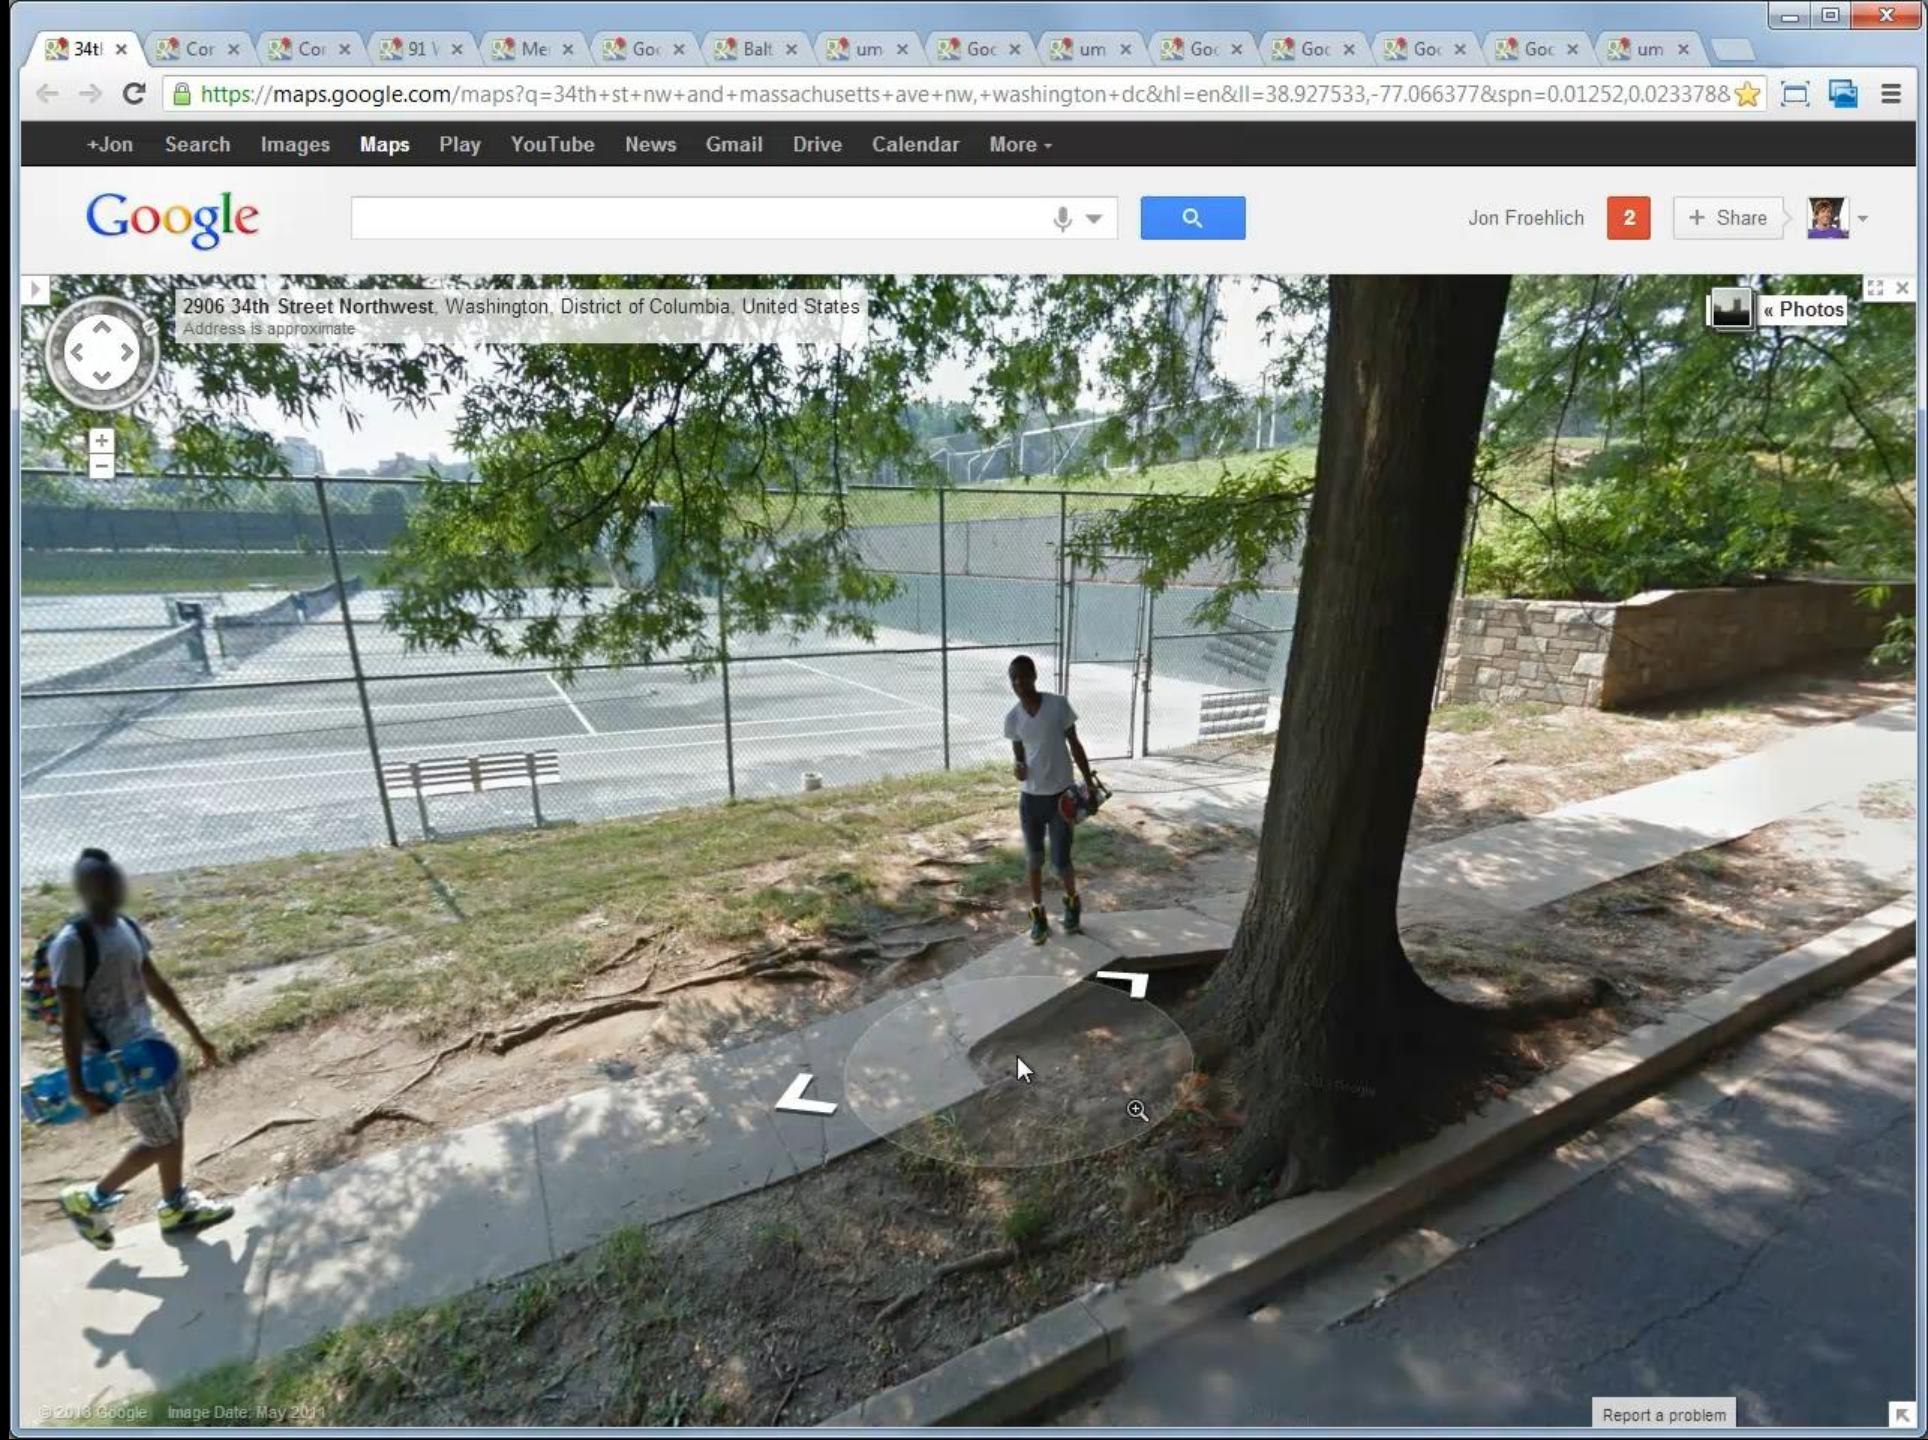
UMD 2012

STARTING AS A PROFESSOR



UMD 2012

CAN WE USE STREET VIEW TO FIND & MAP INACCESSIBLE AREAS OF THE WORLD?



How do **sidewalk barriers** impact **human mobility**?

INTERVIEWS & WORKSHOPS WITH 150 STAKEHOLDERS



Combining Crowdsourcing and Google Street View to Identify Street-level Accessibility Problems

Kotaro Hara, Victoria Le, Jon E. Froehlich
Human-Computer Interaction Lab (HCIL)
Computer Science Department, University of Maryland, College Park

Improving Public Transit Accessibility for Blind Riders by Crowdsourcing Bus Stop Landmark Locations with Google Street View

Kotaro Hara[1], Shiri Azenkot[2], Megan Campbell[2], Cynthia L. Bennett[2], Vicki Le[1], Sean Pannella[1], Robert Moore[1], Kelly Minckler[2], Rochelle H. Ng[2], Jon E. Froehlich[1]
[1]Makeability Lab | HCIL
[2]DUB Group
Computer Science and Engineering

The Design of Assistive Location-based Technologies for People with Ambulatory Disabilities: A Formative Study

Kotaro Hara, Christine Chan, Jon E. Froehlich
Makeability Lab | Human-Computer Interaction Lab (HCIL)
Computer Science Department, University of Maryland, College Park
{kotaro, jonf}@cs.umd.edu; chan2017@umd.edu

ABSTRACT
In this paper, we investigate how people with mobility impairments assess and evaluate accessibility in the built environment and the role of current and emerging location-based technologies therein. We conducted a three-part formative study with 20 mobility impaired participants: a semi-structured interview (Part 1), a design activity (Part 2), and a design [illegible] and 3 actively engaged our [illegible] designing the future of what *based technologies* (ALTs)—that specifically incorporate a[illegible] navigating, searching, and [illegible] Our Part 1 findings highlight provide accessibility benef[illegible] explicitly designed for such help identify and uncover us[illegible] particular, we synthesize 1[illegible] qualities. We conclude with

Author Keywords
Accessibility; assistive te[illegible] location-based technology

ACM Classification Keyw[illegible]

Figure 1. To explore how location-based technologies currently support users with mobility impairments as well as to examine desired future [illegible] we conducted a three-part formative study with 20 [illegible] photos from (a) a semi-structured

"never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk

Chu Li*
University of Washington, USA
chuchuli@cs.washington.edu

Katrina Ma*
University of Washington, USA
ma.katrina91@gmail.com

Michael Saugstad
University of Washington, USA
saugstad@cs.washington.edu

Kie Fujii
[illegible]ckensack Meridian School of Medicine, USA
kie.fujii@hmhn.org

Molly Delaney
University of Illinois at Chicago, USA
mdelan9@uic.edu

Yochai Eisenberg
University of Illinois at Chicago, USA
yeisen2@uic.edu

Delphine Labbé
[illegible]ity of Illinois at Chicago, USA
dlabbe@uic.edu

Judy L. Shanley
Easterseals, USA
jshanley@easterseals.com

Devon Snyder
University of Illinois at Chicago, USA
dsnyde8@uic.edu

Urban Accessibility as a Socio-Political Problem: A Multi-Stakeholder Analysis

MANASWI SAHA, University of Washington
DEVANSHI CHAUHAN, University of Washington
SIDDHANT PATIL, [illegible]

disabilities

MDPI

Article

Multiple-Stakeholder Perspectives on Accessibility Data and the Use of Socio-Technical Tools to Improve Sidewalk Accessibility

Delphine Labbé [1], Yochai Eisenberg [1,*], Devon Snyder [1], Judy Shanley [2], Joy M. Hammel [1,3] and Jon E. Froehlich [4]

[1] Department of Disability and Human Development, University of Illinois-Chicago, Chicago, IL 60608, USA; dlabbe@uic.edu (D.L.); dsnyde8@uic.edu (D.S.); hammel@uic.edu (J.M.H.)
[2] Easterseals, Chicago, IL 60604, USA; jshanley@easterseals.com
[3] Occupational Therapy Department, University of Illinois-Chicago, Chicago, IL 60608, USA
[4] Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA 98195, USA; jonf@cs.washington.edu
* Correspondence: yeisen2@uic.edu

Abstract: For people with disabilities, accessible pedestrian infrastructure can support independence,



Source: Rooted in Rights, https://youtu.be/NRh-4Y3aNeo



Source: Disability Rights Washington, https://youtu.be/LBGDXMTyib4


THE PROBLEM IS
NOT JUST A LACK
OF ACCESSIBLE
INFRASTRUCTURE
A LACK OF DATA

The National Council on Disability notes that there is **no comprehensive information** on "the degree to which sidewalks are accessible" in US cities.

National Council on Disability
The impact of the Americans with Disabilities Act: Assessing the progress toward achieving the goals of the ADA

Google Faculty Research Award, 2012

Combining Crowdsourcing and Computer Vision for Street-level Accessibility

PI: Jon Froehlich (HCI), Co-PI: David Jacobs (Computer Vision)
Department of Computer Science, University of Maryland, College Park
{ jonf, djacobs }@cs.umd.edu
Google Sponsors: Peng Dai (daipeng@google.com) and Sacha Arnoud (sacha@google.com)



(a) Object in Path (b) No Curb Ramp (c) Surface Problem (d) Prematurely Ending Sidewalk

Figure 1: In this proposal, we combine crowdsourcing and computer vision to semi-automatically locate, identify, and assess accessibility problems in city streets, sidewalks, and building facades in online map imagery. The images above show crowd annotations from preliminary experiments on Mechanical Turk where untrained "turkers" were asked to find, label, and rate the severity of accessibility obstacles in Google Street View (GSV) images.

1 Introduction

According to the most recent US Census (2010), roughly 30.6 million adults have physical disabilities that affect their ambulatory activities [8]. 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) [8]. Despite comprehensive civil rights legislation for Americans with disabilities (*e.g.,* [1]), many city streets, sidewalks, and businesses in the US remain inaccessible [6,7]. 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 *a priori*. Indeed, in a recent report, the National Council on Disability noted that they could not find comprehensive information on the "degree to which sidewalks are accessible" across the US [6]. **In this proposal, we describe new, scalable methods for collecting data on street-level accessibility including streets, sidewalks, and building facades by combining crowdsourcing and computer vision (*e.g.,* Figure 1).** Our overarching vision is to transform the way information about street-level accessibility is collected, visualized, and accessed.

Google Faculty Research Award, 2012

Combining Crowdsourcing and Computer Vision for Street-level Accessibility

PI: Jon Froehlich (HCI), Co-PI: David Jacobs (Computer Vision)
Department of Computer Science, University of Maryland, College Park
{ jonf, djacobs }@cs.umd.edu
Google Sponsors: Peng Dai (daipeng@google.com) and Sacha Arnoud (sacha@google.com)



(a) Object in Path (b) No Curb Ramp (c) Surface Problem (d) Prematurely Ending Sidewalk

Figure 1: In this proposal, we combine crowdsourcing and computer vision to semi-automatically locate, identify, and assess accessibility problems in city streets, sidewalks, and building facades in online map imagery. The images above show crowd annotations from preliminary experiments on Mechanical Turk where untrained "turkers" were asked to find, label, and rate the severity of accessibility obstacles in Google Street View (GSV) images.

1 Introduction

According to the most recent US Census (2010), roughly 30.6 million adults have physical disabilities that affect their ambulatory activities [8]. 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) [8]. Despite comprehensive civil rights legislation for Americans with disabilities (*e.g.,* [1]), many city streets, sidewalks, and businesses in the US remain inaccessible [6,7]. 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 *a priori.* Indeed, in a recent report, the National Council on Disability noted that they could not find comprehensive information on the "degree to which sidewalks are accessible" across the US [6]. **In this proposal, we describe new, scalable methods for collecting data on street-level accessibility including streets, sidewalks, and building facades by combining crowdsourcing and computer vision (*e.g.,* Figure 1).** Our overarching vision is to transform the way information about street-level accessibility is collected, visualized, and accessed.

To support the viability of our proposal, we have conducted three preliminary studies. First, because finding, labeling, and assessing the severity of sidewalk accessibility problems in streetscape imagery is a subjective and potentially ambiguous task, we investigated its feasibility with six motivated and diligent labelers: three members of the research team and three wheelchair users. Results showed the labeling approach is reliable, with high intra- and inter-labeler agreement, and allowed us to establish baseline understandings of performance (*i.e.,* what does "good" labeling performance look like for this task) [5]. Second, we piloted a range of labeling interfaces and examined the effect of three interfaces on task performance [2]. Third, we performed preliminary experiments using a manually curated database of 229 Google Street View (GSV) images, demonstrating that *untrained* crowd workers on Mechanical Turk (turkers) could correctly determine the presence of accessibility problems with 81% accuracy. With simple quality control methods, this accuracy increased to 93% [5].

2 Research Plan

Building off this initial momentum, our proposal describes two phases of work: **(i) the development of scalable, map-based techniques and labeling tools** to acquire large amounts of data on street-level accessibility, which will be deployed initially on online labor markets such as Mechanical Turk; **(ii) the integration and use of computer vision algorithms** to further scale our data collection methods by directing the efforts of human labelers to scenes that are likely to contain problems. Though beyond the scope of this proposal, we ultimately plan to deploy these data collection techniques more broadly via a publically accessible website, and to utilize the data and techniques to enable new types of accessible-aware navigation tools.; For example, *RouteAssist* (Figure 2) will personalize route suggestions based on a user's reported mobility level.

2.1 Crowdsourcing Streetscape Labeling

Our labeling work will focus on interfaces and experiments for practical, scalable collection of street-level accessibility data.

Improving labeling scalability: Our current prototype labeling system relies on a manually curated database of images selected by the research team. This was sufficient to demonstrate feasibility but ignored important practical aspects such as locating the GSV camera in geographic space and selecting an optimal viewpoint. These challenges clearly need to be solved

NSF Medium Grant, 2012

HCC: Medium: Combining Crowdsourcing and Computer Vision for Street-level Accessibility



(a) Object in Path (b) No Curb Ramp (c) Surface Problem (d) Prematurely Ending Sidewalk

Figure 1: In this proposal, we outline a research agenda that combines crowdsourcing, online map imagery, and computer vision to semi-automatically locate, identify, and assess accessibility problems in city streets, sidewalks, and building facades. The images above show crowd annotations from preliminary experiments on Mechanical Turk where untrained "turkers" were asked to find, label, and rate the severity of sidewalk accessibility obstacles in Google Street View (GSV) images.

1 Introduction

According to the most recent US Census (2010), roughly 30.6 million adults have physical disabilities that affect their ambulatory activities [128]. 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) [128]. Despite comprehensive civil rights legislation for Americans with disabilities (e.g., [9,75]), many city streets, sidewalks, and businesses in the US remain inaccessible [90,96,120]. 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 *a priori.* Indeed, in a recent report, the National Council on Disability noted that they could not find comprehensive information on the "degree to which sidewalks are accessible" across the US [96]. In this proposal, we describe new scalable methods for collecting data on street-level accessibility including streets, sidewalks, and building facades (*e.g.,* Figure 1) and a set of new accessible-aware map-based tools to take advantage of this data (*e.g.,* Figure 2). Our overarching vision is to transform the way in which information is collected and visualized about street-level accessibility.

The lack of street-level accessibility information can have a significant negative impact on the independence and mobility of citizens [99,120]. For example, in our own initial 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). Although maintenance issues such as buckled or cracked sidewalks can pose significant accessibility challenges so too do larger, more permanent infrastructural issues such as utility poles or fire hydrants directly in sidewalk paths or the lack of curb ramps at intersections or sidewalks (see Figure 1). These issues are significant. In a precedent-setting 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 that "without curb cuts, people with ambulatory disabilities simply cannot navigate the city" [75].

Sidewalk and street accessibility problems become even more significant in light of the unprecedented growth of the senior (65+) population in the US [97] and, commensurately, the number of people in the US with some level of impairment that affects their mobility. According to the 2010 US Census, nearly *half* of the US population with some form of ambulatory limitation is 65 or older (15.2 million) [128]. Because the attenuated reflexes and physical limitations of older adults might prohibit them from driving automobiles, they are more likely to rely on public transit or walking than other adults [57]. In this way, inaccessible public rights-of-way have cascading effects—they do not only impede trips to local areas but can also prevent access to and from public transit stops (*e.g.,* subway stations), further restricting a person's radius of travel [128].

“This is a good idea…”

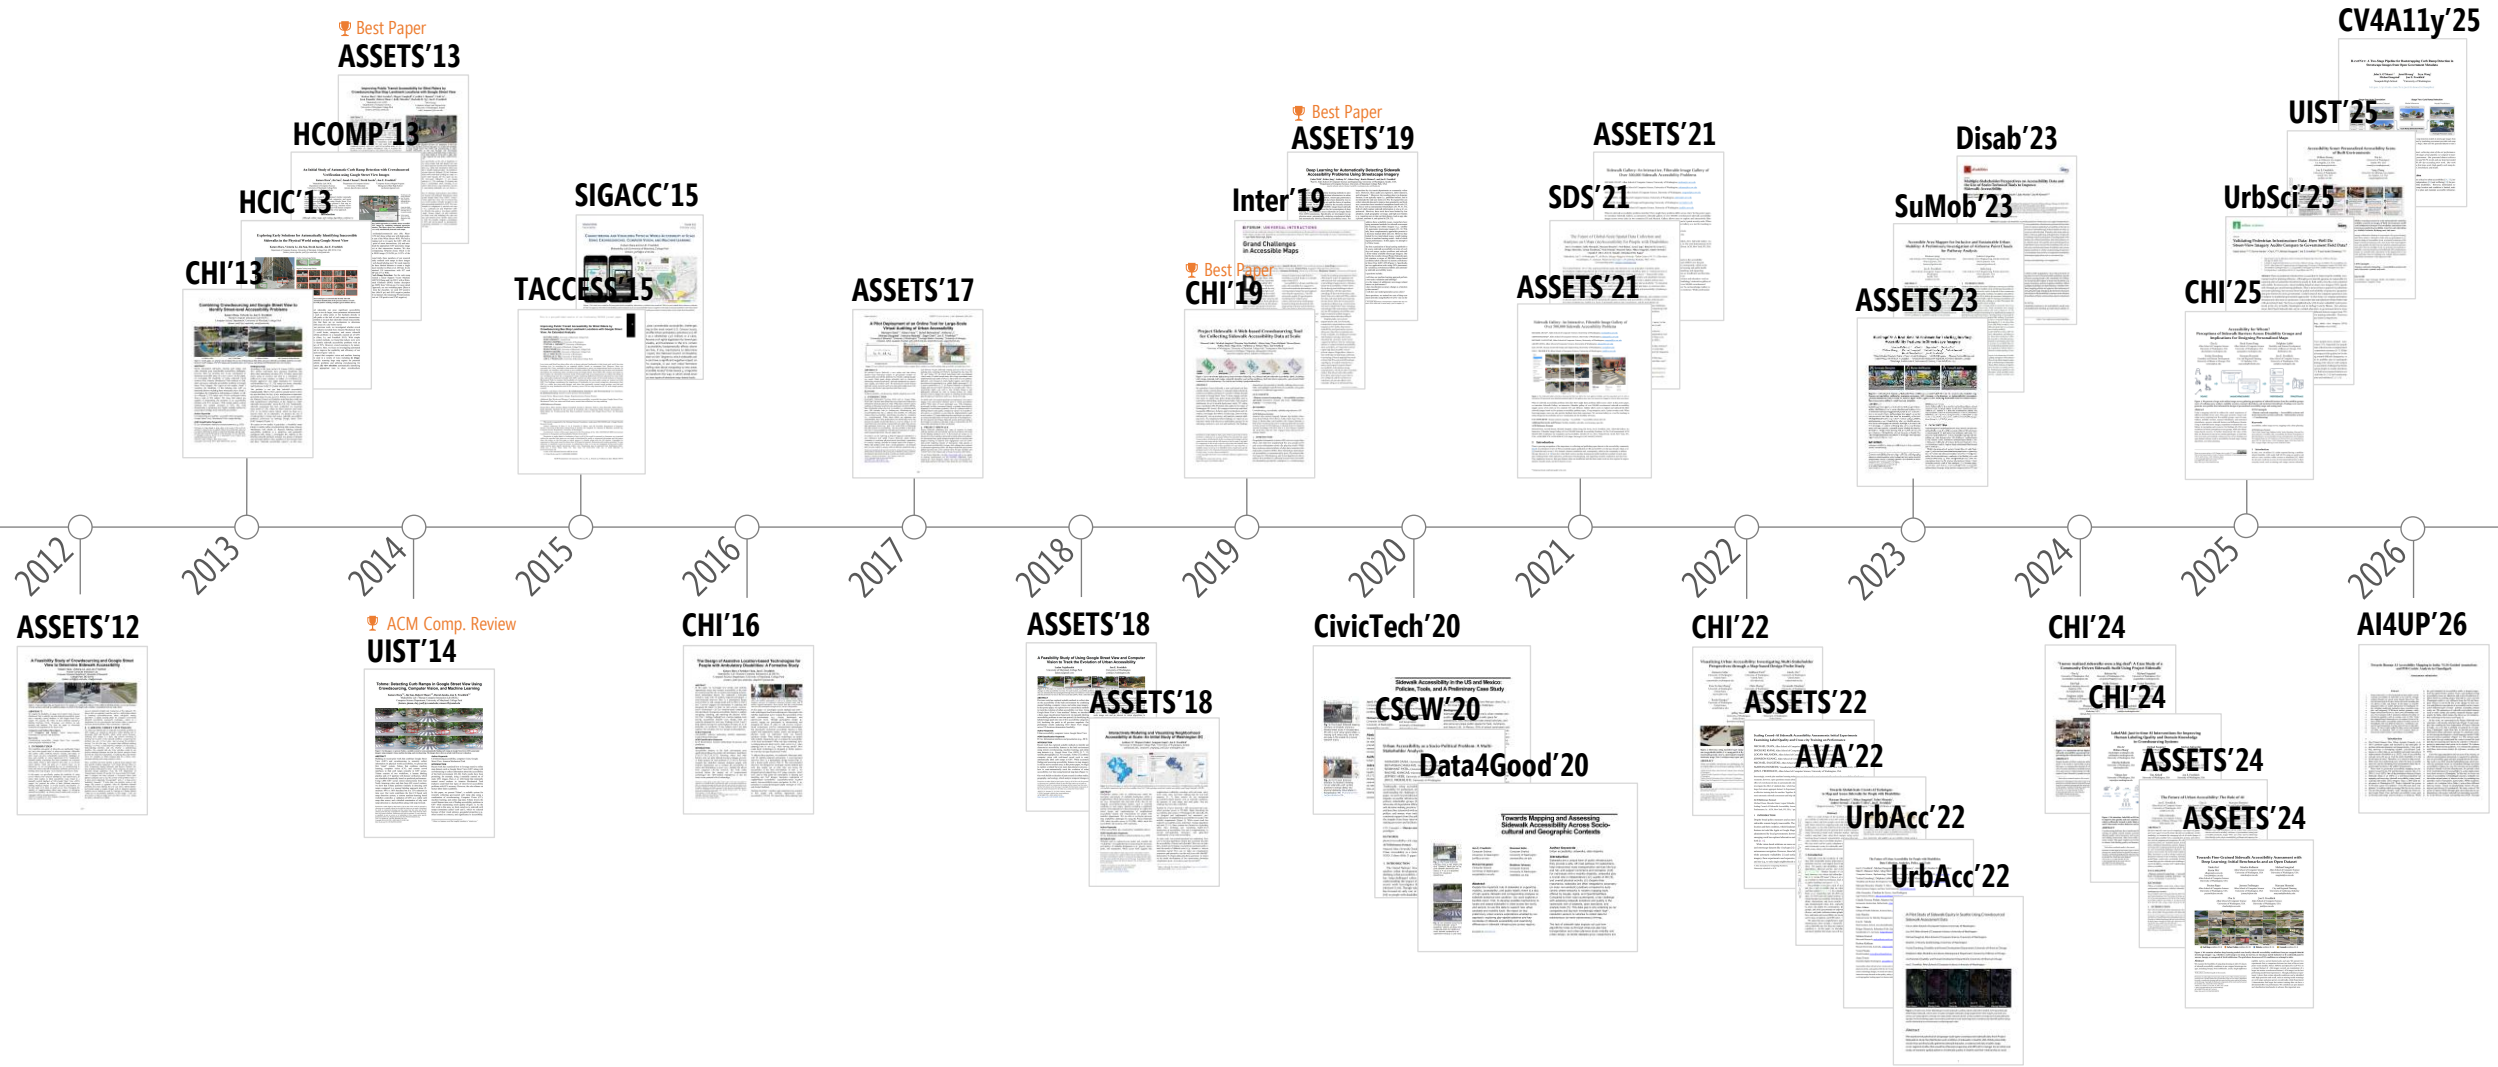
-SENIOR FACULTY

“This is a good idea but **not** a tenure-level idea.”

-SENIOR FACULTY

URBAN ACCESSIBILITY

RESEARCH



URBAN ACCESSIBILLITY

FOUR INTERNATIONAL WORKSHOPS



Gratitude! 🙏 Thank you attendees and presenters for the thoughtful discussions around the future of urban accessibility. Here are the workshop slides with links to the Miro board.

Update: join us for UrbanAccess'24 in October 2024!

Welcome 👋,

This is the webpage for our workshop entitled *"The Future of Urban Accessibility for People with Disabilities: Data Collection, Analytics, Policy, and Tools"*, which will be held virtually via **Zoom on Mon, Oct 17th** as part of the ASSETS'22 conference. You can read more about all five of the accepted ASSETS'22



🍂 Paper notifications are out. Congrats to all accepted authors! 🎉 🥳

The workshop will be held virtually over Zoom on Oct 24th starting at 7AM Pacific (GMT -7) / 3PM London (GMT +1). See schedule. And don't forget to register! All attendees/speakers must register.

ABOUT

The Future of Urban Accessibility Workshop: The Role of AI.

TRY IT!

PROJECT SIDEWALK

http://projectsidewalk.org



Kotaro Hara
UMD CS PhD'16



Manaswi Saha
UW CSE PhD'22



Mikey Saugstad
UMD CS MS'18
UW CSE



Chu Li
UW CSE PhD



Yochai Eisenberg
UIC Professor
Disability & Human Development

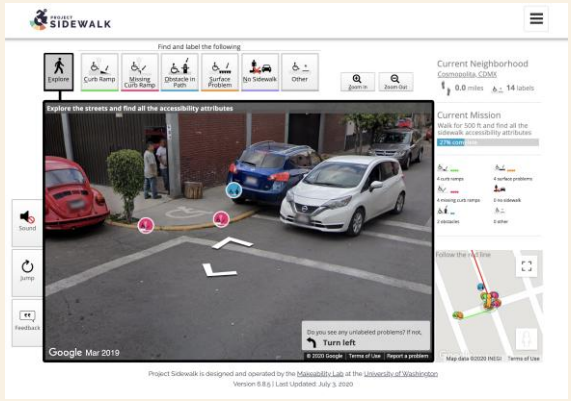
ONLINE MAP IMAGERY



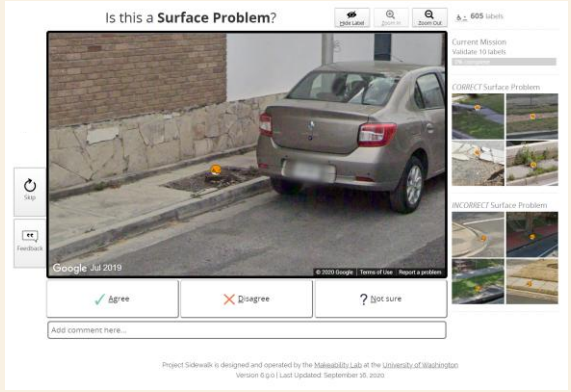




REMOTE CROWDSOURCING INTERFACES



Labeling missions



Validation missions

HUMAN LABELS

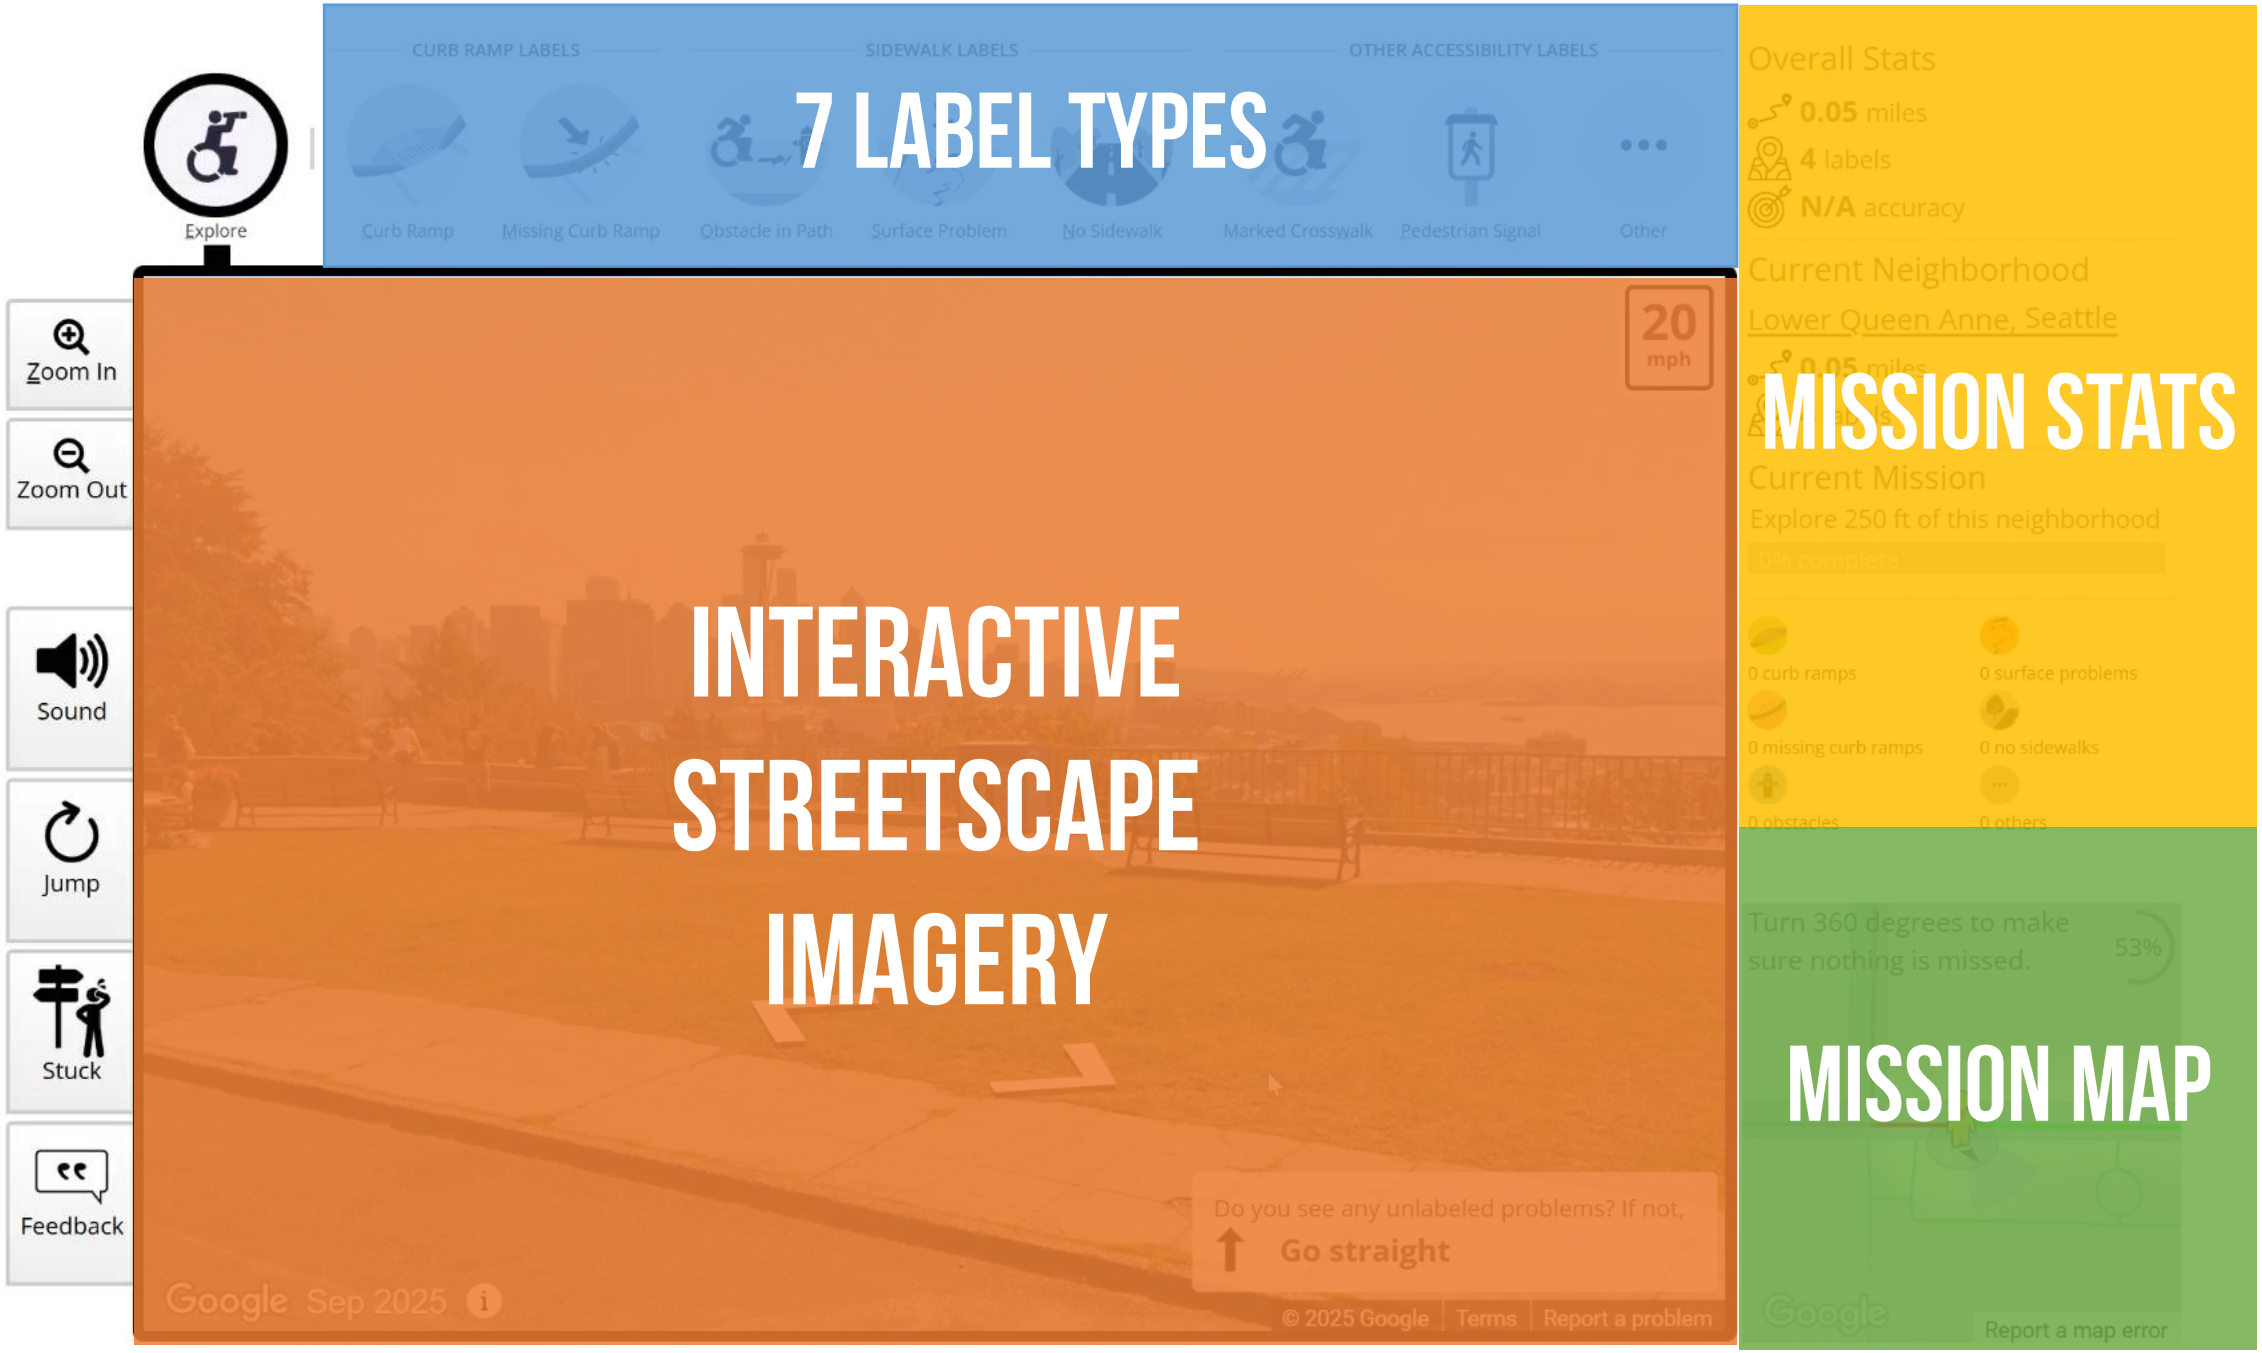
MACHINE LEARNING



PROJECT SIDEWALK
HTTP://PROJECTSIDEWALK.IO

MAP, ASSESS, & VISUALIZE EVERY SIDEWALK IN THE WORLD

PROJECTSIDEWALK.ORG


Explore
CURB RAMP LABELS
SIDEWALK LABELS
OTHER ACCESSIBILITY LABELS
Curb Ramp
Missing Curb Ramp
Obstacle in Path
Surface Problem
No Sidewalk
Marked Crosswalk
Pedestrian Signal
Other
7 LABEL TYPES
Zoom In
Zoom Out
Sound
Jump
Stuck
Feedback
20 mph
INTERACTIVE STREETSCAPE IMAGERY
Do you see any unlabeled problems? If not,
Go straight
Google Sep 2025
© 2025 Google
Terms
Report a problem
Overall Stats
0.05 miles
4 labels
N/A accuracy
Current Neighborhood
Lower Queen Anne, Seattle
MISSION STATS
Current Mission
Explore 250 ft of this neighborhood
0 curb ramps
0 surface problems
0 missing curb ramps
0 no sidewalks
0 obstacles
0 others
Turn 360 degrees to make sure nothing is missed.
53%
MISSION MAP
Google
Report a map error

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

ASSETS'12



CHI'13



HCIC'13



HCOMP'13



Best Paper

ASSETS'13



ACM Comp. Review

UIST'14



SIGACC'15

TACCESS'15

ASSETS'17

ASSETS'18

CHI'16

Best Paper

CHI'19

Best Paper

ASSETS'19

Inter'19

CivicTech'20

CSCW'20

Data4Good'20

ASSETS'21

SDS'21

ASSETS'21

CHI'22

ASSETS'22

AVA'22

UrbAcc'22

UrbAcc'22

Disab'23

SuMob'23

ASSETS'23

CHI'24

CHI'24

ASSETS'24

ASSETS'24

CV4A11y'25

UIST'25

UrbSci'25

CHI'25

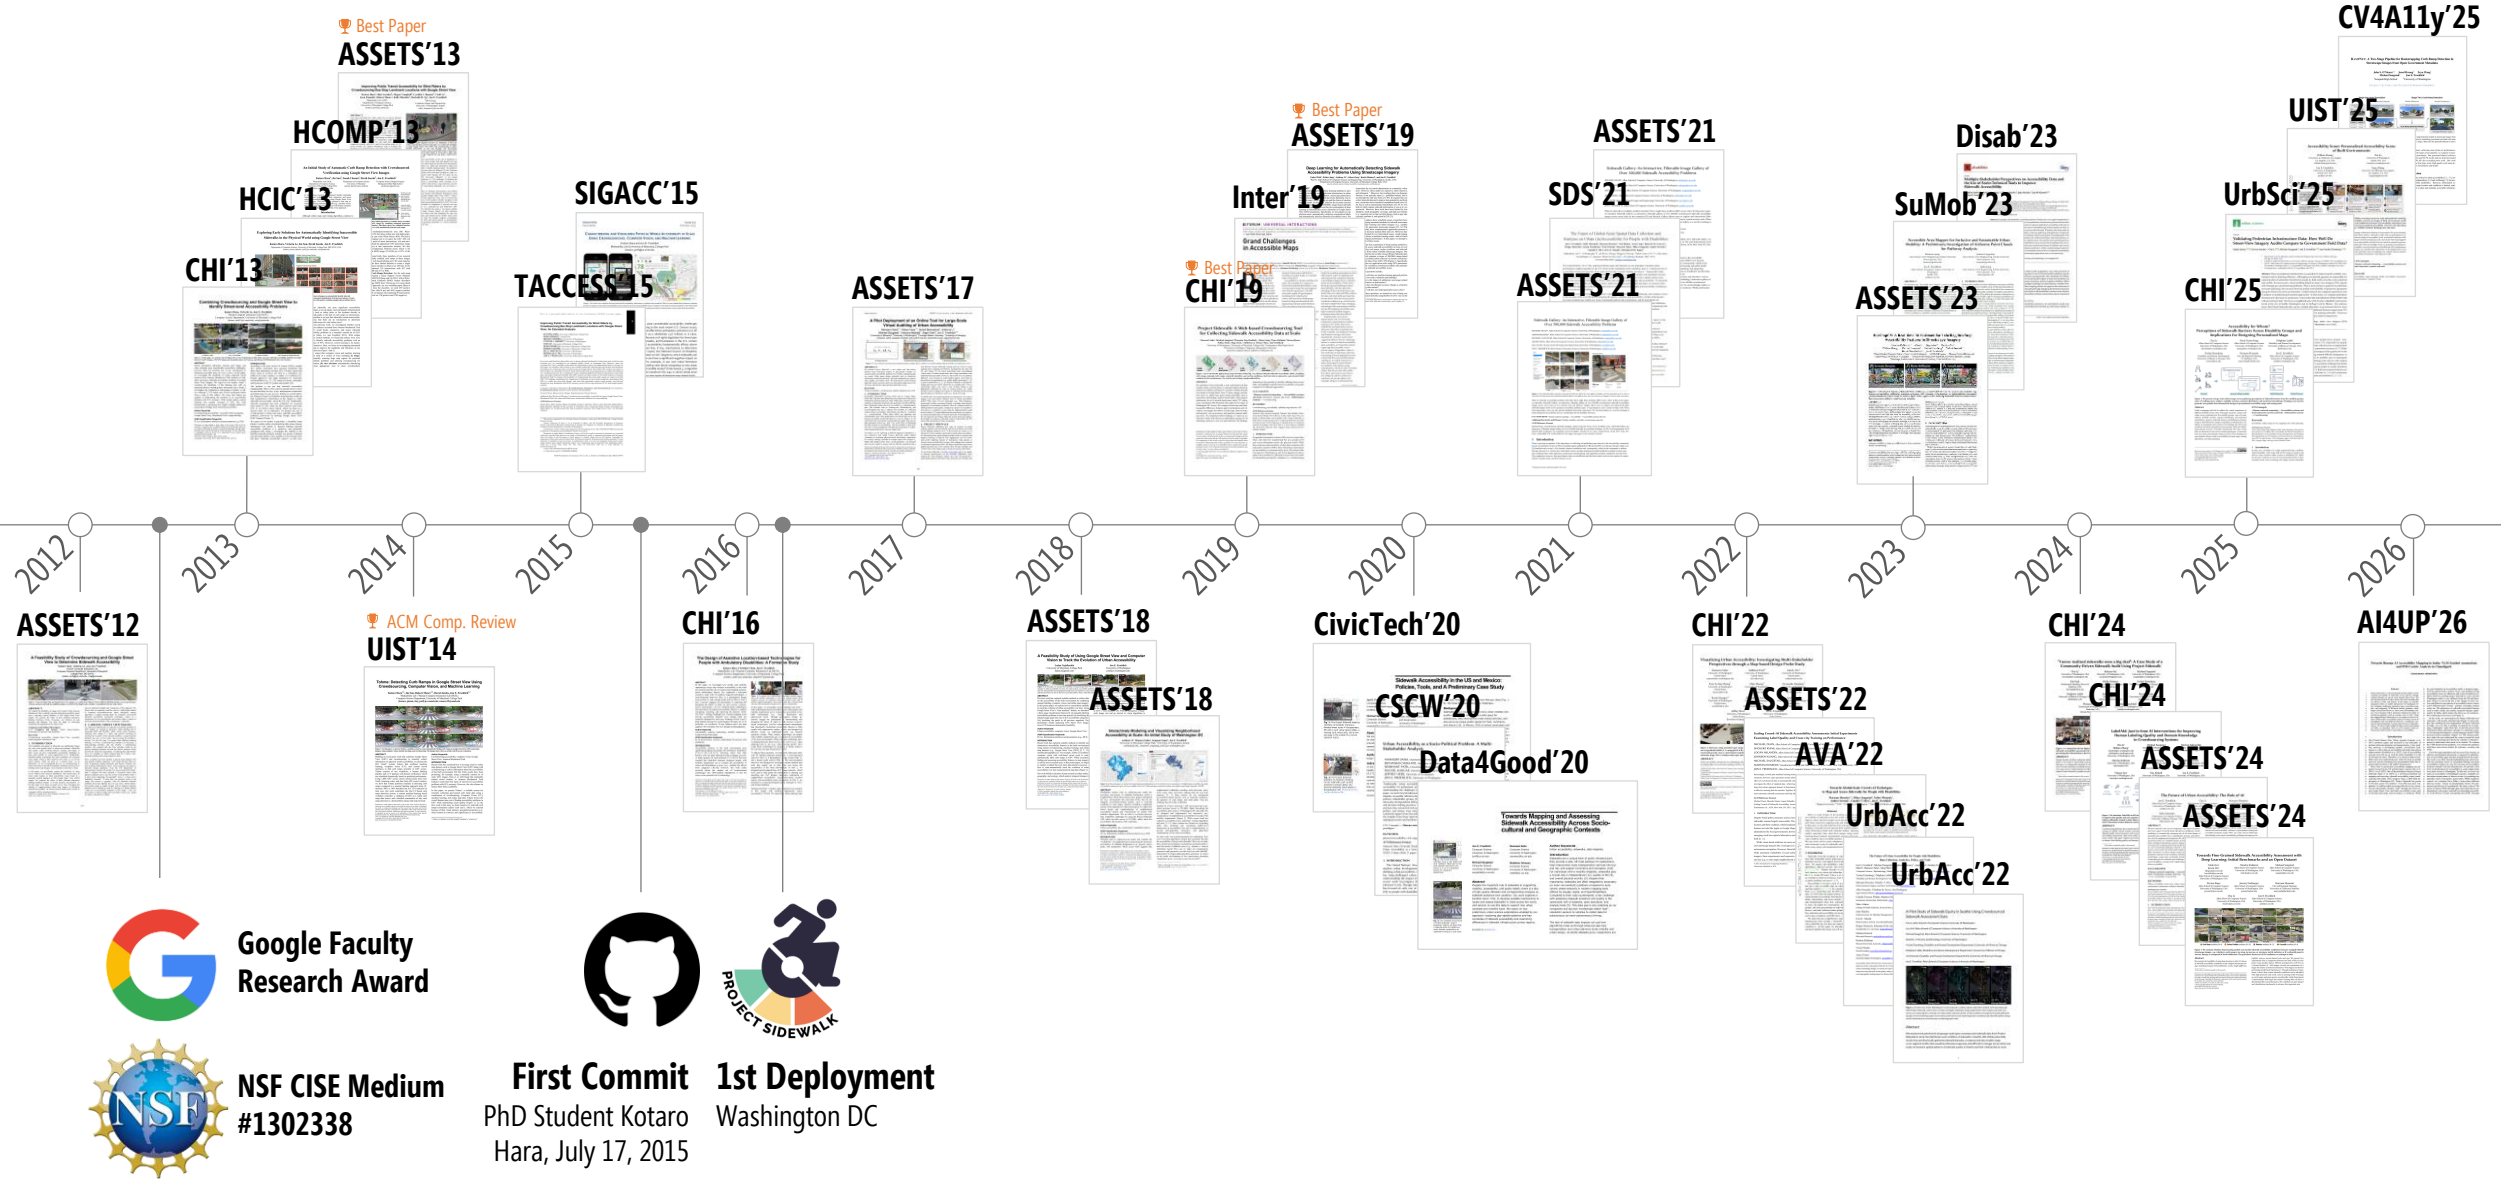
AI4UP'26

Google Faculty Research Award

NSF CISE Medium #1302338

“It’s good research but **don’t spend your time** on **open sourcing** and **deploying**. It’s not worth the effort & **won’t be** recognized by academia.”

-SENIOR FACULTY


Best Paper
ASSETS'13
HCOMP'13
HCIC'13
CHI'13
SIGACC'15
TACCESS'15
ASSETS'17
Best Paper
CHI'19
Best Paper
ASSETS'19
Inter'19
ASSETS'21
SDS'21
ASSETS'21
Disab'23
SuMob'23
ASSETS'23
CV4A11y'25
UIST'25
UrbSci'25
CHI'25
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
ASSETS'12
ACM Comp. Review
UIST'14
CHI'16
ASSETS'18
ASSETS'18
CivicTech'20
CSCW'20
Data4Good'20
CHI'22
ASSETS'22
AVA'22
UrbAcc'22
UrbAcc'22
CHI'24
CHI'24
ASSETS'24
ASSETS'24
AI4UP'26
Google Faculty Research Award
NSF CISE Medium #1302338
First Commit
PhD Student Kotaro Hara, July 17, 2015
1st Deployment
Washington DC

PROJECT SIDEWALK

55 CITIES, 12 COUNTRIES, 7 LANGUAGES



Deployments	Streets Explored	Labels	Validations
55 Cities	31,004 km	1.7 million	2.3 million

路緣坡道標籤 | 人行道標籤 | 其他可及性標籤

探索E | 路緣斜坡C | 無路緣斜坡M | 障礙物O | 鋪面問題S | 無實體人行道N | 有標記的行人穿越道W | 行人號誌P | 其他

放大Z

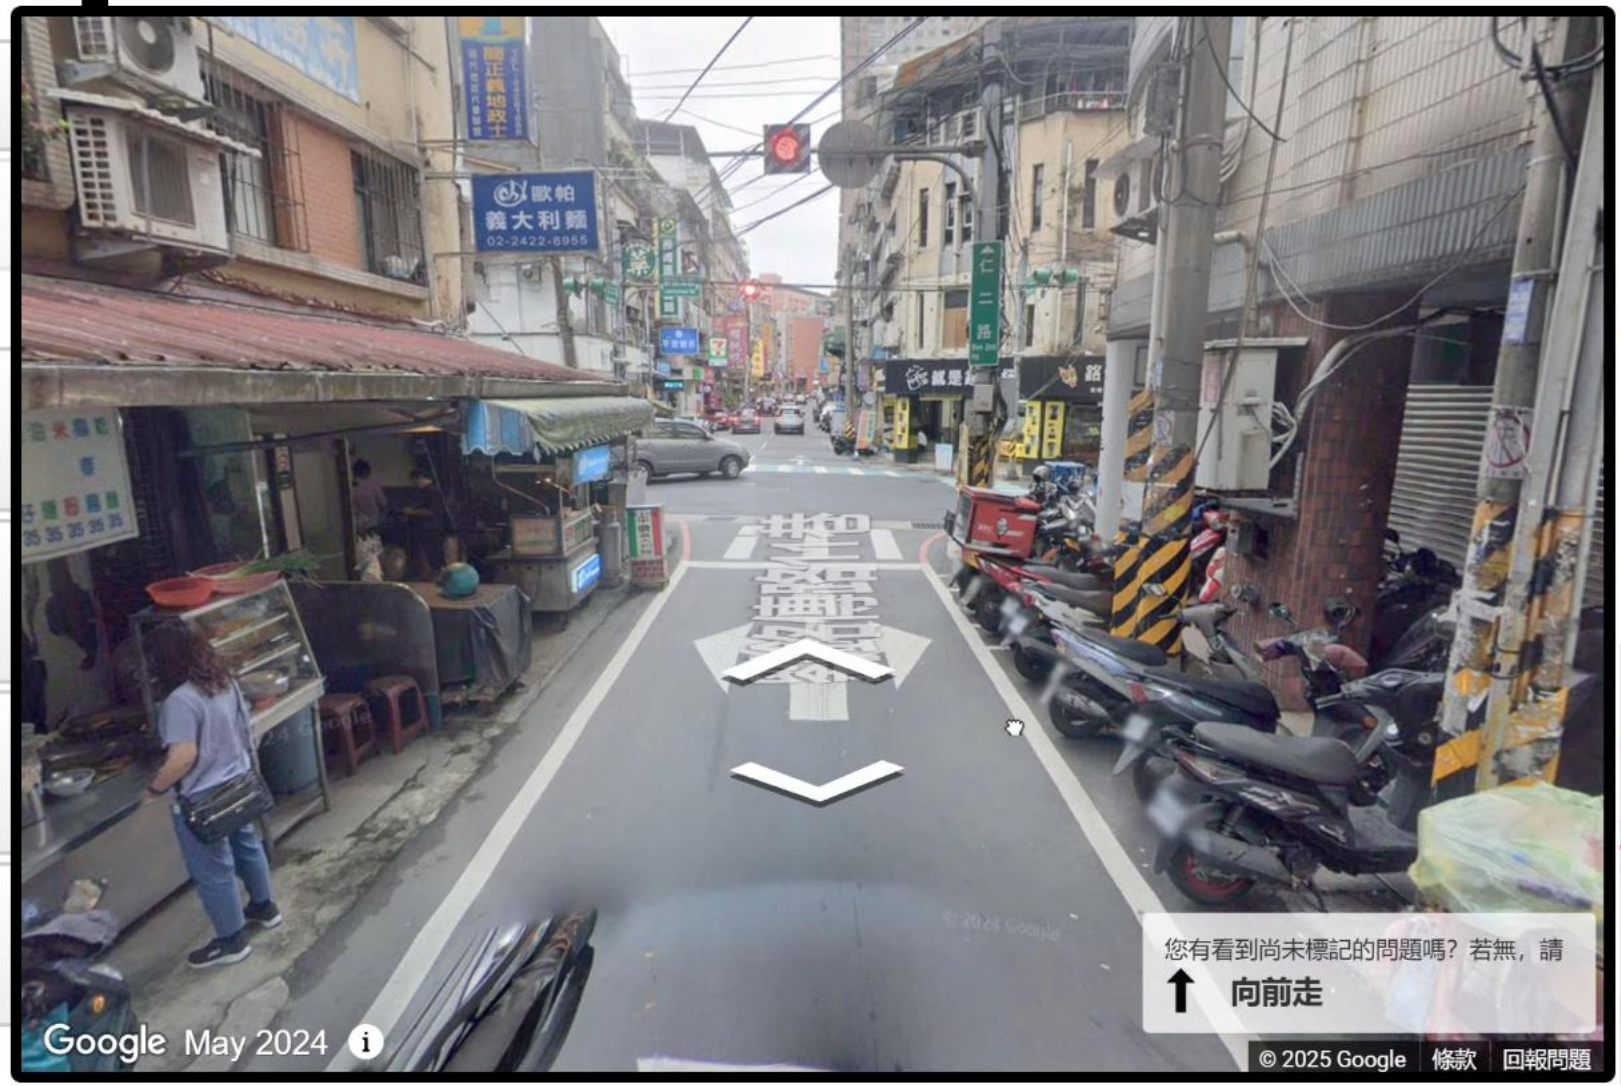
縮小

音效

跳過

卡住

回饋



整體統計數據

0.11 公里

50 標記

74.5% 精準度

目前所在的鄰里

水錦里, 基隆市

0.03 公里

0 標記

目前的任務

探索75 公尺並找到所有人行道可及性選項

9% 完成

0 路緣斜坡

0 鋪面問題

0 缺路緣斜坡

0 無實體人行道

0 有障礙物

0 其他



TROTTOIR OPRIT LABELS | TROTTOIR LABELS | ANDERE TOEGANKELIJKHEIDSLABELS

Ontdek | Trottoir Oprit (C) | Ontbrekende Trottoir Oprit | Obstakel in het Pad | Oppervlakteprobleem (S) | Geen Trottoir | Gemarkeerde oversteekplaats (W) | Verkeerslicht (P) | Anders

Zoom In

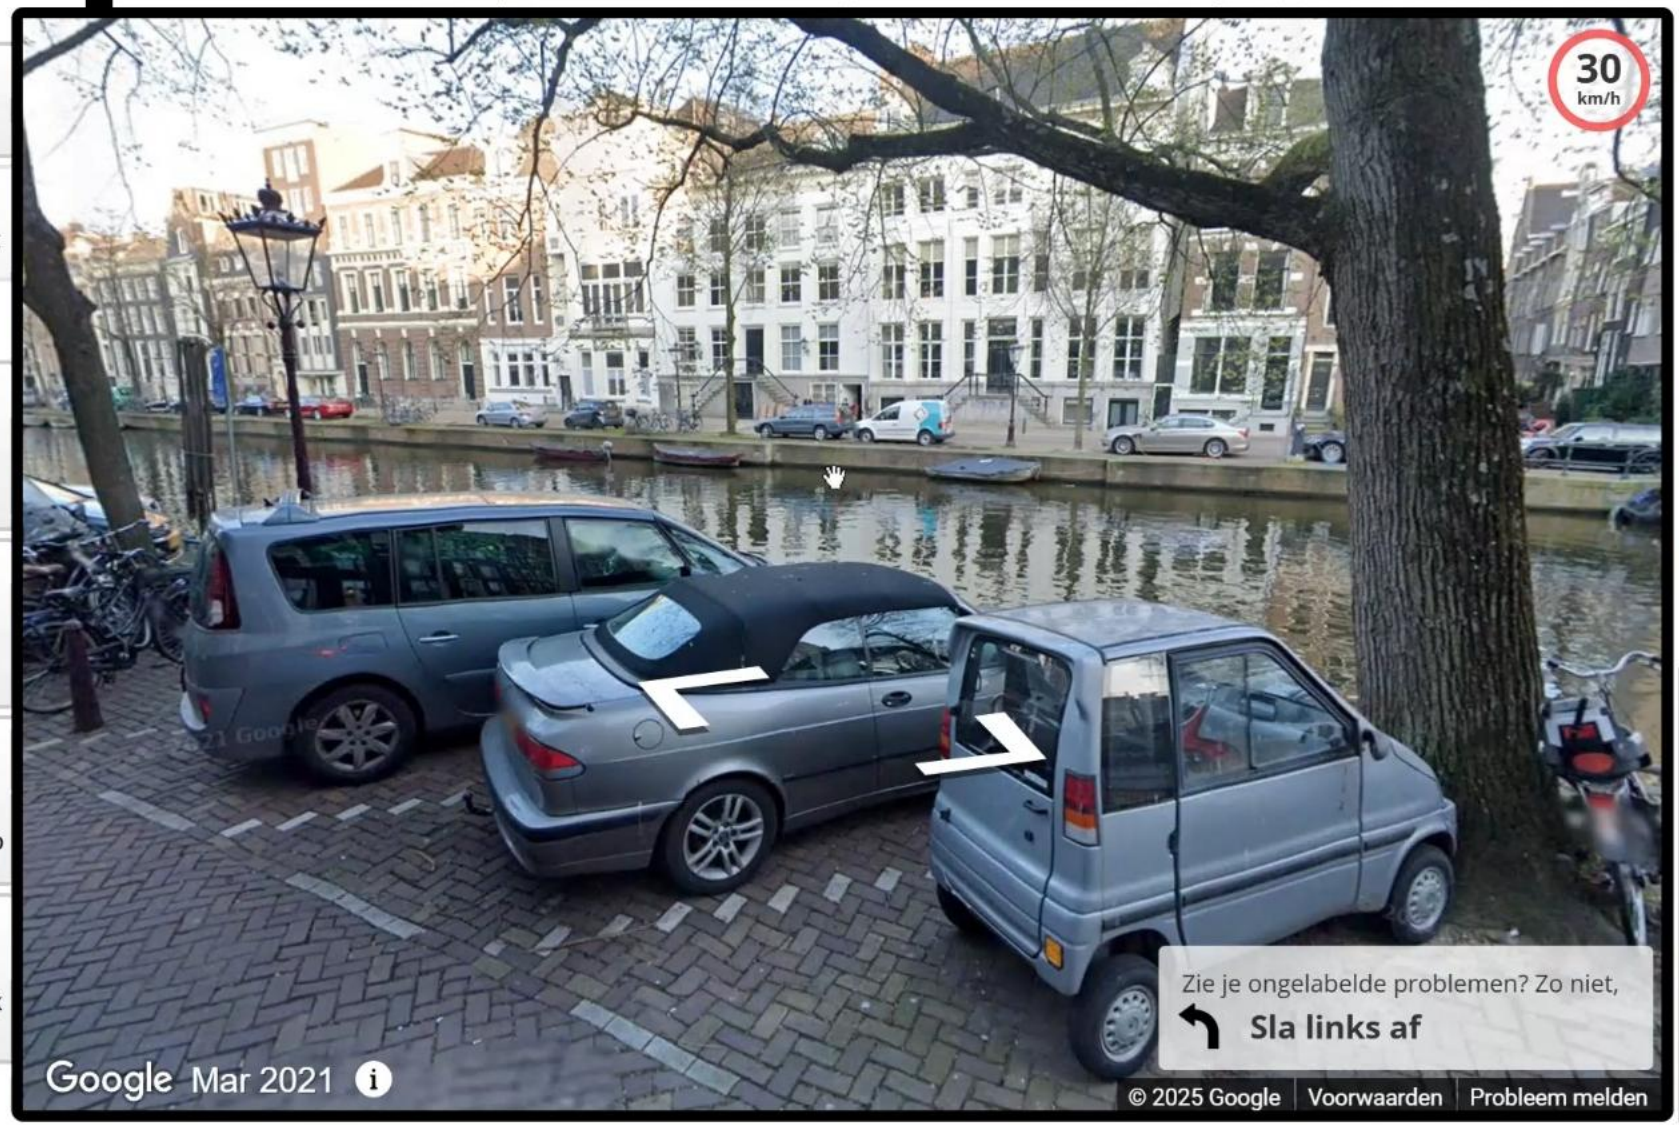
Zoom Uit

Geluid

Spring

Vastgelopen

Feedback



Algemene statistieken

3,06 kilometer

526 labels

74,4% nauwkeurigheid

Huidige Wijk

Leliegracht e.o., Amsterdam

0,95 kilometer

235 labels

Huidige Missie

Zoek 150 m voor deze wijk

0% voltooid

0 stoep oprit

0 oppervlakte problemen

0 ontbrekende stoep opritten

0 geen trottoirs

0 obstakels

0 anders




World Enabled
Ciudad Fácil
WPI
PLAKSHA UNIVERSITY
ACCESS LIVING
CITIES FOR ALL
GLOBAL CAMPAIGN ON INCLUSIVE AND ACCESSIBLE CITIES
CITY OF MENDOTA ILLINOIS
City of St. Louis
Oradell Girl Scouts
Hackensack Meridian School of Medicine
CREATE
LIGA PEATONAL
AIMCIL
Achieving Independence and Mobility
Center for Independent Living
scistarter
People-powered science
EL COLEGIO DE MICHOACÁN
UNIVERSIDAD DEL AZUAY
ACTIVE TRANSPORTATION ALLIANCE
Chicago Metropolitan Agency for Planning
OPEN COLUMBUS
GREATER FORT WAYNE INC.
METRO CHAMBER ALLIANCE
trailnet
STREETS FOR ALL
UNIVERSIDAD AUTÓNOMA DE SINALOA
MetropolitanPlanningCouncil
shaping a more equitable, sustainable and prosperous Chicago region since 1934
SFU
SIMON FRASER UNIVERSITY
BIKE PGH!
UNIVERSIDAD DE MONTERREY
HOMO HOMINIS IN MINISTERIO PERFICITUR
WSDOT
City of Amsterdam
San Pedro Garza García
GOBIERNO MUNICIPAL
City of Burnaby
City of Newberg
USP
Universidade de São Paulo
amsterdam intelligence
Regional Transportation Authority
PACIFIC NW TRANSPORTATION CONSORTIUM
PACTRANS
REGION 10
PROGRESS CENTER
for independent living
Chicago Metropolitan Agency for Planning
UNIVERSITY OF SOUTH CAROLINA
Bellevue School District
DC
UW Disability & D/deaf Cultural Center
Vitruvius
give it strength
make it useful
deliver it beautifully
WindSwell
FOUNDATION
UIC
UNIVERSITY OF ILLINOIS CHICAGO
Universität Zürich[UZH]
Metropolitan Mayors Caucus
MS
Disability Rights WASHINGTON
UF UNIVERSITY of FLORIDA
College of Health and Human Performance
Tecnológico de Monterrey
Teaneck High School
公民幫推
People Nudge
DENNY
Denny International Middle School
CIVIC INNOVATION LAB
Great Lakes ADA Center
A Member of the ADA National Network
Nondrivers Alliance
WEEK WITHOUT DRIVING

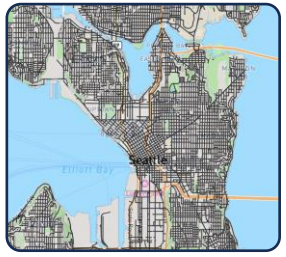
1 PREPROCESSING

Create labeling routes

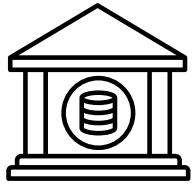
OpenStreetMap



Filtered Street Network



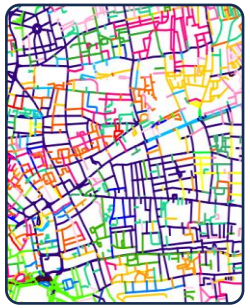
Open Gov Data



Neighborhood Boundaries



Labeling Mission Routes



2 LABEL SIDEWALKS

Collect geo-located sidewalk assessments

Streetscape Imagery APIs

Google Street View

Meta Mapillary

Panoramax

iNovitas infra3D



Community Partners

Labeling Missions



AI Labeler



Custom Trained ConvNeXt V2

CURB RAMPS



MISSING RAMP



SURFACE PROB



OBSTACLE



NO SIDEWALK



CROSSWALK



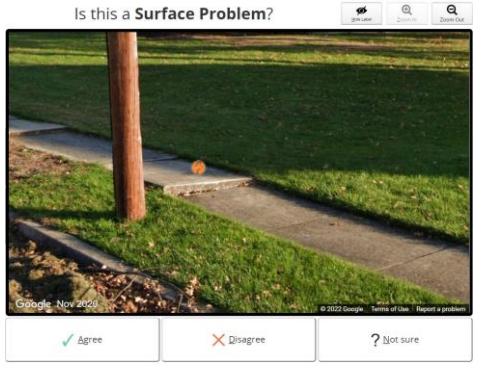
PED SIGNAL



3 VALIDATE LABELS

Human-AI quality control

Human Validation UI



AI-Assisted Human Validation



Custom Trained DINOv2

AI Validation

Agree



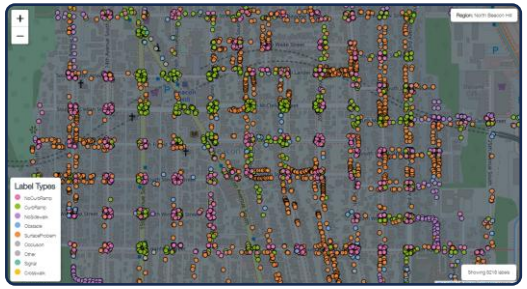
Disagree



4 POSTPROCESSING

Prepare data for apps

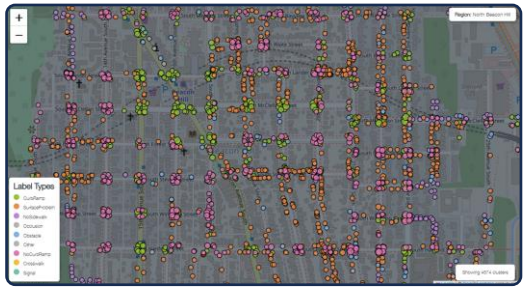
Raw Labels



Remove poor-quality labelers

Geo-cluster proximal labels

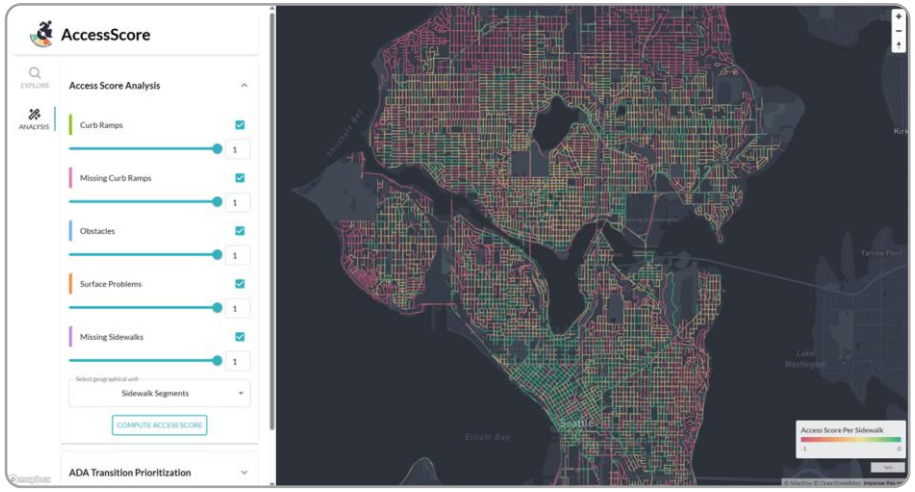
Label Clusters



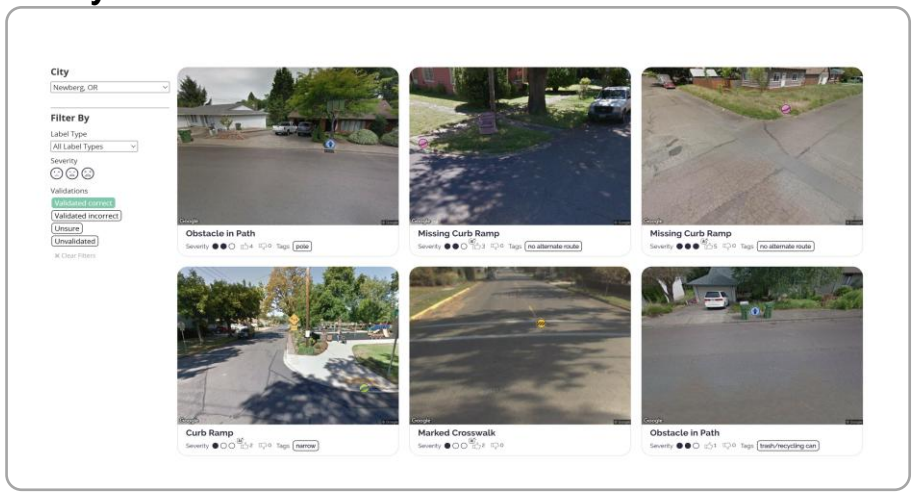
5 APPLICATIONS

Analytics & mobility tools

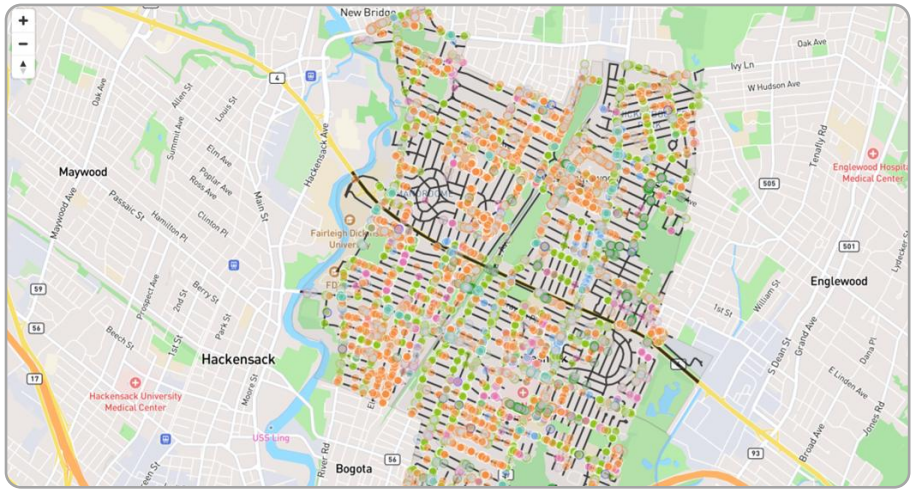
AccessScore



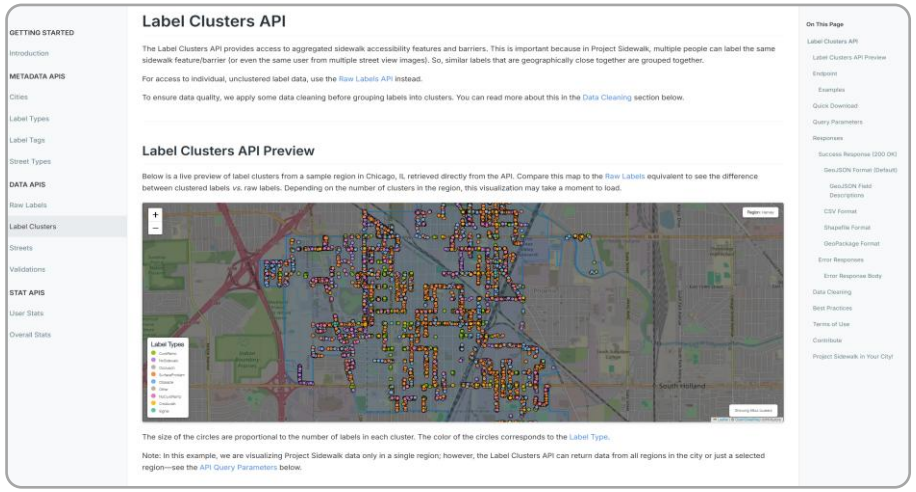
Gallery



LabelMap



API



2 LABEL SIDEWALKS

Collect geo-located sidewalk assessments

Streetscape Imagery APIs

Google Street View

Meta Mapillary

Panoramax

iNovitas infra3D

Community Partners

Labeling Missions



AI Labeler



Custom Trained ConvNeXt V2

CURB RAMPS



MISSING RAMP



SURFACE PROB



OBSTACLE



NO SIDEWALK



CROSSWALK



PED SIGNAL



3 VALIDATE LABELS

Human-AI quality control

Human Validation UI



AI-Assisted Human Validation



Custom Trained DINOv2

AI Validation

Agree



Disagree



4 POSTPROCESS

Prepare data for apps

Raw Labels

Remove poor-quality labelers

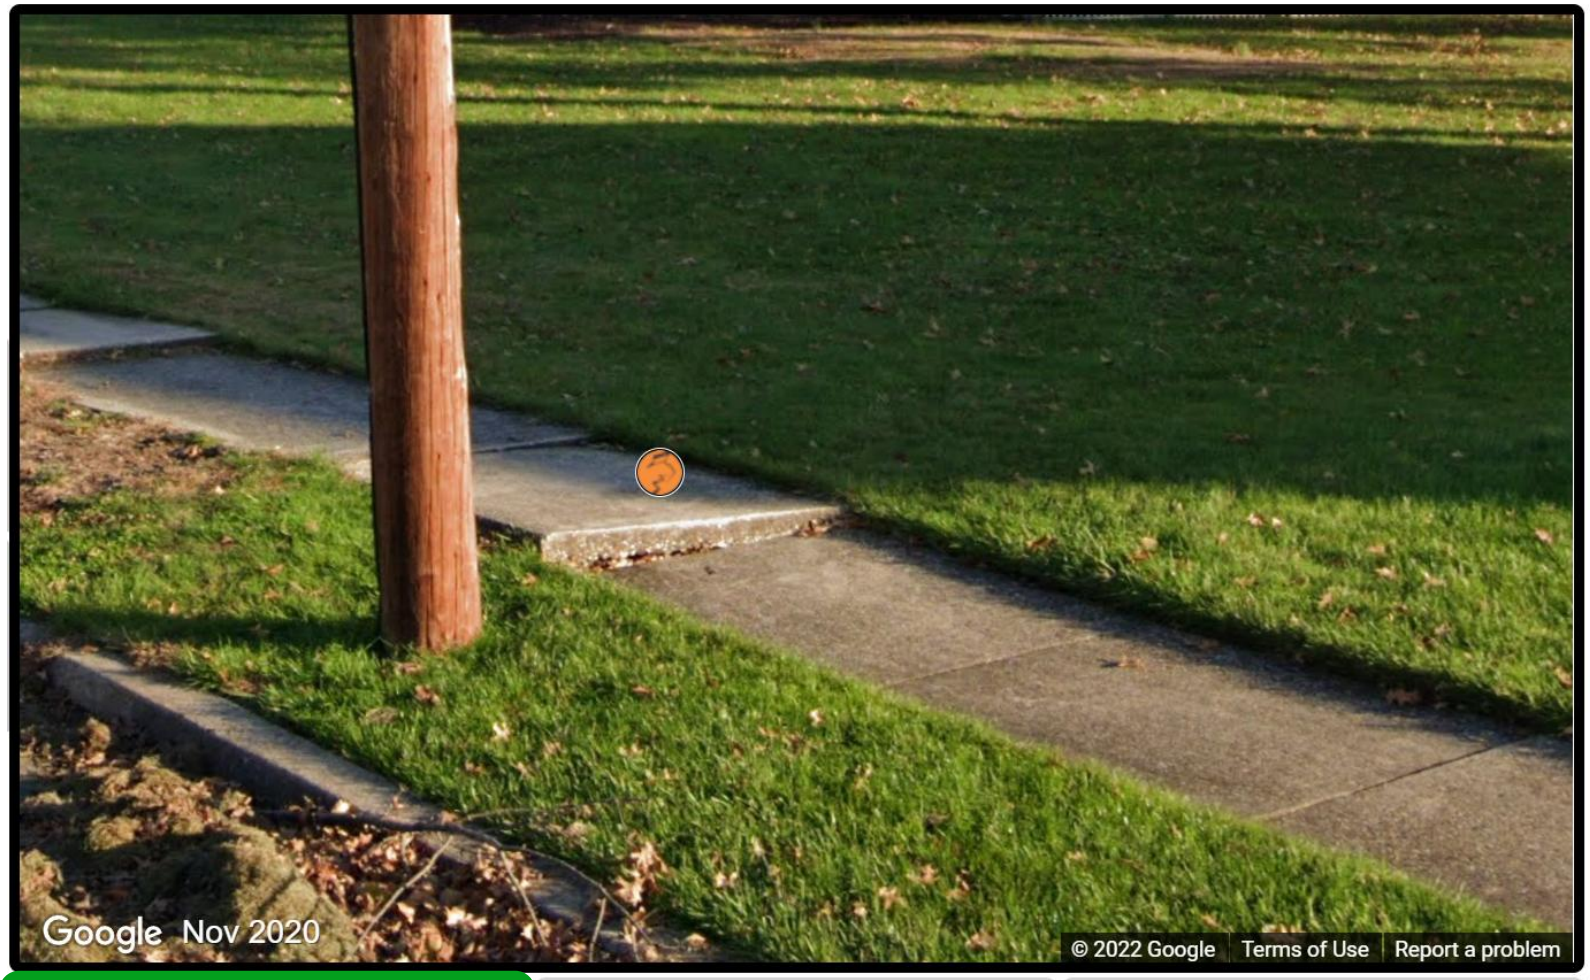
Geo-cluster proximal labels

Label Clusters

Is this a **Surface Problem**?

Hide Label | Zoom In | Zoom Out

3961 labels



✓ Agree

✗ Disagree

? Not sure

Add comment here...

Current Mission
Validate 10 labels
0% complete

Correct Examples



Incorrect Examples



Is this a **Missing Curb Ramp**?

Hide Label | Zoom In | Zoom Out

1113 labels



Current Mission
Validate 10 labels
50% complete

CORRECT Missing Curb Ramp



INCORRECT Missing Curb Ramp



✓ Agree | ✗ Disagree | ? Not sure

Add comment here...

Hide Label

Is this a **Curb Ramp**?

Zoom In | Zoom Out

0 labels



✓ Agree | ✗ Disagree | ? Not sure

Current Mission

Validate 10 labels

0% complete

Curb Ramp



NOT a Curb Ramp



Hide Label

Is this a **Curb Ramp**?

Zoom In | Zoom Out



✓ Agree | ✗ Disagree | ? Not sure

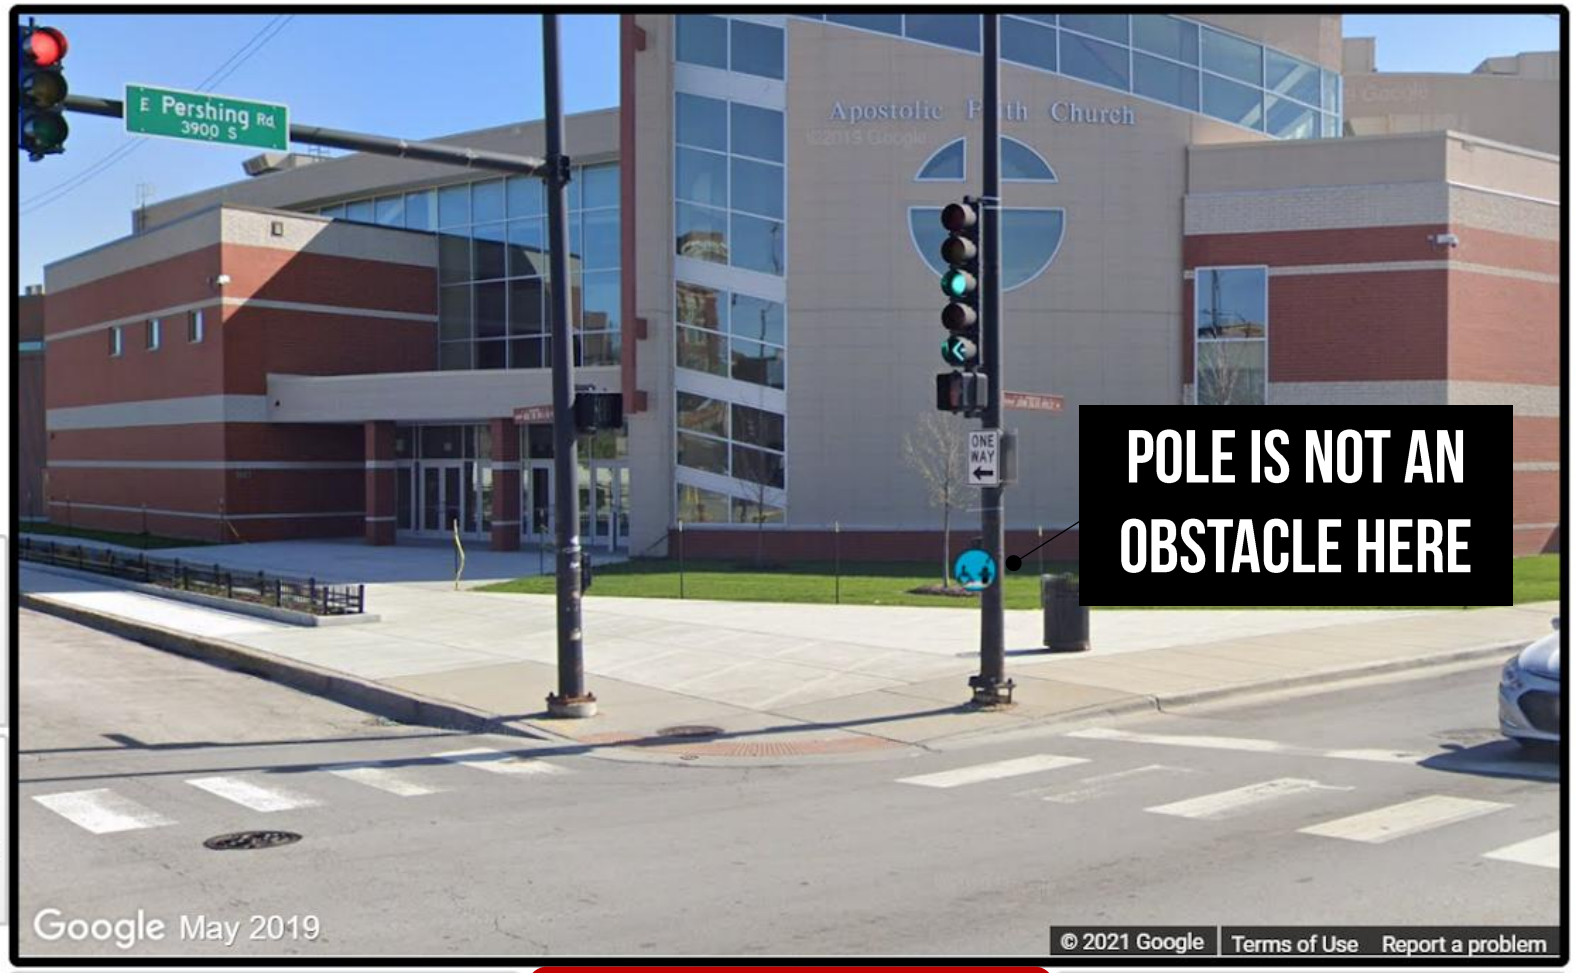
Labeled **every driveway** as a **curb ramp**



Is this an **Obstacle**?

Hide Label | Zoom In | Zoom Out

934 labels



Skip

Feedback

✓ Agree | ✗ Disagree | ? Not sure

Add comment here...

Current Mission
Validate 10 labels
0% complete

Correct Examples



Incorrect Examples



Hide Label

Is this a **Surface Problem**?

Zoom In | Zoom Out

9 labels



Skip

Feedback

✓ Agree

✗ Disagree

? Not sure

Current Mission

Validate 10 labels

90% complete

Surface Problem



NOT a Surface Problem

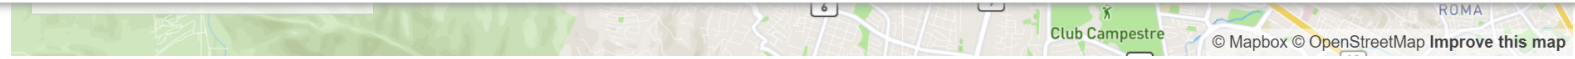




Image source: montgomeryplanning.org


Image source: montgomeryplanning.org
ONLY MARK BARRIERS IN THE PEDESTRIAN PATH
Frontage Zone
Pedestrian Zone
Planting/ Furnishing
Curb Zone
Bike Facility
Buffer

We also try to make Project Sidewalk **fun** and **educational**



Your missions	Distance	Labels	Validations	Accuracy
140	2.03 mi	568	1249	90.7%

Achievements

Missions

Congratulations, you've earned all mission badges!



Distance

Thanks for helping! **2.97 more miles** until your next achievement.



Labels

Great job! **432 more labels** until your next achievement.



Validations

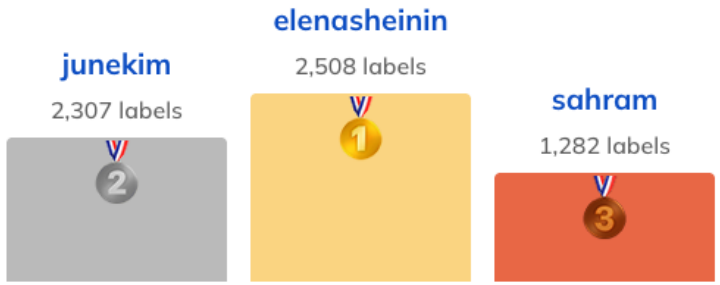
Amazing work! **3751 more validations** until your next achievement.



COMMUNITY
Leaderboard
YOU
Dashboard
Settings

Top of the board

Paterson's all-time leaders. 🏆



Paterson all-time

The top all-time mappers in Paterson ranked by overall score.

▶ How is the ranking computed?

#	Mapper	Labels	Missions	Distance	Accuracy ⓘ
1	elenasheinin	2,508	327	26.3 mi	96%
2	junekim	2,307	366	30.6 mi	96%
3	sahram	1,282	307	19.5 mi	98%
4	evakan	1,277	149	12.2 mi	96%
5	carcha	1,216	128	10.5 mi	96%

On This Page
Top of the board
Paterson all-time
Paterson this week
Your place in Paterson this week
Teams in Paterson
All cities, all-time
Community impact

PROJECT SIDEWALK

REMOTE CROWDSOURCING

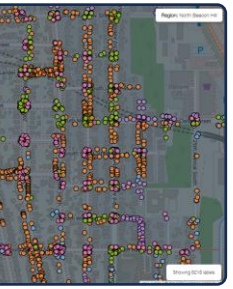


LABELING MEXICO CITY FROM GERMANY!

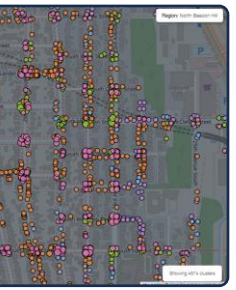
Source: Edgar Martínez, Liga Peatonal

ROCESSING

a for apps



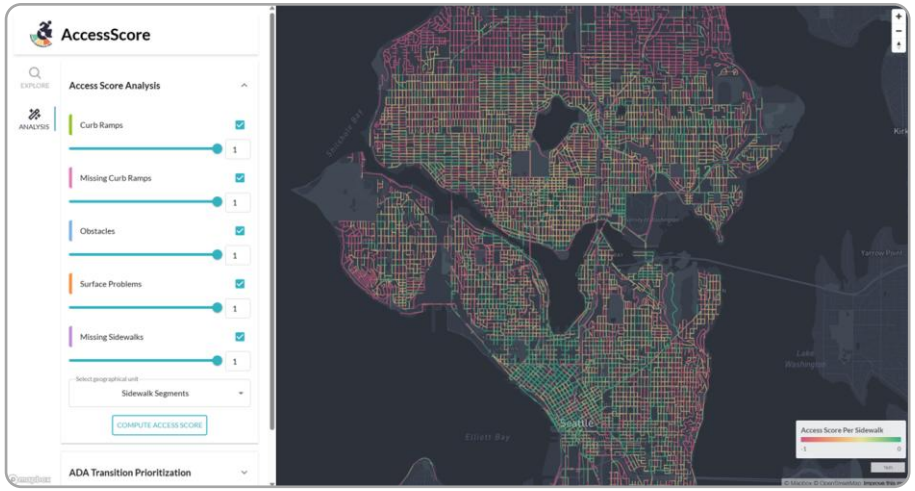
Geo-cluster
roximal labels



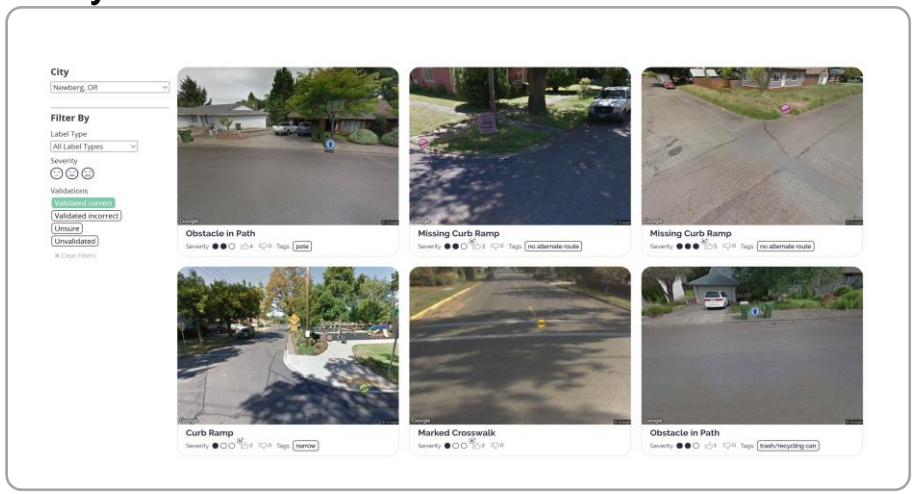
5 APPLICATIONS

Analytics & mobility tools

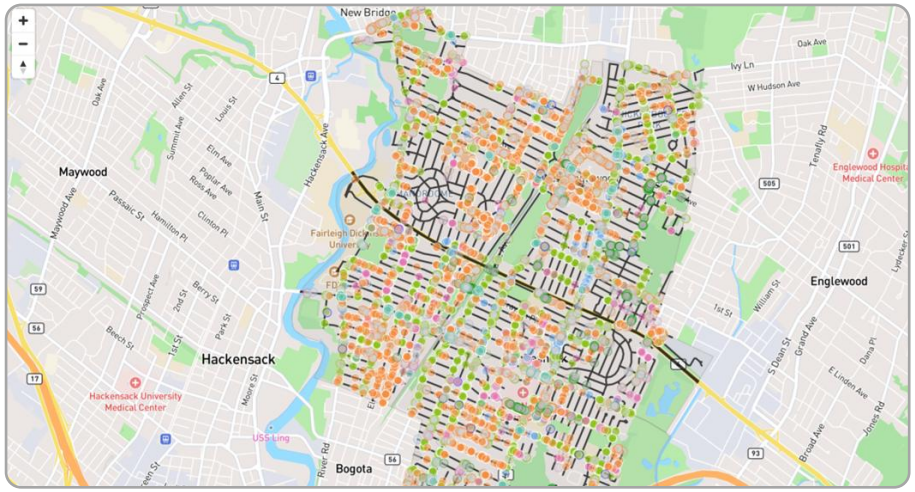
AccessScore



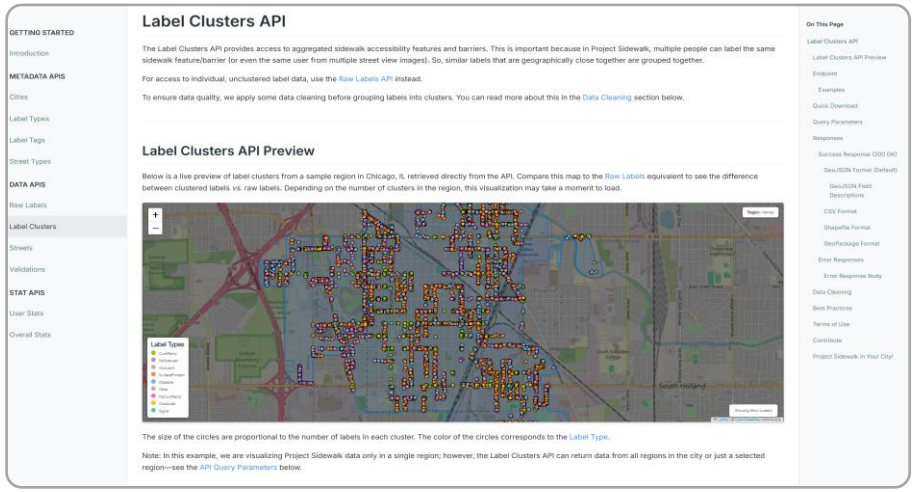
Gallery



LabelMap



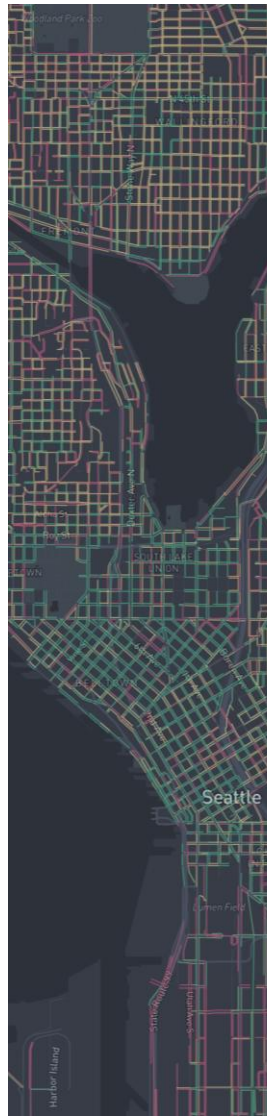
API

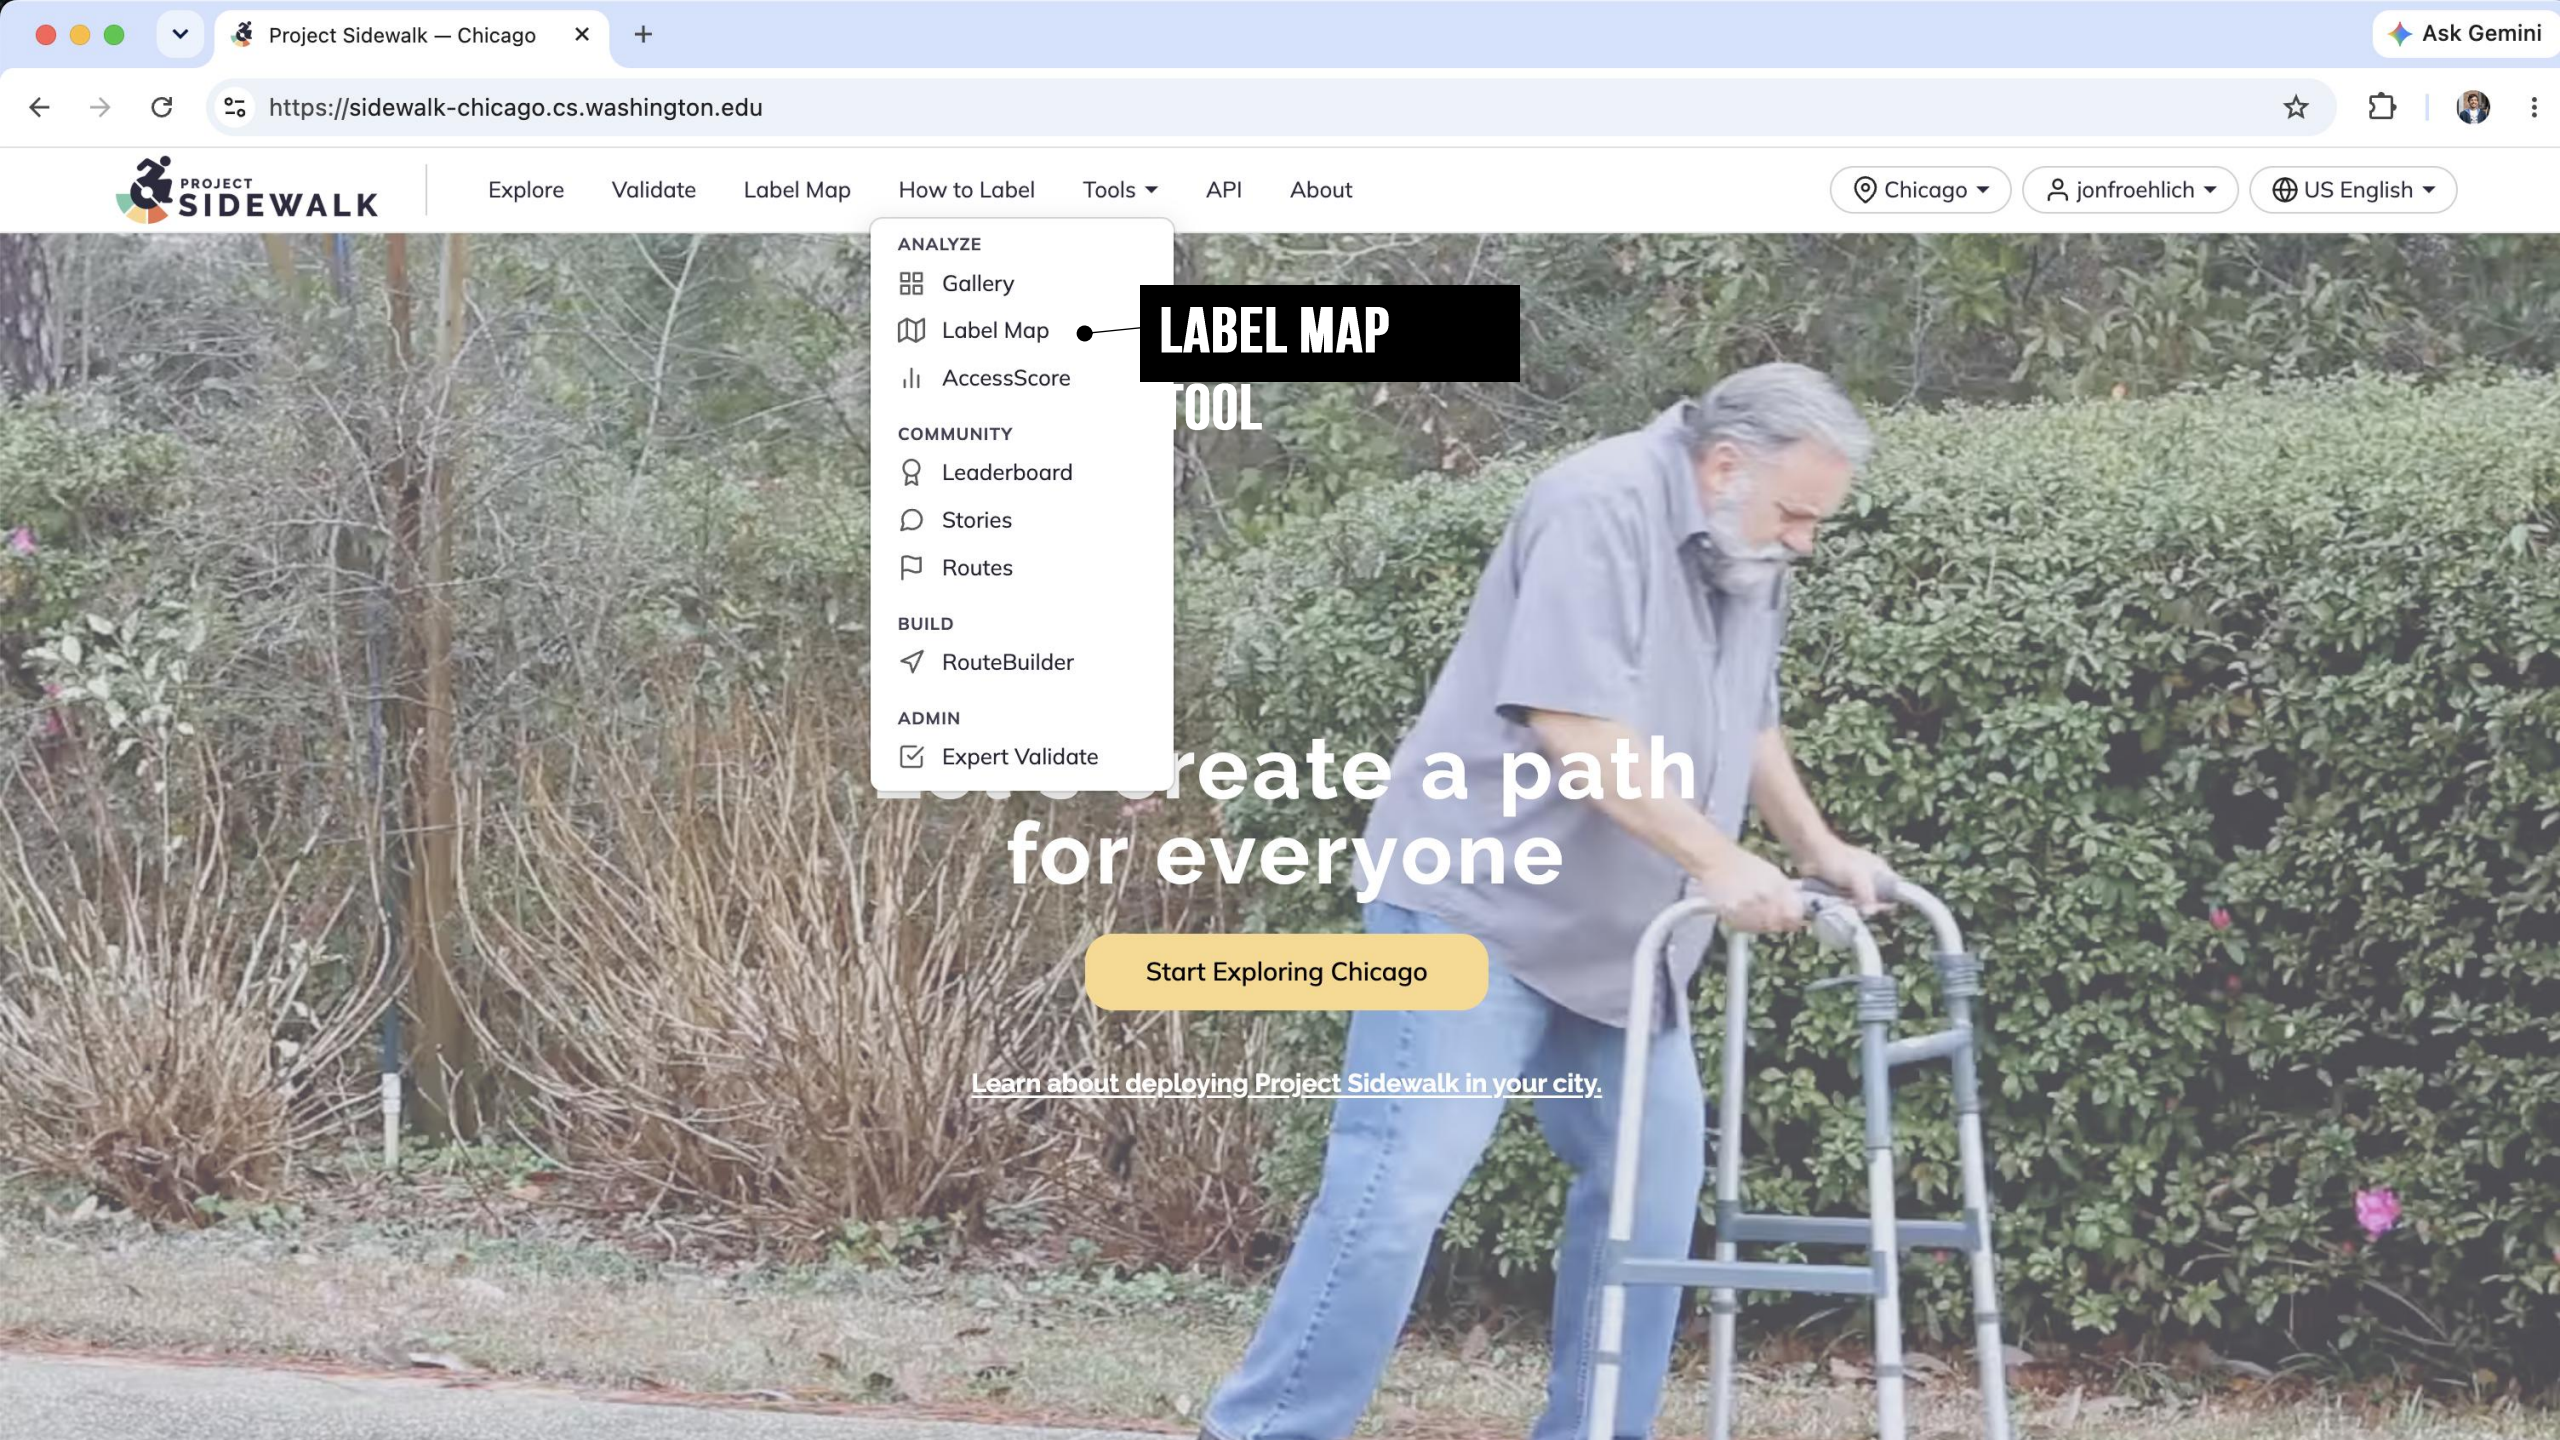


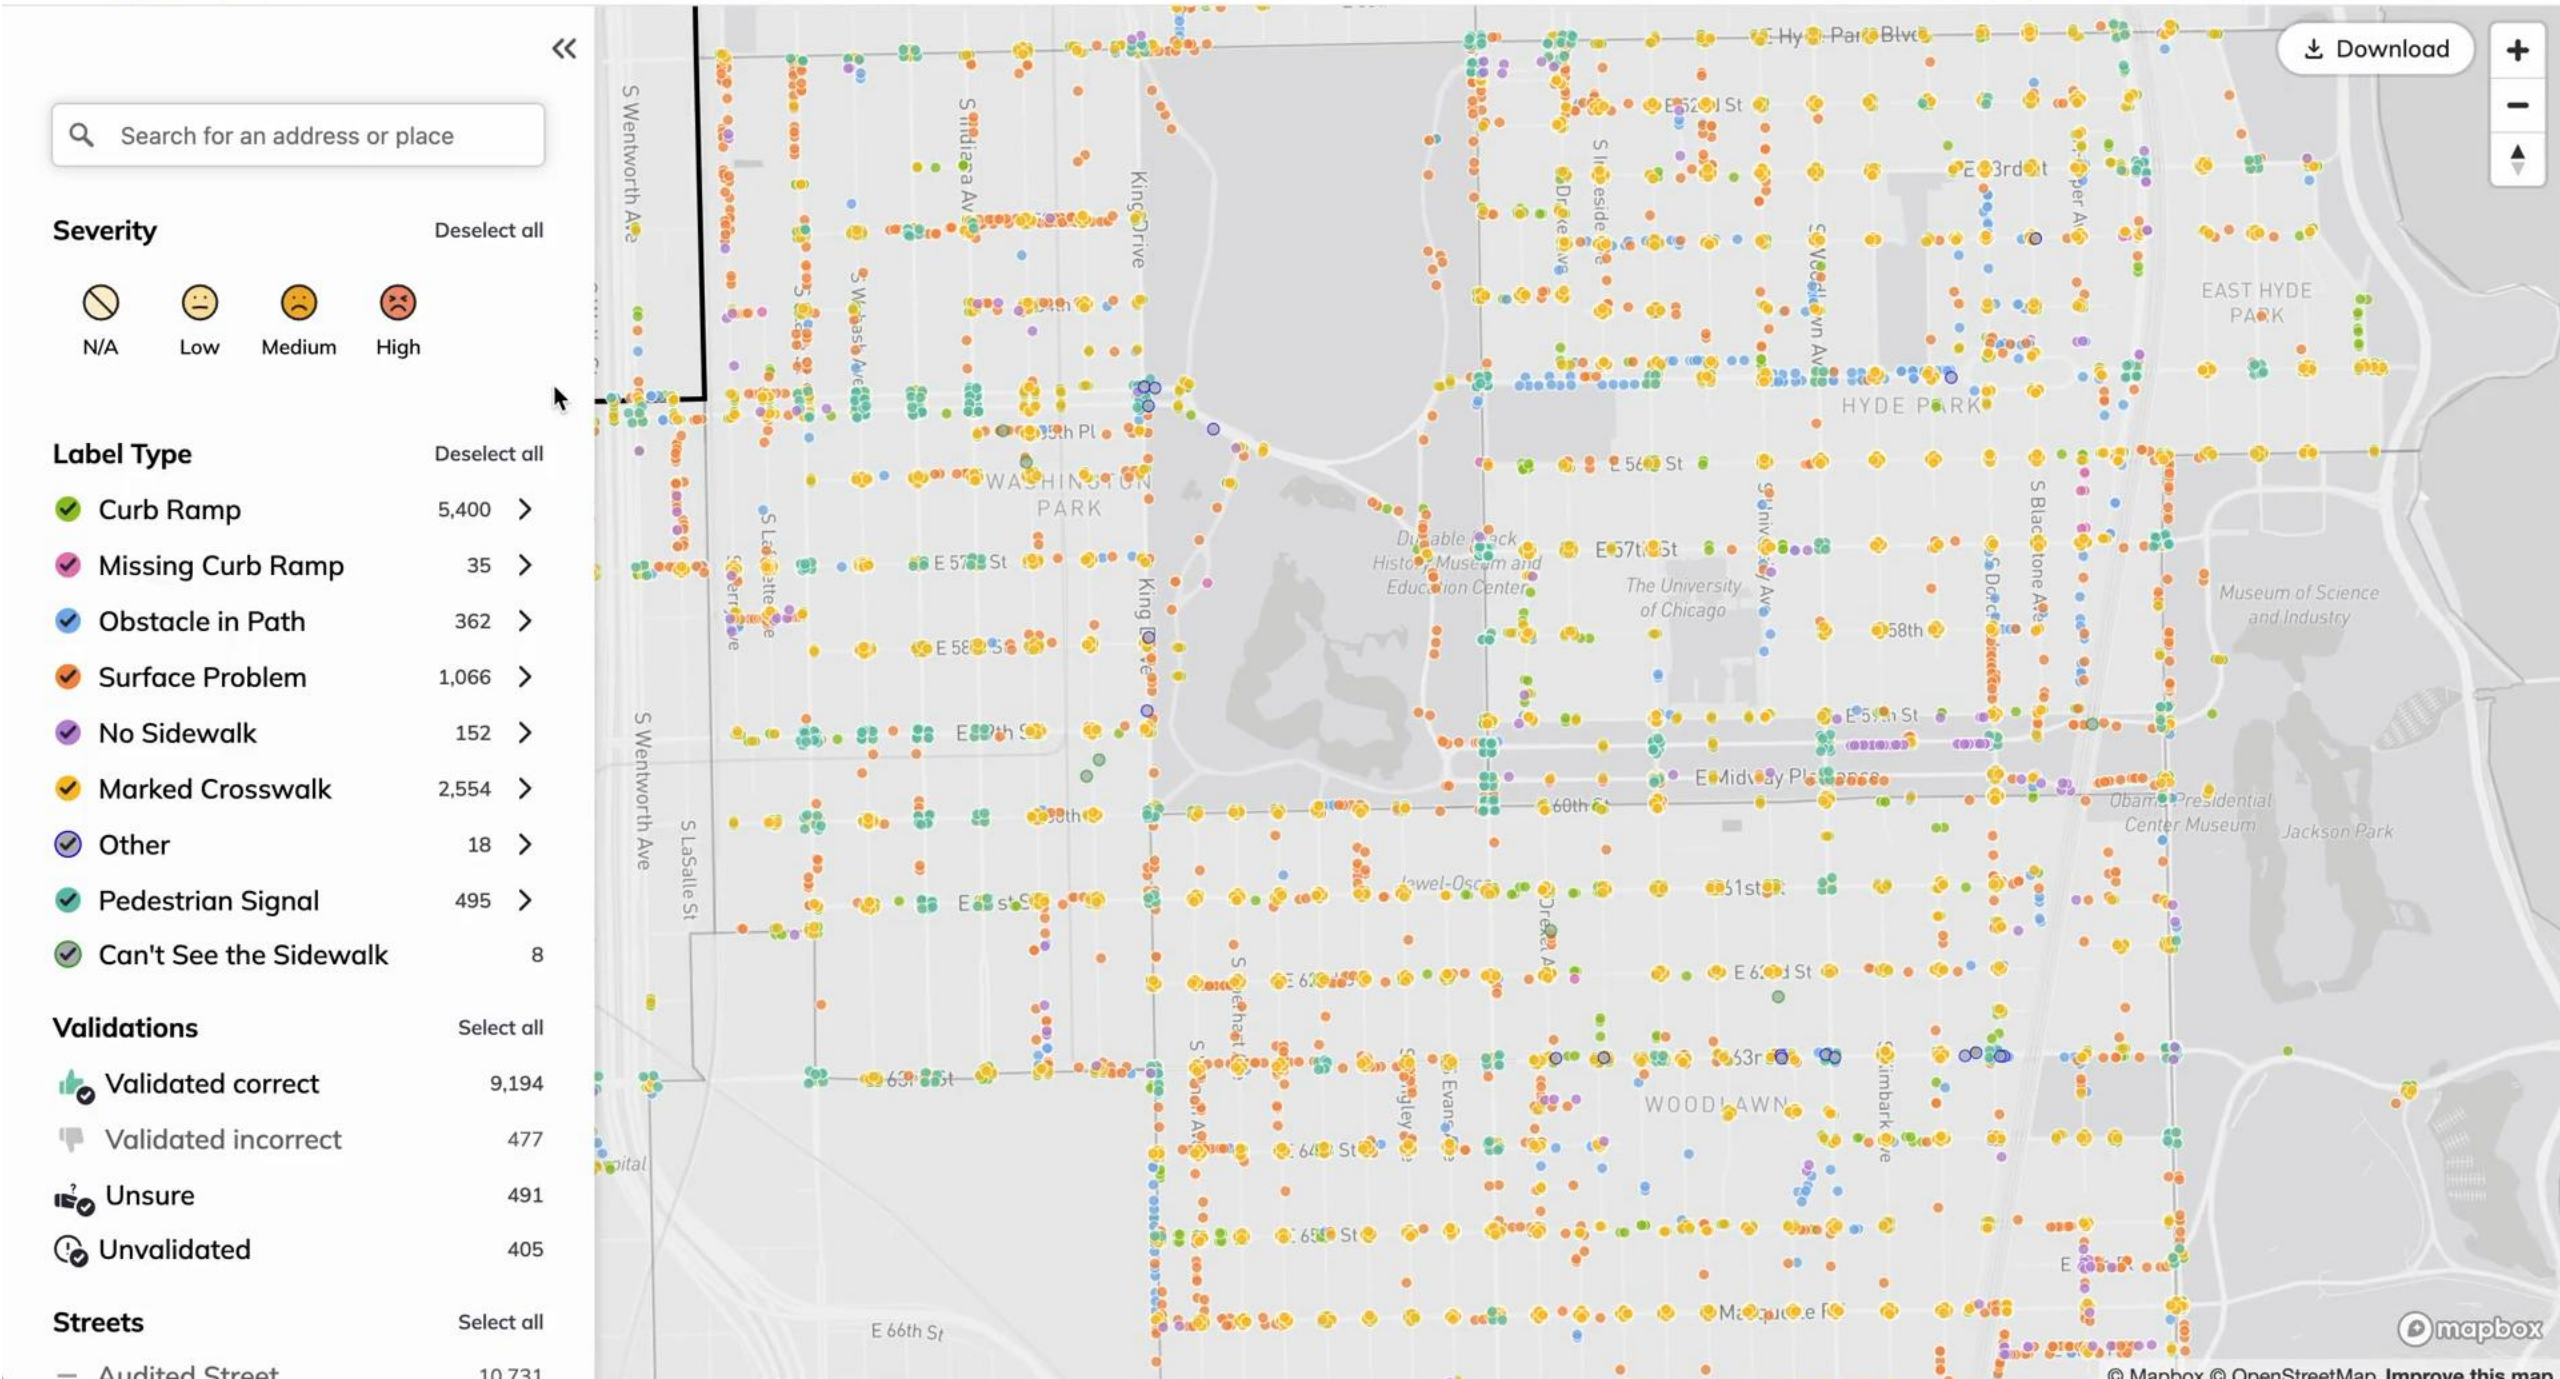
6 OUTCOM

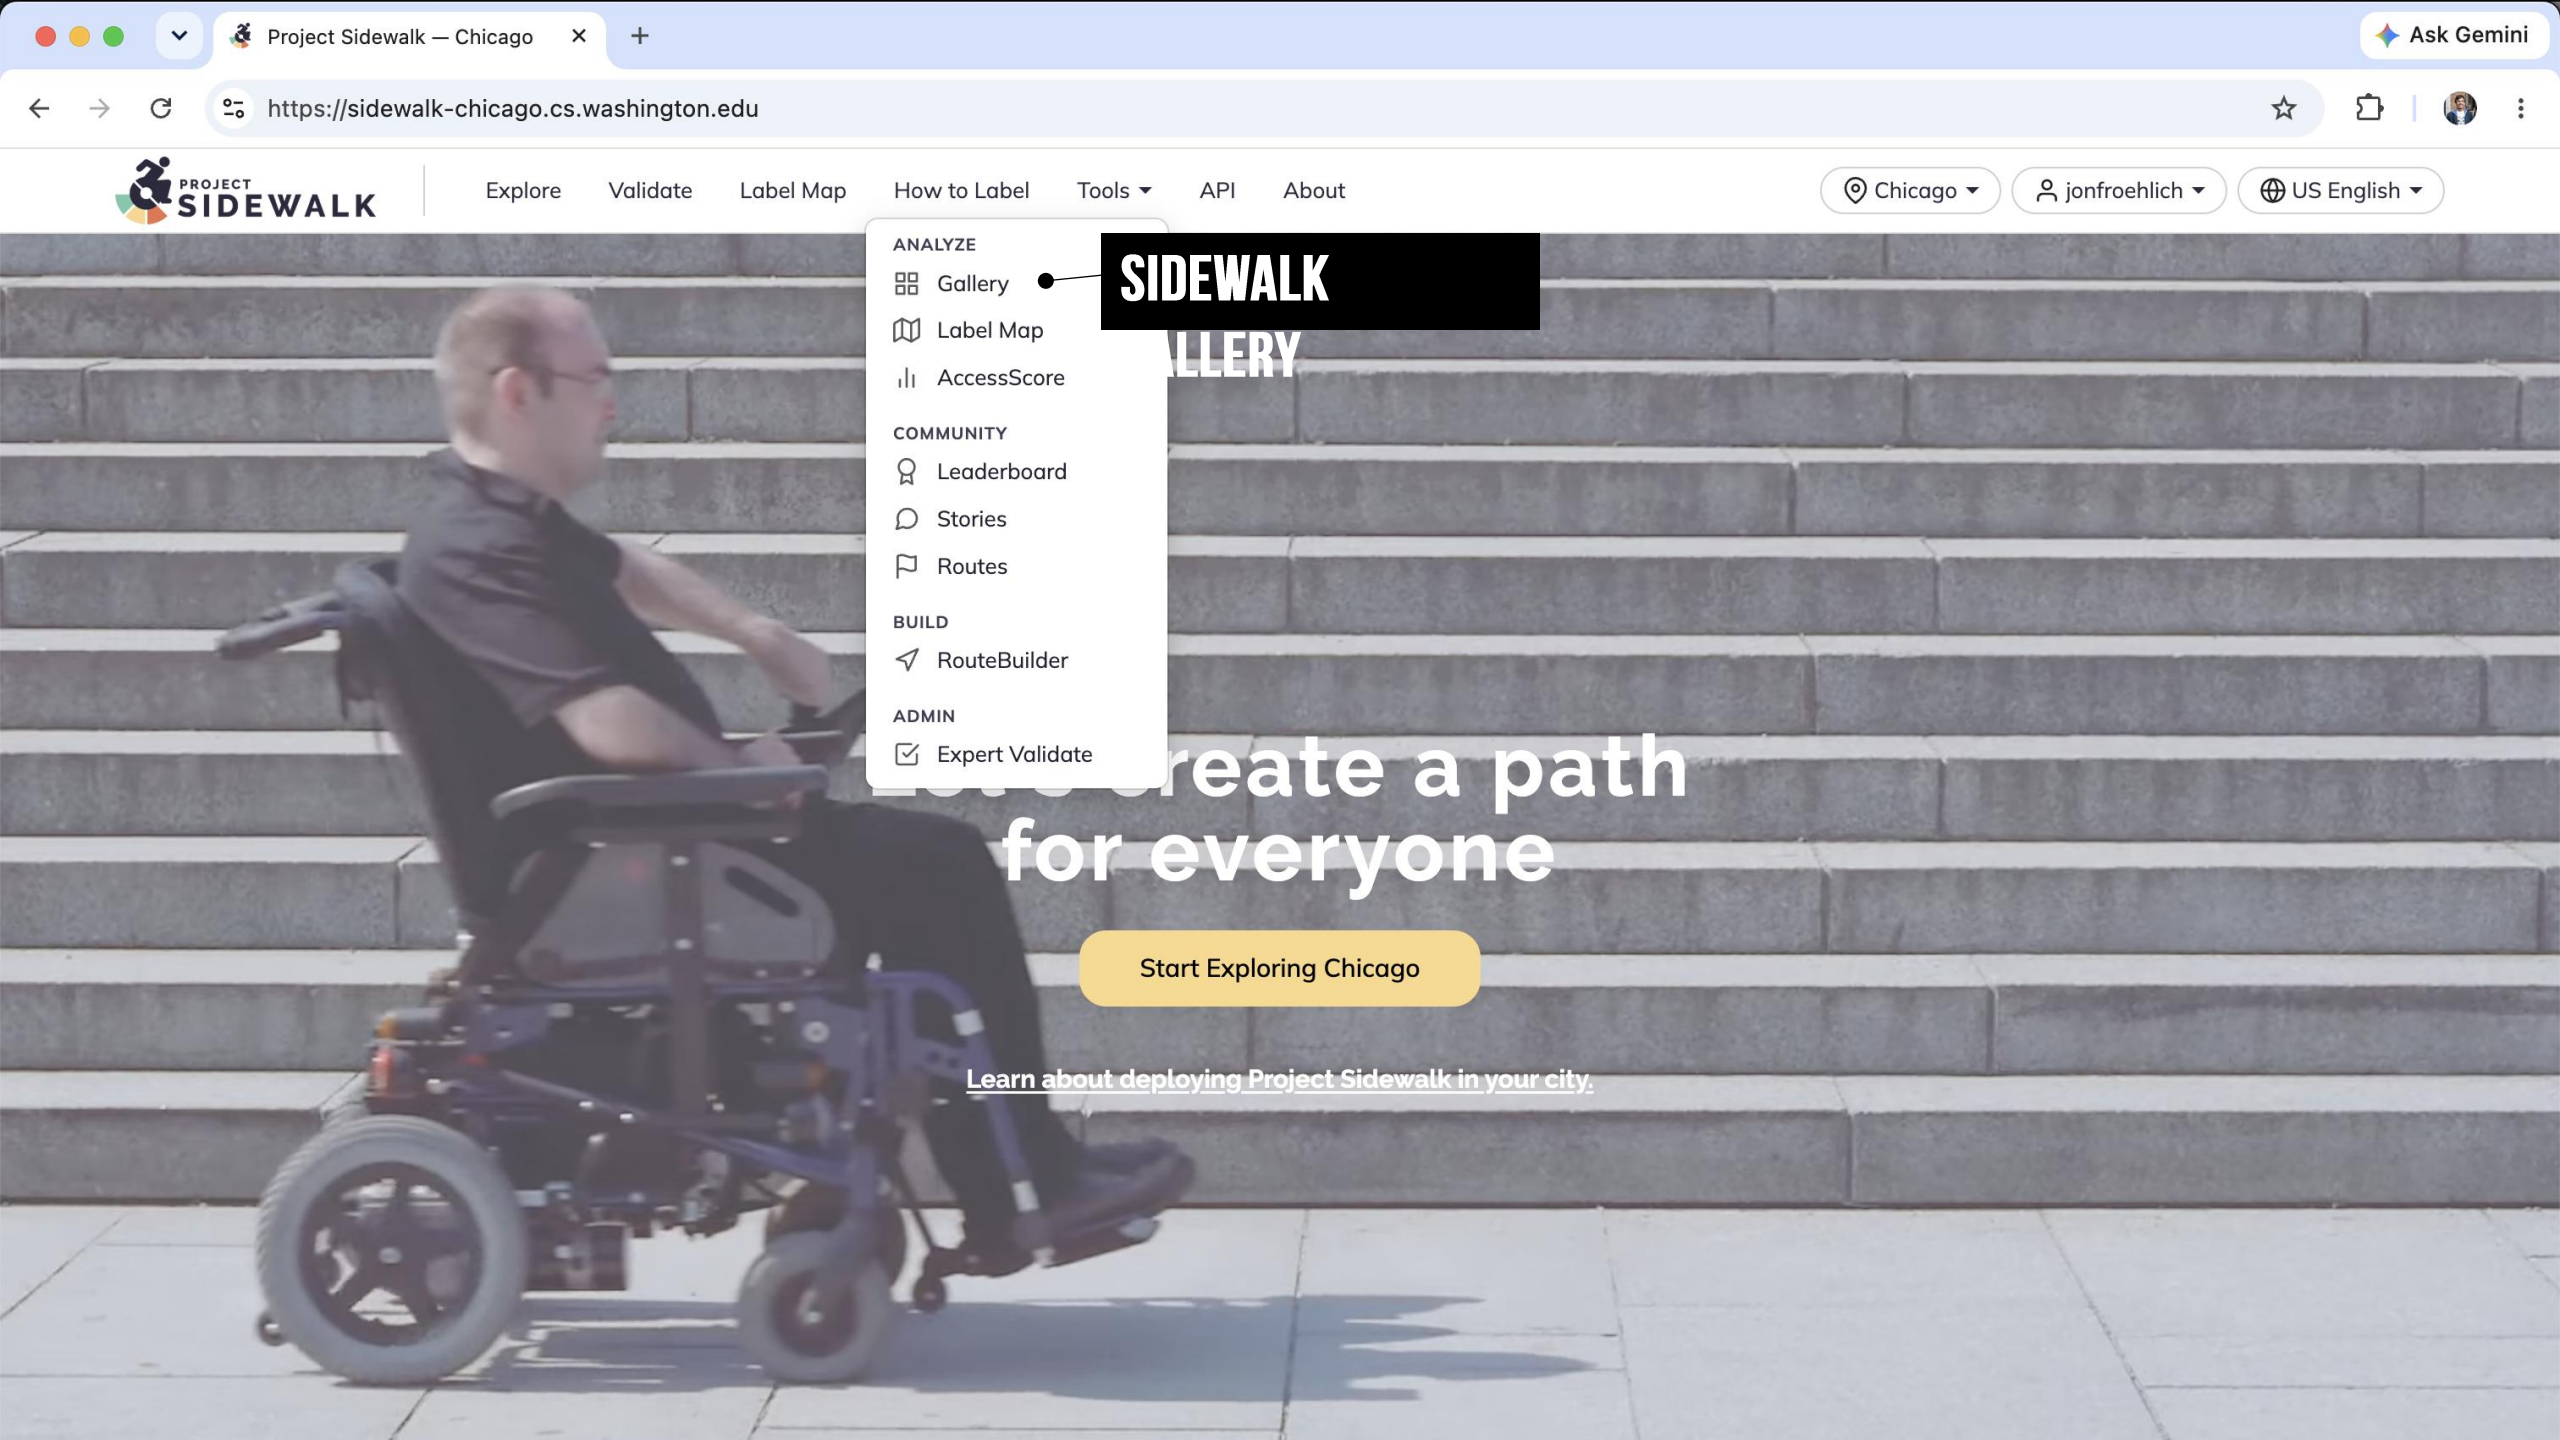
Research to

Research




Project Sidewalk — Chicago
Ask Gemini
https://sidewalk-chicago.cs.washington.edu
PROJECT SIDEWALK
Explore
Validate
Label Map
How to Label
Tools
API
About
Chicago
jonfroehlich
US English
ANALYZE
Gallery
Label Map
AccessScore
COMMUNITY
Leaderboard
Stories
Routes
BUILD
RouteBuilder
ADMIN
Expert Validate
LABEL MAP
TOOL
reate a path
for everyone
Start Exploring Chicago
Learn about deploying Project Sidewalk in your city.


Download
Search for an address or place
Severity
Deselect all
N/A
Low
Medium
High
Label Type
Deselect all
Curb Ramp 5,400
Missing Curb Ramp 35
Obstacle in Path 362
Surface Problem 1,066
No Sidewalk 152
Marked Crosswalk 2,554
Other 18
Pedestrian Signal 495
Can't See the Sidewalk 8
Validations
Select all
Validated correct 9,194
Validated incorrect 477
Unsure 491
Unvalidated 405
Streets
Select all
Audited Street 10,731
S Wentworth Ave
S Indiana Ave
King Drive
S Drexel Ave
S Lakeside
S Woodlawn Ave
EAST HYDE PARK
HYDE PARK
WASHINGTON PARK
E 57th St
E 58th St
E 59th St
E 60th St
E 61st St
E 63rd St
E 64th St
E 65th St
E 66th St
S LaSalle St
S Blackstone Ave
S Dorchester Ave
S University Ave
Dusable Black History Museum and Education Center
The University of Chicago
Museum of Science and Industry
Obama Presidential Center Museum
Jackson Park
WOODLAWN
S Cornell Ave
S Evans Ave
S Kimbark Ave
mapbox
© Mapbox © OpenStreetMap Improve this map


Project Sidewalk — Chicago
Ask Gemini
https://sidewalk-chicago.cs.washington.edu
PROJECT SIDEWALK
Explore
Validate
Label Map
How to Label
Tools
API
About
Chicago
jonfroehlich
US English
ANALYZE
Gallery
Label Map
AccessScore
COMMUNITY
Leaderboard
Stories
Routes
BUILD
RouteBuilder
ADMIN
Expert Validate
SIDEWALK
reate a path
for everyone
Start Exploring Chicago
Learn about deploying Project Sidewalk in your city.

Filter By

✖ Clear Filters

Severity

Deselect all

N/A | Low | Medium | High

Label Type

Deselect all

- Curb Ramp
- Missing Curb Ramp
- Obstacle in Path
- Surface Problem
- No Sidewalk
- Marked Crosswalk
- Other
- Pedestrian Signal
- Can't See the Sidewalk

Validations

Select all

- Validated correct
- Validated incorrect
- Unsure
- Unvalidated



Curb Ramp — Shadyside

Quality ●○○ 2 0



Curb Ramp — Homewood North

Quality ●○○ 1 0



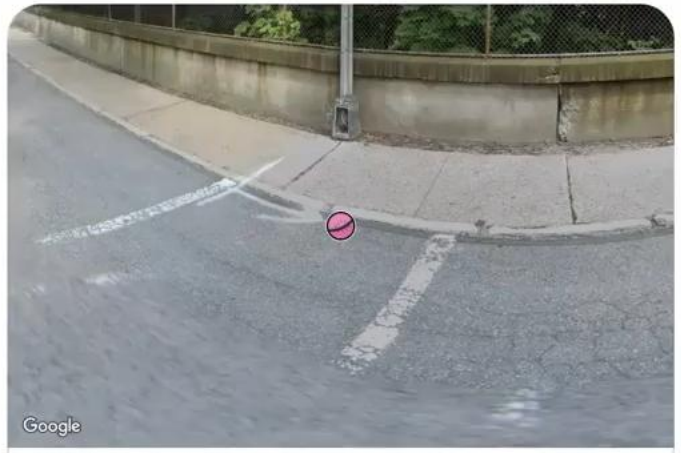
Curb Ramp — North Shore

Quality ●○○ 1 0 Tags points into cross traffic



Missing Curb Ramp — Swissvale

Severity ●●○ 3 0 Tags no alternate route



Missing Curb Ramp — North Oakland

Severity ●●● 5 3 Tags no alternate route



Missing Curb Ramp — Millvale

Severity ●●● 1 0 Tags no alternate route







How can **AI help**?

Google Faculty Research Award, 2012

Combining Crowdsourcing and Computer Vision for Street-level Accessibility

PI: Jon Froehlich (HCI), Co-PI: David Jacobs (Computer Vision)
Department of Computer Science, University of Maryland, College Park
{ jonf, djacobs }@cs.umd.edu
Google Sponsors: Peng Dai (daipeng@google.com) and Sacha Arnoud (sacha@google.com)



(a) Object in Path (b) No Curb Ramp (c) Surface Problem (d) Prematurely Ending Sidewalk

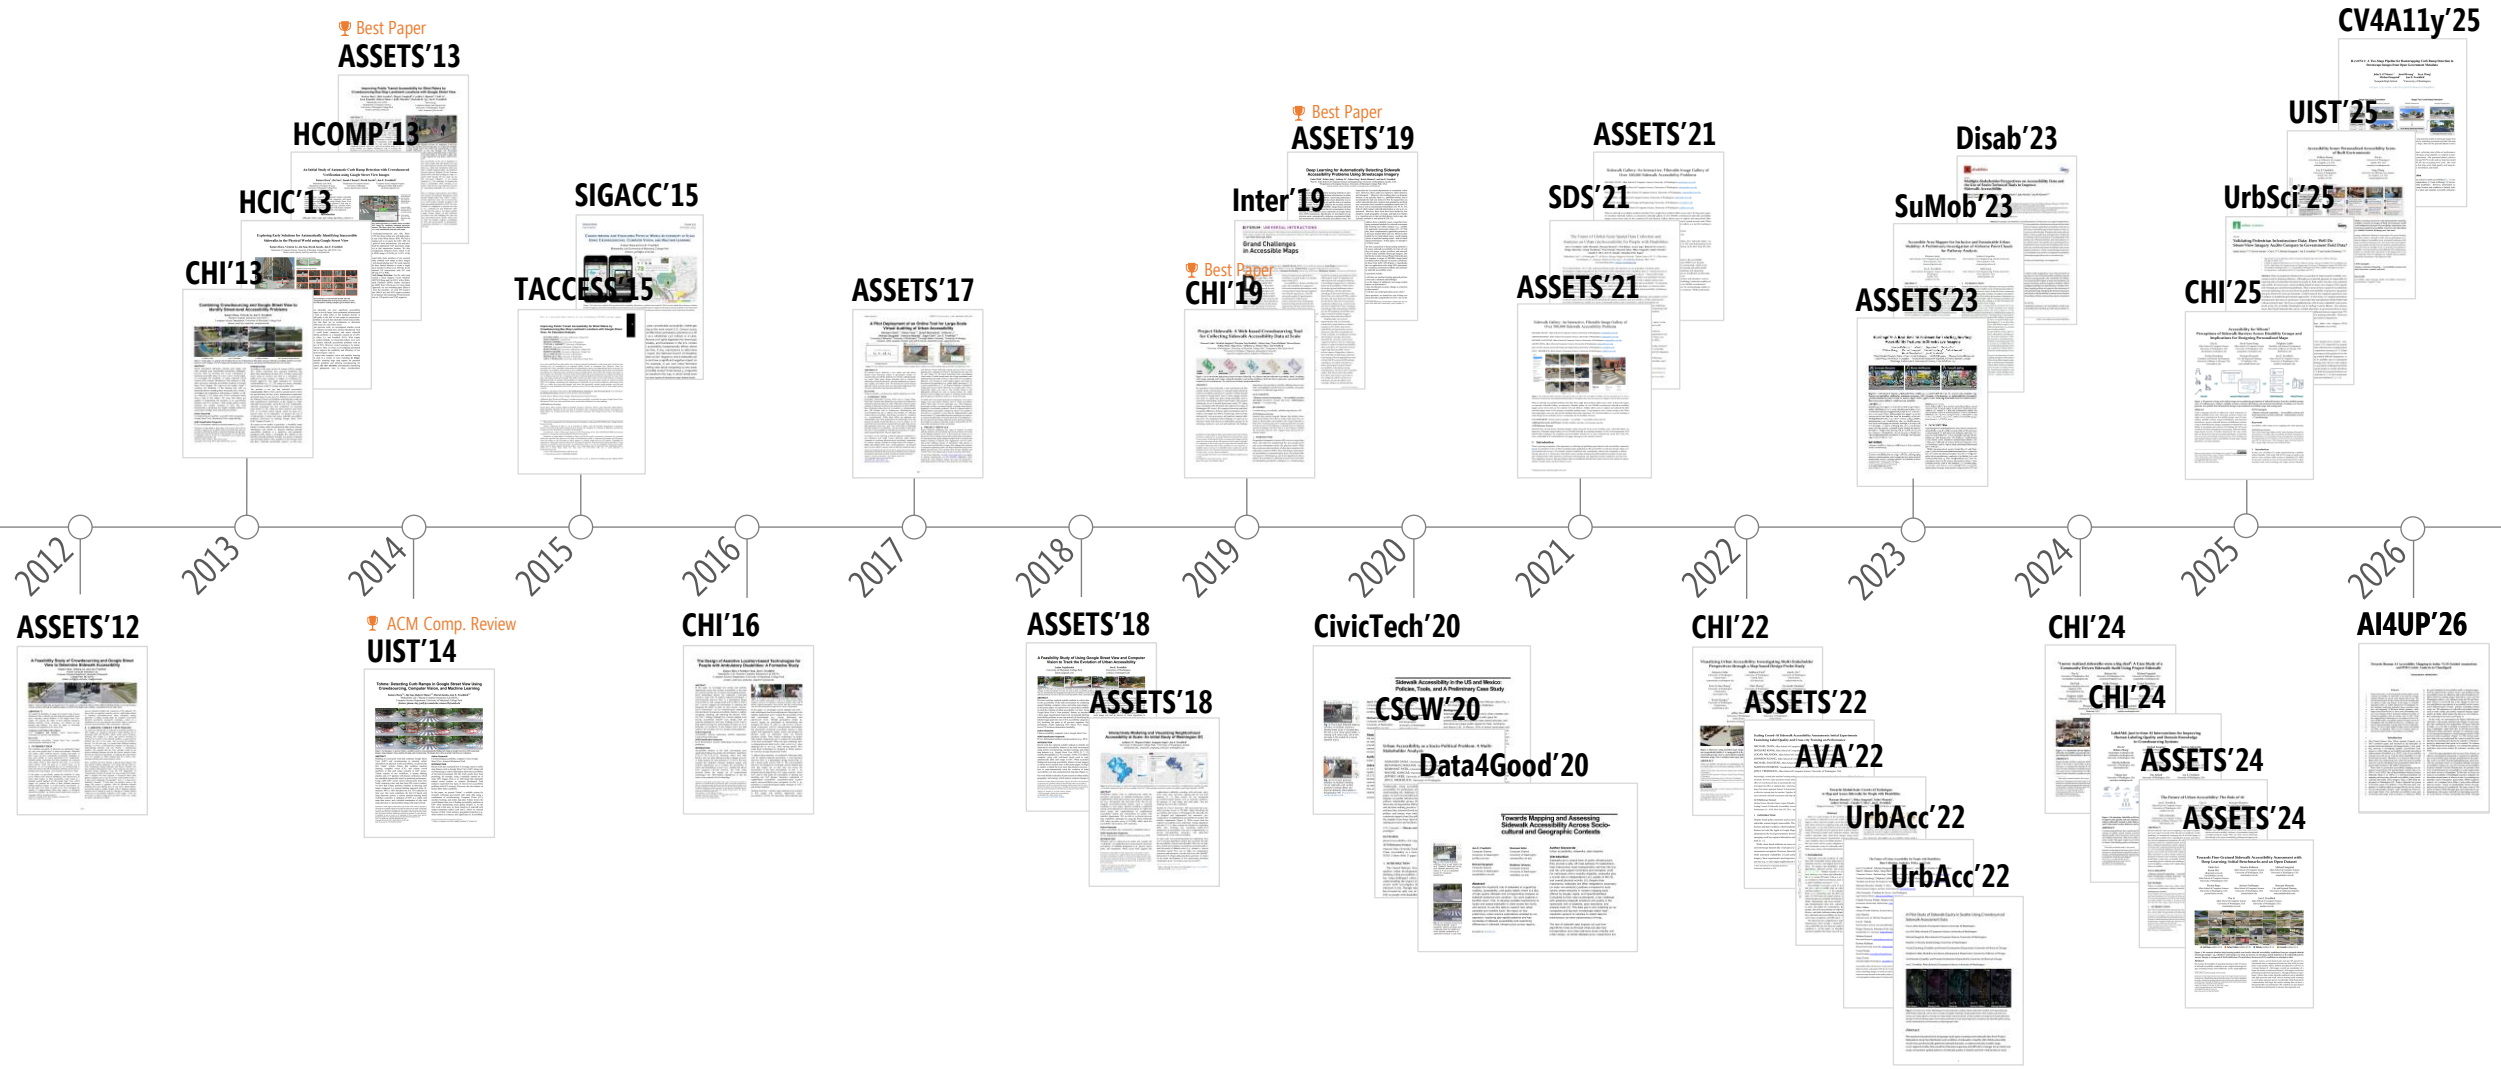
Figure 1: In this proposal, we combine crowdsourcing and computer vision to semi-automatically locate, identify, and assess accessibility problems in city streets, sidewalks, and building facades in online map imagery. The images above show crowd annotations from preliminary experiments on Mechanical Turk where untrained "turkers" were asked to find, label, and rate the severity of accessibility obstacles in Google Street View (GSV) images.

1 Introduction

According to the most recent US Census (2010), roughly 30.6 million adults have physical disabilities that affect their ambulatory activities [8]. 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) [8]. Despite comprehensive civil rights legislation for Americans with disabilities (*e.g.,* [1]), many city streets, sidewalks, and businesses in the US remain inaccessible [6,7]. The problem is not just that sidewalk

URBAN ACCESSIBILITY

RESEARCH TIMELINE



FORMATIVE & REFLECTIVE

Formative Studies

CHI'13



CHI'16



CSCW'20



CHI'22



CHI'25



Thought Pieces

Inter'19



SDS'21



AVA'22



UrbAcc'22



ASSETS'24



DATA COLLECTION TECHNIQUES

Manual Crowdsourcing

ASSETS'12



CHI'13



Best Paper

ASSETS'13



TACCESS'15



Best Paper

CHI'19



UrbSci'25



Automated Techniques

HCIC'13



HCOMP'13



Best Paper

ASSETS'19



SuMob'23



ASSETS'24



CV4A11y'25



Human-AI Techniques

ACM Comp. Review

UIST'14



ASSETS'18



ASSETS'21



ASSETS'22



ASSETS'23



CHI'24



UIST'25



AI4UP'26



ACCESS TOOLS

Vis Analytics & Mobility Tools

SIGACC'15



CHI'16



ASSETS'18



ASSETS'21



CHI'22



CHI'25



IMPACT STUDIES

Urban Science

CivicTech'20



UrbAcc'22



UrbSci'25



AI4UP'26



Deployment Studies

ASSETS'17

Best Paper

CHI'19

Data4Good'20

Disab'23

CHI'24

AI4UP'26

DATA COLLECTION TECHNIQUES

Manual Crowdsourcing

ASSETS'12 | CHI'13 | Best Paper ASSETS'13 | TACCESS'15 | Best Paper CHI'19 | UrbSci'25



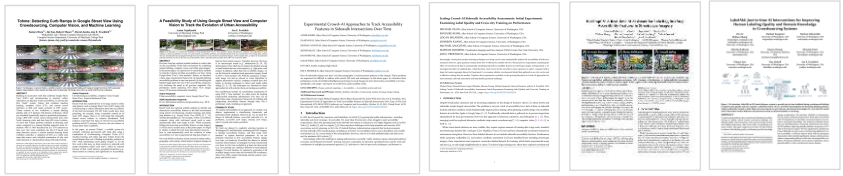
Automated Techniques

HCIC'13 | HCOMP'13 | Best Paper ASSETS'19 | SuMob'23 | ASSETS'24 | CV4A11y'25



Human-AI Techniques

ACM Comp. Review UIST'14 | ASSETS'18 | ASSETS'21 | ASSETS'22 | ASSETS'23 | CHI'24



UIST'25 | AI4UP'26



How can we...

Use AI techniques to improve **data collection efficiency**, **quality**, and **human performance**.

2 LABEL SIDEWALKS

Collect geo-located sidewalk assessments

Labeling Mission Routes

Streetscape Imagery APIs

Google Street View

Meta Mapillary

Panoramax

iNovitas infra3D

Community Partners

Labeling Missions

AI Labeler

Custom Trained ConvNeXt V2

CURB RAMPS

MISSING RAMP

SURFACE PROB

OBSTACLE

NO SIDEWALK

CROSSWALK

PED SIGNAL

3 VALIDATE LABELS

Human-AI quality control

Human Validation UI

AI-Assisted Human Validation

Custom Trained DINOv2

AI Validation

Agree

Disagree

4 POSTPROCESSING

Prepare data for apps

Raw Labels

Remove poor-quality labelers

Geo-cluster proximal labels

Label Clusters

Can AI automatically find and label **sidewalk features** and **problems**?



RESEARCH

AUTO-LABELING

HCIC'13



Exploring Early Solutions for Automatically Identifying Inaccessible Sidewalks in the Physical World using Google Street View

Kotaro Hara, Victoria Le, Jin Sun, David Jacobs, Jon E. Froehlich

Department of Computer Science, University of Maryland, College Park, USA

{kotaro, jinsun, dwj, jonf}@cs.umd.edu, vicle@umd.edu

HCOMP'13



An Initial Study of Automatic Curb Ramp Detection with Crowdsourced Verification using Google Street View Images

Kotaro Hara[1], Jin Sun[2], Jonah Chazan[1], David Jacobs[2], Jon E. Froehlich[1]

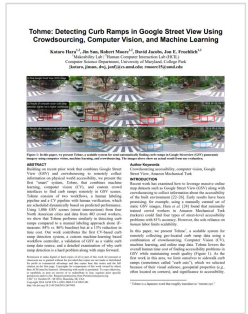
[1]Department of Computer Science, University of Maryland

[2]Computer Science Department, Montgomery Blair High School

{kotaro,jonf}@cs.umd.edu

UIST'14

ACM Comp. Review



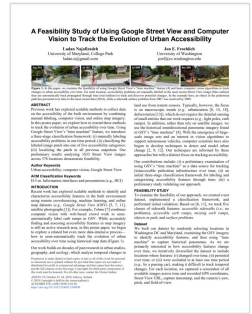
Tohme: Detecting Curb Ramps in Google Street View Using Crowdsourcing, Computer Vision, and Machine Learning

Kotaro Hara[1,2], Jin Sun, Robert Moore[1,2], David Jacobs, Jon E. Froehlich[1,2]

[1]Human-Computer Interaction Lab, [2]Computer Science Department, University of Maryland, College Park

{kotaro, jinsun, dwj, jonf}@cs.umd.edu rmoore1@umd.edu

ASSETS'18



A Feasibility Study of Using Google Street View and Computer Vision to Track the Evolution of Urban Accessibility

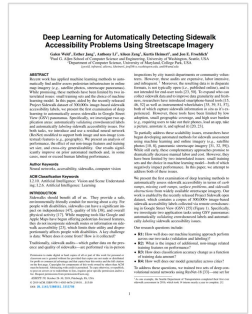
Lukas Negahdar[1], Jon E. Froehlich[2]

[1]University of Maryland, College Park

[2]University of Washington

ASSETS'19

Best Paper



Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery

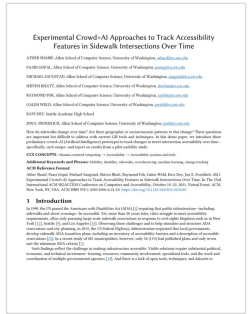
Galen Weld[1], Esther Jang[1], Anthony Li[1], Aileen Zeng[1], Kurtis Heimerl[2], and Jon E. Froehlich[1]

[1]Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

[2]Department of Computer Science, University of Maryland, College Park, USA

{galenweld, aileenz, kheimerl, jonf}@cs.washington.edu, esther@uw.edu

ASSETS'21



Experimental Crowd-AI Approaches to Track Accessibility Features in Sidewalk Intersections Over Time

AITHER HAMRE, Allen School of Computer Science, University of Washington, USA

PAMELYN FAL, Allen School of Computer Science, University of Washington, USA

MICHAEL SAUGSTAD, Allen School of Computer Science, University of Washington, USA

SHIVEN BHATT, Allen School of Computer Science, University of Washington, USA

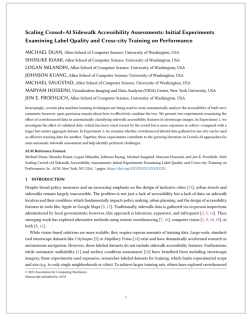
RAYMOND FOK, Allen School of Computer Science, University of Washington, USA

GALEN WELD, Allen School of Computer Science, University of Washington, USA

KAITLIN?, Seattle Academy High School

JON E. FROEHLICH, Allen School of Computer Science, University of Washington, USA

ASSETS'22



Scaling Crowd-AI Sidewalk Accessibility Assessments: Initial Experiments Examining Label Quality and Cross-city Training on Performance

MICHAEL FAN, Allen School of Computer Science, University of Washington, USA

SHYOSUKE KAMI, Allen School of Computer Science, University of Washington, USA

LOGAN MELANDES, Allen School of Computer Science, University of Washington, USA

JOHNSON KUANG, Allen School of Computer Science, University of Washington, USA

MICHAEL SAUGSTAD, Allen School of Computer Science, University of Washington, USA

MARYAM HOSSEINI, Visualization Imaging and Data Analysis (VIDA) Center, New York University, USA

JON E. FROEHLICH, Allen School of Computer Science, University of Washington, USA

ASSETS'23

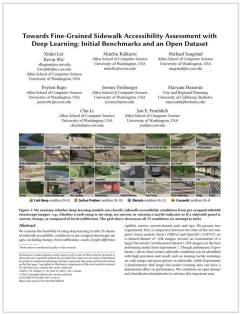


BusStopCV: A Real-time AI Assistant for Labeling Bus Stop Accessibility Features in Streetscape Imagery

Mineha Kulkarni[1], Michael Saugstad[1], Zhihan Zhang[2], Michael Saugstad[1], Yochai Eisenberg[3], Valerie Novack[3], Brett Chamberlain[3], Jon E. Froehlich[1]

[1]Allen School of Computer Science & Engineering, University of Washington [2]Paul G. Allen School of Computer Science & Engineering, University of Washington [3]Smithers Oasis Company

ASSETS'24

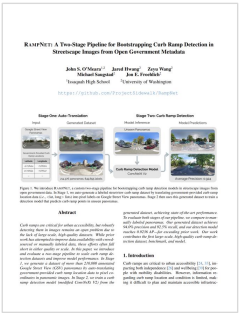


Towards Fine-Grained Sidewalk Accessibility Assessment with Deep Learning: Initial Benchmarks and an Open Dataset

Kinjal Gor[1], Mineha Kulkarni[1], Michael Saugstad[1], Jon E. Froehlich[1]

[1]Allen School of Computer Science & Engineering, University of Washington, USA

CV4A11y'25



RAMPNET: A Two-Stage Pipeline for Bootstrapping Curb Ramp Detection in Streetscape Images from Open Government Metadata

John S. O'Meara[1,2] Jared Hwang[1] Zeyu Wang[1]

Michael Saugstad[1] Jon E. Froehlich[1]

[1]Issaquah High School [2]University of Washington

https://github.com/ProjectSidewalk/SvmNet

RESEARCH

AUTO-LABELING



GALEN WELD

Best Paper

ASSETS'18

A Feasibility Study of Using Google Street View and Computer Vision to Track the Evolution of Urban Accessibility

Ladan Najafizadeh
University of Maryland, College Park
ladan.n@gmail.com

Jon E. Froehlich
University of Washington
jonf@cs.washington.edu



Figure 1. In this paper, we examine the feasibility of using Google Street View's "time machine" feature [4] and basic computer vision algorithms to track changes in urban accessibility over time. For each location, accessibility problems are manually labeled in the most recent Street View image (blue outline) then are automatically back propagated through time (red outlines) to track and discover potential changes. In the example here, an object in the pedestrian path has persisted over time to the most recent data (2014), while a sidewalk surface problem from 2007 was resolved by 2009.

ABSTRACT

Previous work has explored scalable methods to collect data on the accessibility of the built environment by combining manual labeling, computer vision, and online map imagery. In this poster paper, we explore how to extend these methods to track the evolution of urban accessibility over time. Using Google Street View's "time machine" feature, we introduce a three-stage classification framework: (i) manually labeling accessibility problems in one time period; (ii) classifying the labeled image patch into one of five accessibility categories; (iii) localizing the patch in all previous snapshots. Our preliminary results analyzing 1633 Street View images across 376 locations demonstrate feasibility.

Author Keywords

Urban accessibility; computer vision; Google Street View

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI)

INTRODUCTION

Recent work has explored scalable methods to identify and characterize accessibility features in the built environment using remote crowdsourcing, machine learning, and online map datasets (*e.g., Google Street View* (GSV) [5, 7, 11], satellite photographs [1]). For example, *Tohme* [7] combines computer vision with web-based crowd work to semi-automatically label curb ramps in GSV. While accurately finding and assessing accessibility features in map imagery is still an active research area, in this poster paper, we begin to explore a related but even more data-intensive process—how to semi-automatically track the evolution of urban accessibility over time using historical map data (Figure 1).

Our work builds on decades of past research in urban studies, geography, and ecology, which analyze temporal changes in land use from remote sensors. Typically, however, the focus is on macroscopic trends (*e.g.,* urbanization [8, 14, 18], deforestation [13]), which do not require the detailed sensing of small entities that our work requires (*e.g.,* light poles, curb ramps). In addition, rather than rely on satellite images, we use the historical omnidirectional panoramic imagery found in GSV's "time machine" [4]. With the emergence of large-scale image sets and an interest in vision algorithms to support autonomous vehicles, computer scientists have also begun to develop techniques to detect and model urban change [2, 9, 12]. Our techniques are informed by these approaches but with a distinct focus on tracking accessibility.

Our contributions include: (i) a preliminary examination of using GSV's "time machine" as a data source for tracking (in)accessible pedestrian infrastructure over time; (ii) an initial three-stage classification framework for labeling and categorizing accessibility features through time; (iii) a preliminary study validating our approach.

FEASIBILITY STUDY

To examine the feasibility of our approach, we created a test dataset, implemented a classification framework, and performed initial validation. Based on [6, 11], we track five classes of sidewalk features: *accessible sidewalks* (*i.e.*, no problems), *accessible curb ramps*, *missing curb ramps*, *objects in path*, and *surface problems*.

Dataset

We built our dataset by randomly selecting locations in Washington DC and Maryland, examining the GSV imagery to identify accessibility features, and then using "time machine" to capture historical panoramas. As we are primarily interested in how accessibility features change over time, we iteratively diversified the dataset to include locations where features: (i) changed over time; (ii) persisted over time; or (iii) were occluded in at least one time period (*e.g.*, by a passing car), making it difficult to track temporal changes. For each location, we captured a screenshot of all available images across time and recorded GPS coordinates, Street View URL, capture timestamp, and the camera's yaw, pitch, and field-of-view.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

ASSETS '18, October 22–24, 2018, Galway, Ireland

© 2018 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-5650-3/18/10.
https://doi.org/10.1145/3234695.3240999

ASSETS'19

Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery

Galen Weld[1], Esther Jang[1], Anthony Li[2], Aileen Zeng[1], Kurtis Heimerl[1], and Jon E. Froehlich[1]
[1]Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, USA
[2]Department of Computer Science, University of Maryland, College Park, USA
{gweld, infrared, aileenz, kheimerl, jonf}@cs.washington.edu, amli@umd.edu

ABSTRACT

Recent work has applied machine learning methods to automatically find and/or assess pedestrian infrastructure in online map imagery (*e.g.,* satellite photos, streetscape panoramas). While promising, these methods have been limited by two interrelated issues: small training sets and the choice of machine learning model. In this paper, aided by the recently released Project Sidewalk dataset of 300,000+ image-based sidewalk accessibility labels, we present the first examination of deep learning to automatically assess sidewalks in Google Street View (GSV) panoramas. Specifically, we investigate two application areas: automatically *validating* crowdsourced labels and automatically *labeling* sidewalk accessibility issues. For both tasks, we introduce and use a residual neural network (ResNet) modified to support both image and non-image (contextual) features (*e.g.,* geography). We present an analysis of performance, the effect of our non-image features and training set size, and cross-city generalizability. Our results significantly improve on prior automated methods and, in some cases, meet or exceed human labeling performance.

Author Keywords

Neural networks, accessibility, sidewalks, computer vision

ACM Classification Keywords

I.2.10. Artificial Intelligence: Vision and Scene Understanding; I.2.6. Artificial Intelligence: Learning

INTRODUCTION

Sidewalks should benefit all of us. They provide a safe, environmentally-friendly conduit for moving about a city. For people with disabilities, sidewalks can have a significant impact on independence [47], quality of life [38], and overall physical activity [17]. While mapping tools like Google and Apple Maps have begun offering pedestrian-focused features, they do not incorporate sidewalk routes or information on sidewalk accessibility [23], which limits their utility and disproportionately affects people with disabilities. A key challenge is data: Where does it come from? How is it collected?

Traditionally, sidewalk audits—which gather data on the presence and quality of sidewalks—are performed via in-person inspections by city transit departments or community volunteers. However, these audits are expensive, labor intensive, and infrequent. [1] Moreover, the resulting data is in disparate formats, is not typically open (*i.e.,* published online), and is not intended for end-user tools [23, 50]. To expand who can collect sidewalk data and to improve data granularity and freshness, researchers have introduced smartphone-based tools [15, 46, 52] as well as instrumented wheelchairs [35, 39, 51, 57], both of which capture sidewalk information *in situ* as it's experienced. However, these tools have been limited by low adoption, small geographic coverage, and high user burden (*e.g.,* requiring users to take out their phones, load an app, take a picture, annotate it, and upload it) [20, 23].

To partially address these scalability issues, researchers have begun developing automated methods for sidewalk assessment using machine learning and online imagery (*e.g.,* satellite photos [10, 8], panoramic streetscape imagery [31, 32, 59]). While still early, these complementary approaches promise to dramatically decrease manual labor and cost. However, they have been limited by two interrelated issues: small training sets and the choice in machine learning model—both of which negatively impact performance. In this paper, we attempt to address both of these issues.

We present the first examination of deep learning methods to automatically assess sidewalk accessibility in terms of *curb ramps*, *missing curb ramps*, *surface problems*, and *sidewalk obstructions* from widely available streetscape imagery. Our work is enabled by the recently released Project Sidewalk open dataset, which contains a corpus of 300,000+ image-based sidewalk accessibility labels collected via remote crowdsourcing in Google Street View (GSV) [55] (Figure 1). Specifically, we investigate two application tasks using GSV panoramas: automatically *validating* crowdsourced labels and automatically *labeling* sidewalk accessibility issues.

Our research questions include:

- **R1:** How well does our machine learning approach perform across our two tasks (validation and labeling)?
- **R2:** What is the impact of additional, non-image related training features on performance?
- **R3:** How does classification accuracy change as a function of training data amount?
- **R4:** How well does our model generalize across cities?

To address these questions, we trained two sets of deep convolutional neural networks using ResNet-18 [33]—one set for

[1]As one example, the Seattle Department of Transportation completed their first ever sidewalk assessment in 2016, which took 14 interns nearly a year to complete. [1]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASSETS '19, October 28–30, 2019, Pittsburgh, PA, USA

© 2019 ACM. ISBN 978-1-4503-6676-2/19/10…$15.00

DOI: 10.1145/3308561.3353798

ASSETS'21

Experimental Crowd+AI Approaches to Track Accessibility Features in Sidewalk Intersections Over Time

ATHER SHARIF, Allen School of Computer Science, University of Washington, asharif@cs.uw.edu

PAARI GOPAL, Allen School of Computer Science, University of Washington, paarig@cs.uw.edu

MICHAEL SAUGSTAD, Allen School of Computer Science, University of Washington, saugstad@cs.uw.edu

SHIVEN BHATT, Allen School of Computer Science, University of Washington, shivbat@cs.uw.edu

RAYMOND FOK, Allen School of Computer Science, University of Washington, rayfok@cs.uw.edu

GALEN WELD, Allen School of Computer Science, University of Washington, gweld@cs.uw.edu

KAVI DEY, Seattle Academy High School

JON E. FROEHLICH, Allen School of Computer Science, University of Washington, jonf@cs.uw.edu

How do sidewalks change over time? Are there geographic or socioeconomic patterns to this change? These questions are important but difficult to address with current GIS tools and techniques. In this demo paper, we introduce three preliminary crowd+AI (Artificial Intelligence) prototypes to track changes in street intersection accessibility over time—specifically, curb ramps—and report on results from a pilot usability study.

CCS CONCEPTS • Human-centered computing → Accessibility → Accessibility systems and tools

Additional Keywords and Phrases: Mobility, disability, sidewalks, crowdsourcing, machine learning, change tracking

ACM Reference Format:

Ather Sharif, Paari Gopal, Michael Saugstad, Shiven Bhatt, Raymond Fok, Galen Weld, Kavi Dey, Jon E. Froehlich. 2021. Experimental Crowd+AI Approaches to Track Accessibility Features in Sidewalk Intersections Over Time. In The 23rd International ACM SIGACCESS Conference on Computers and Accessibility, October 18–22, 2021, Virtual Event. ACM, New York, NY, USA. ACM ISBN 978-1-4503-8306-6/21/10. https://doi.org/10.1145/3441852.3476549

1 Introduction

In 1990, the US passed the *Americans with Disabilities Act* (ADA) [1] requiring that public infrastructure—including sidewalks and street crossings—be accessible. Yet, more than 30 years later, cities struggle to meet accessibility requirements, often only pursuing large-scale sidewalk renovations in response to civil rights litigation such as in New York [12], Seattle [9], and Los Angeles [18]. Observing these challenges and to help stimulate and structure ADA renovations and city planning, in 2015, the *US Federal Highway Administration* requested that local governments develop sidewalk ADA transition plans, including an inventory of accessibility barriers and a description of accessible renovations [25]. In a recent study of 401 municipalities, however, only 54 (13%) had published plans and only seven met the minimum ADA criteria [7].

Such findings reflect the challenge in making infrastructure accessible. Viable solutions require substantial political, economic, and technical investment—training, resources, community involvement, specialized tools, and the work and coordination of multiple governmental agencies [19]. And there is a lack of open tools, techniques, and datasets to

ASSET

Scaling Crowd+
Examining Lab

MICHAEL DUAN,
SHOSUKE KIAMI,
LOGAN MILANDI
JOHNSON KUANG
MICHAEL SAUGS
MARYAM HOSSEI
JON E. FROEHLIC

Increasingly, crowds plus
ronments; however, open
effect of crowdsourced da
investigate the effect of va
larger but noisier aggregat
as effective training data
semi-automatic sidewalk

ACM Reference Format
Michael Duan, Shosuke K
Scaling Crowd+AI Sidew
Performance. In . ACM, N

1 INTRODUCTION

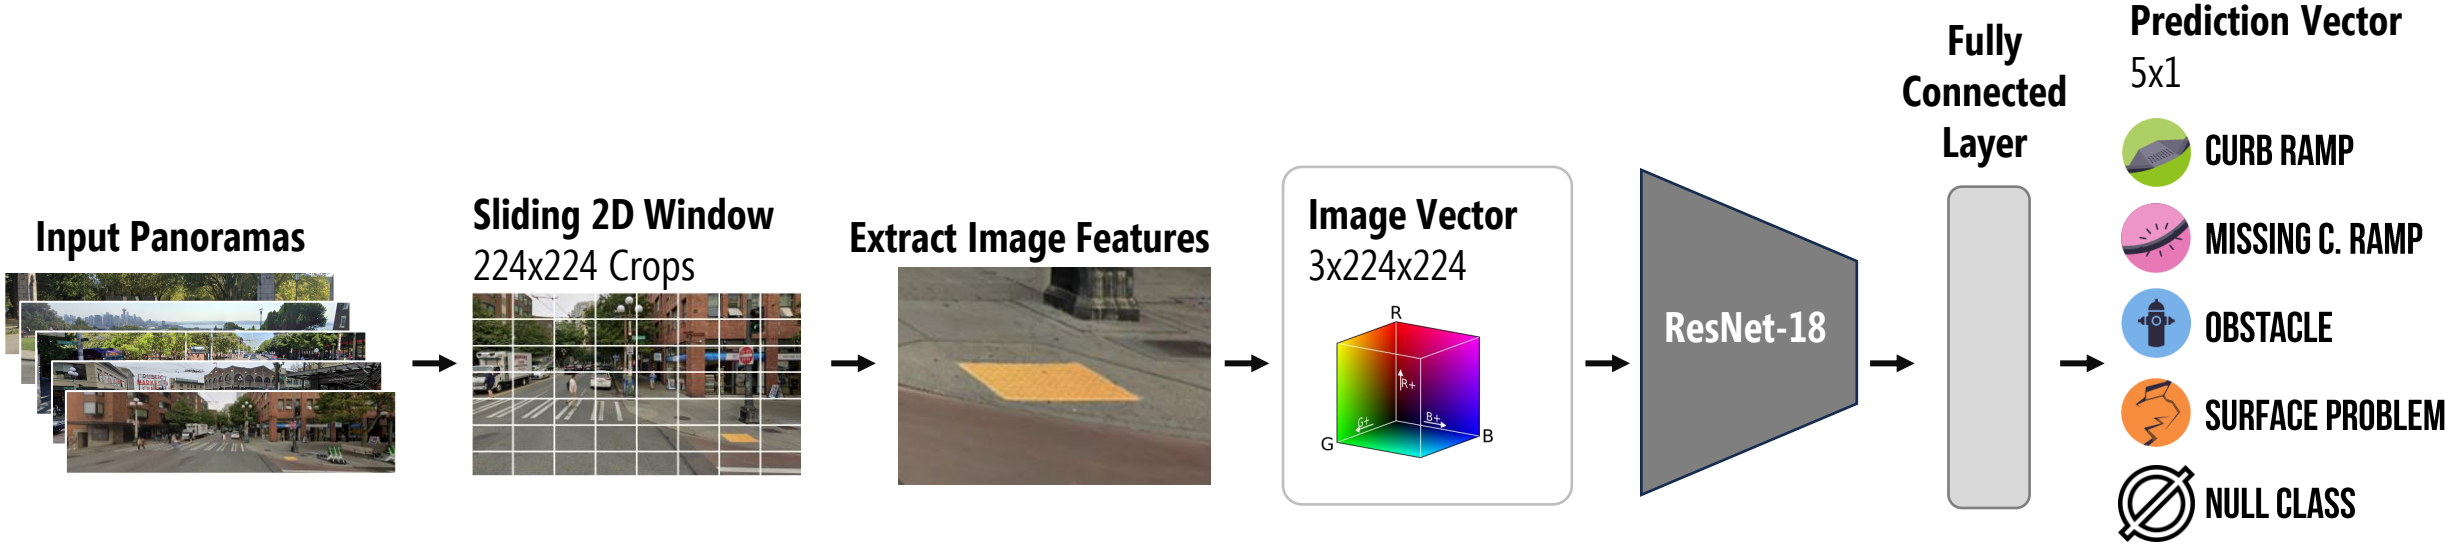
Despite broad policy
sidewalks remain large
location and their condi
features in tools like A
administered by local
emerging work has exp
both [8, 11].

While vision-based
ized streetscape dataset
autonomous navigatio
while automatic walk
imagery, these experim
and size (e.g., to only si

© 2022 Association for Comp
Manuscript submitted to AC

AUTO-LABELING

MODIFIED RESNET-18 ARCHITECTURE

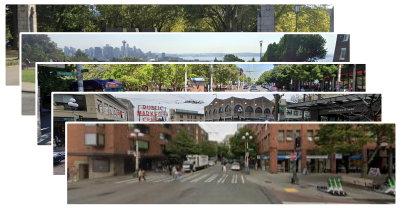


Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

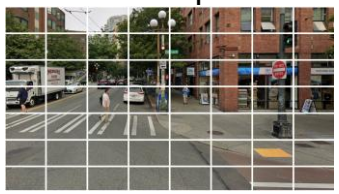
AUTO-LABELING

MODIFIED RESNET-18 ARCHITECTURE

Input Panoramas



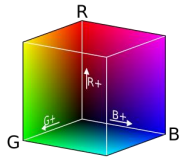
Sliding 2D Window
224x224 Crops



Extract Image Features



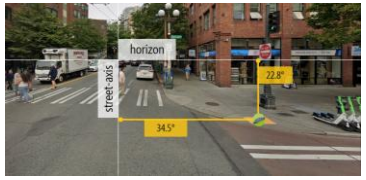
Image Vector
3x224x224



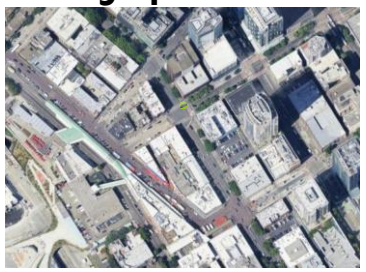
ResNet-18

Fully Connected Layer

Extract Positional Features



Extract Geographic Features



Metadata Vector
12x1

```
Depth
<x,y> pixel pos
Street dir
Block pos
Distance from
city center
…
```

Prediction Vector
5x1

- CURB RAMP
- MISSING C. RAMP
- OBSTACLE
- SURFACE PROBLEM
- NULL CLASS

Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

AUTO-LABELING

TRAINING DATASET

46.5K

PANORAMAS





120K

CURB RAMPS





15.7K

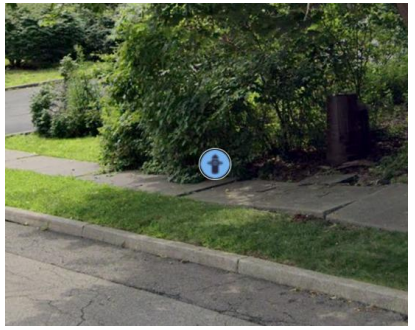
MISSING C. RAMPS





17.5K

OBSTACLES





7.0K

SURFACE PROBS.

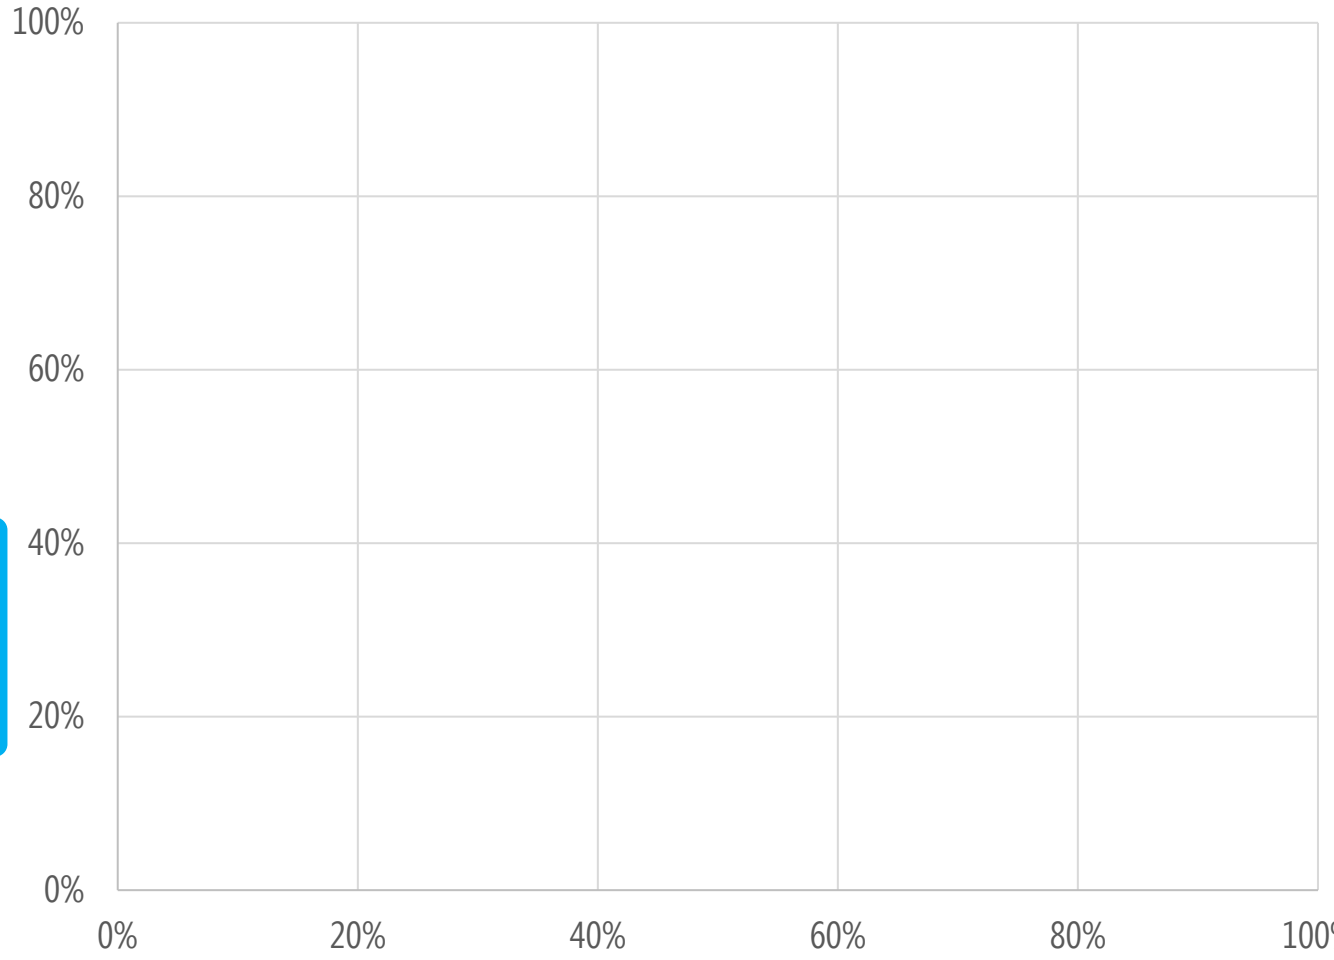




Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

AUTO-LABELING

PERFORMANCE RESULTS



Precision

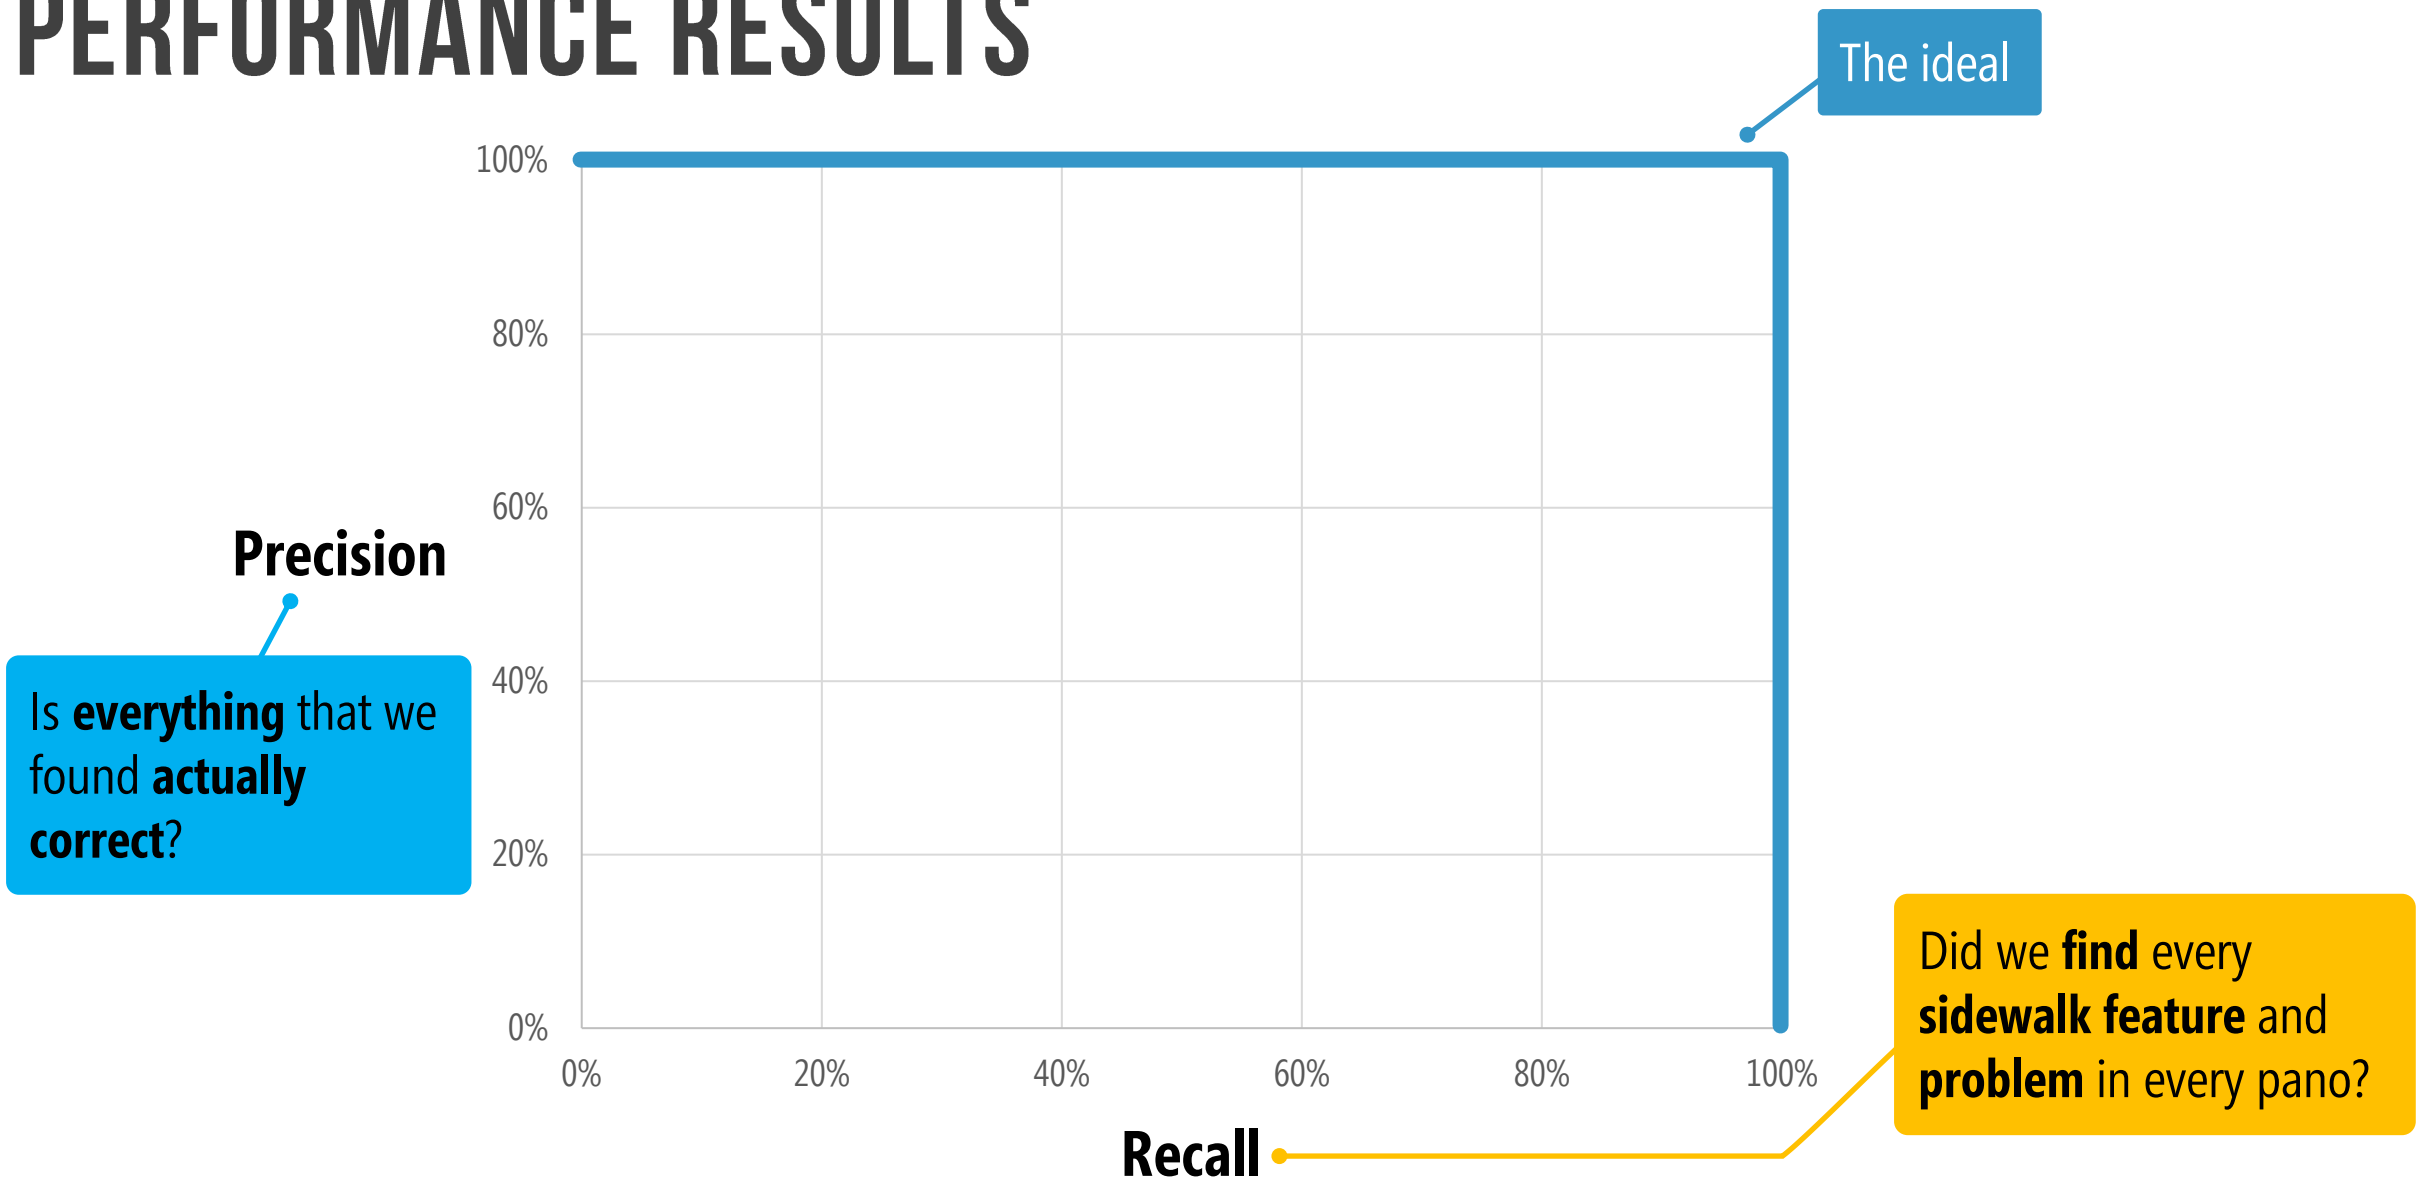
Is **everything** that we found **actually correct**?

Recall

Did we **find** every **sidewalk feature** and **problem** in every pano?

AUTO-LABELING

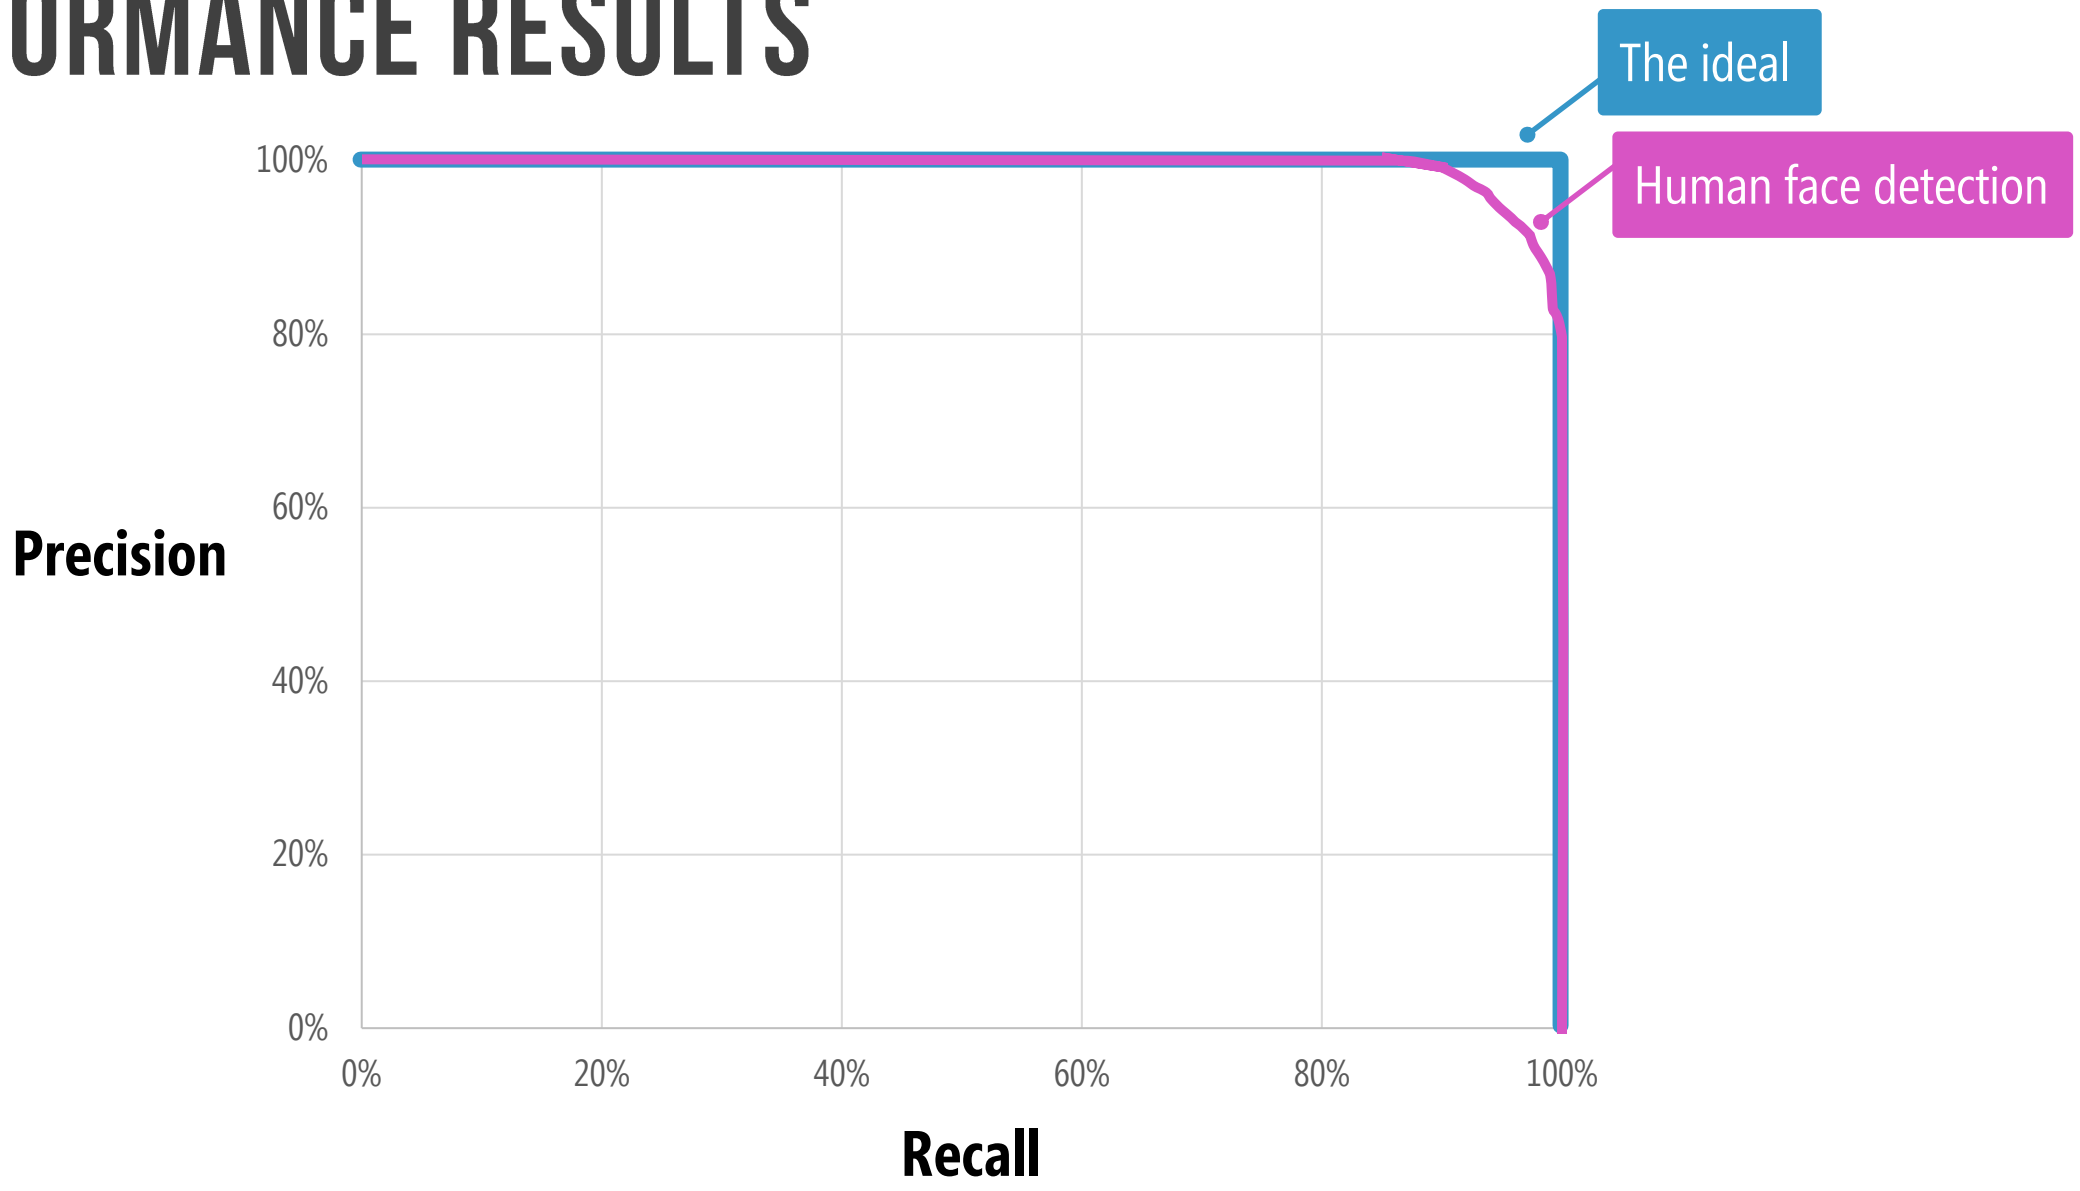
PERFORMANCE RESULTS



Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

AUTO-LABELING

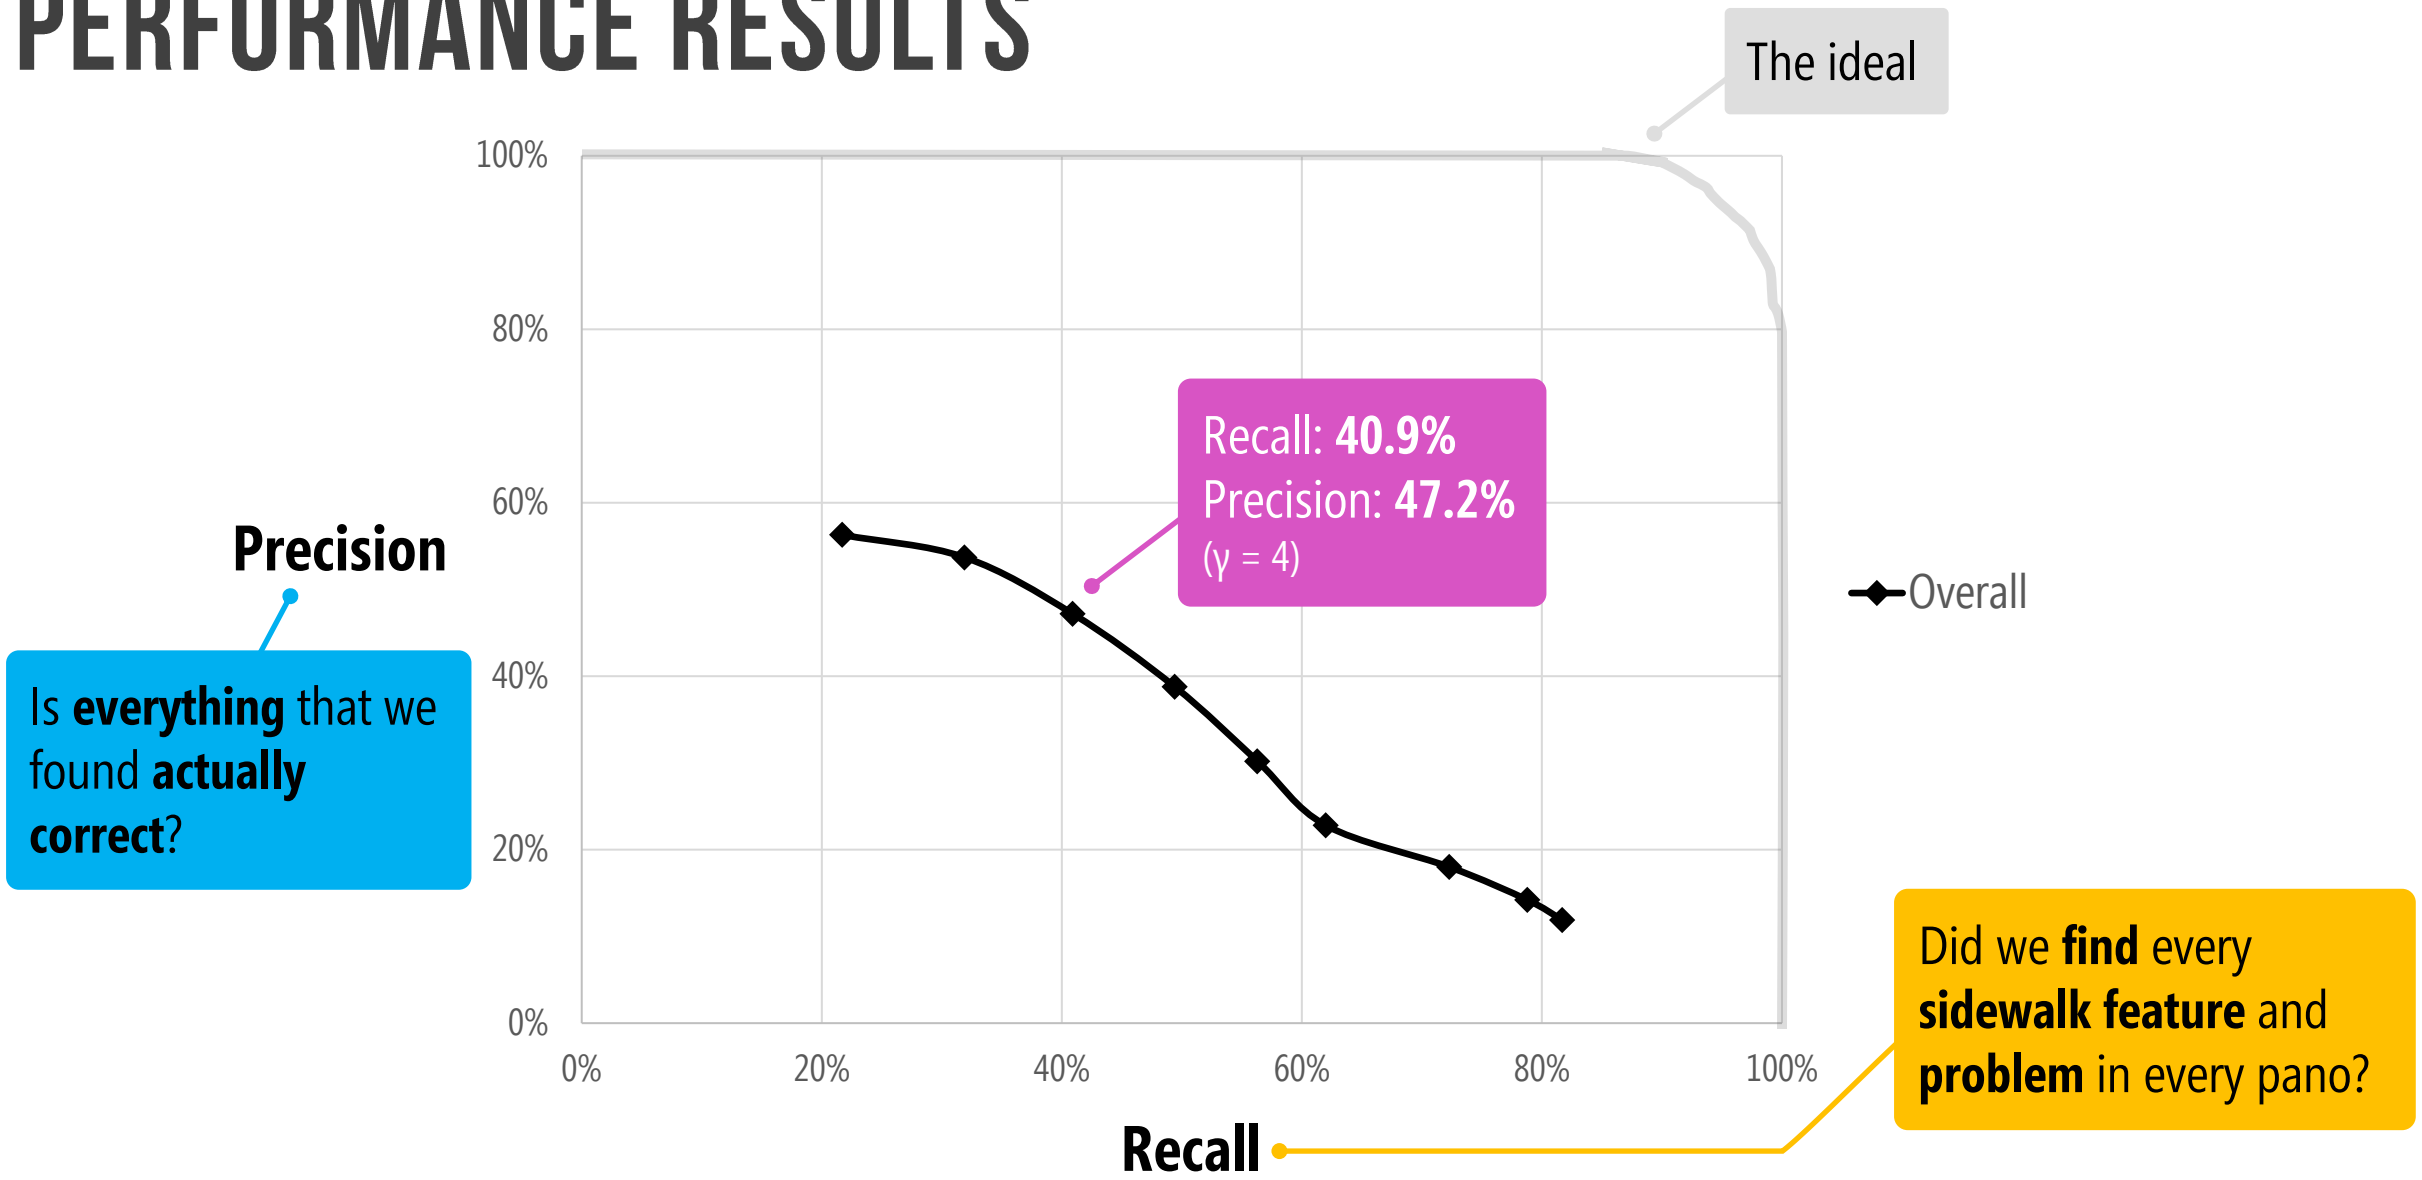
PERFORMANCE RESULTS



Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

AUTO-LABELING

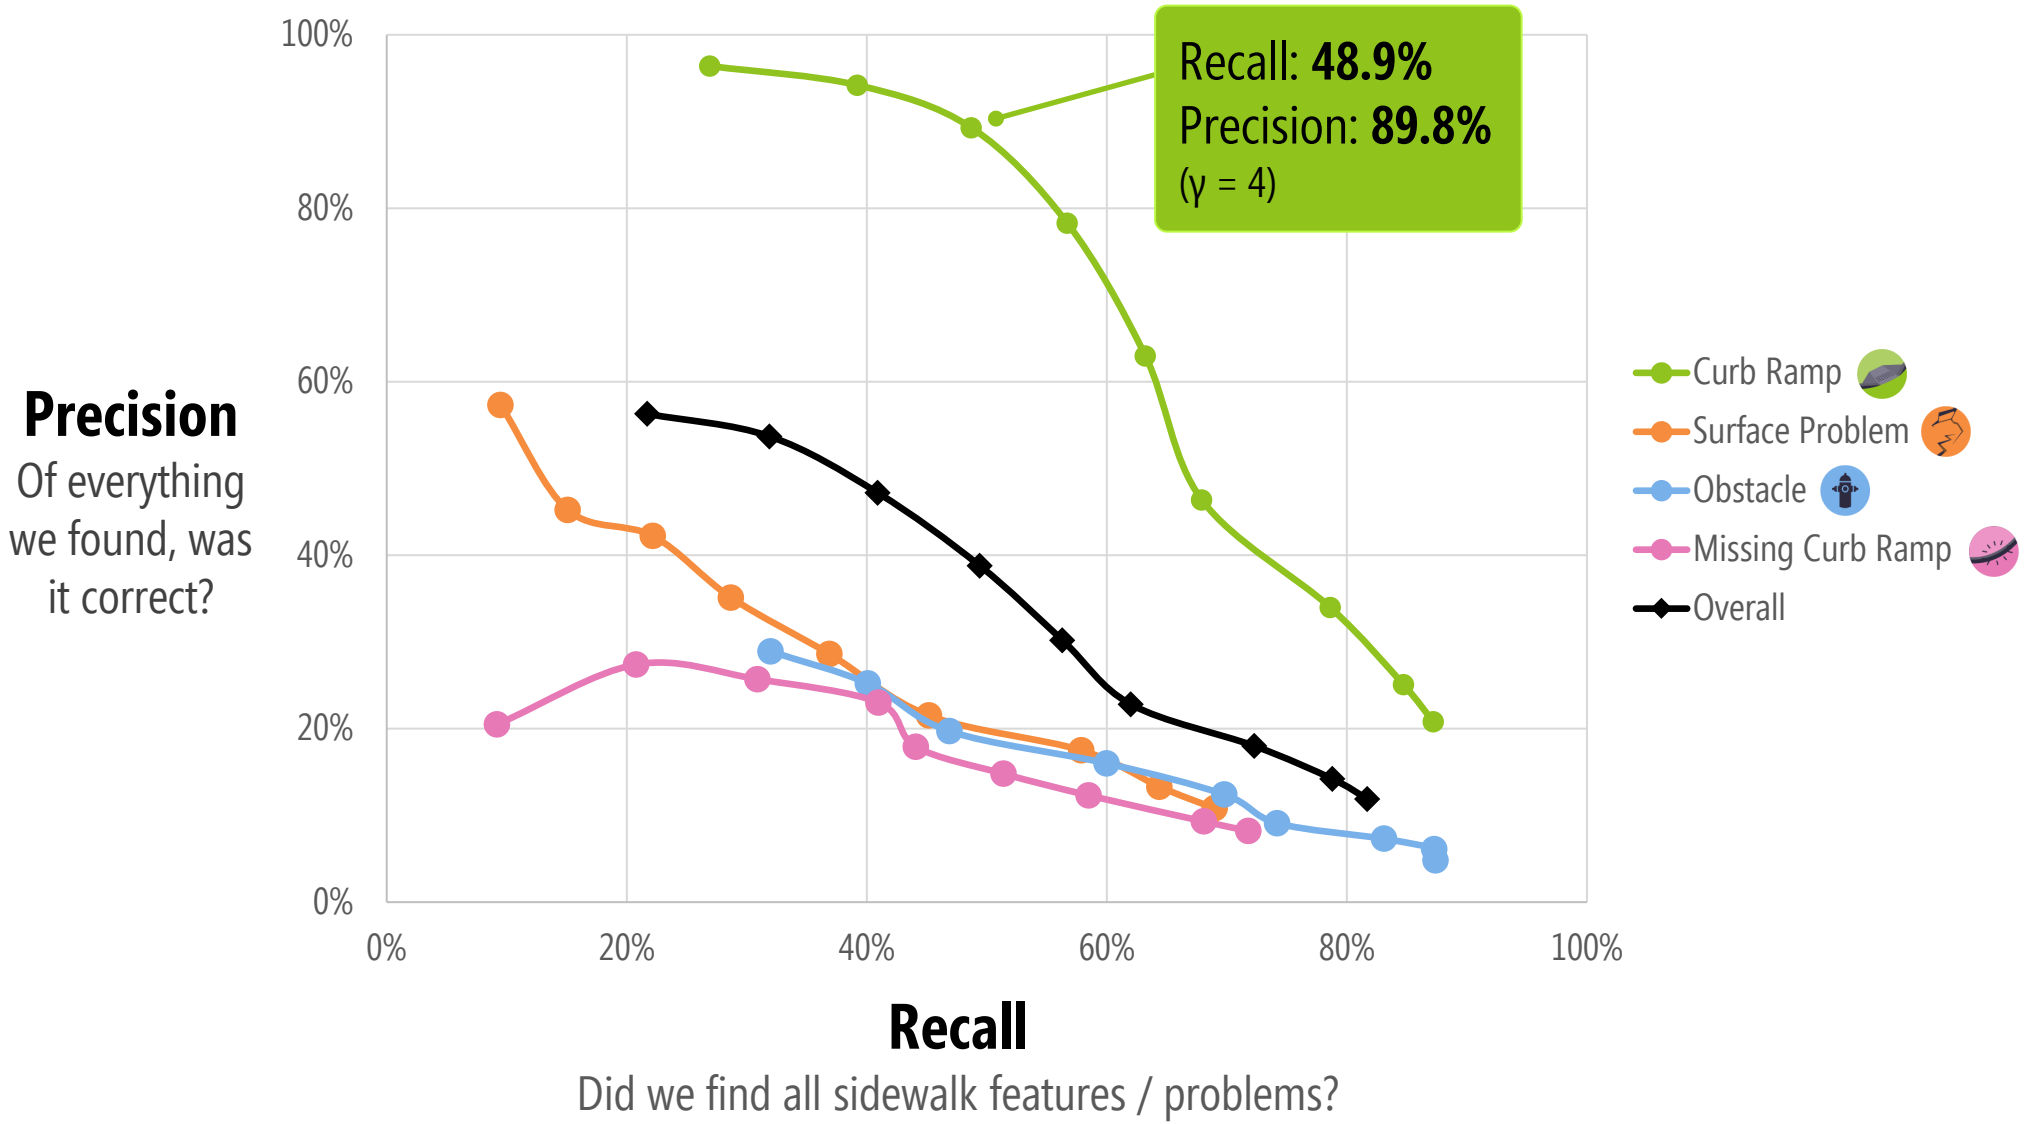
PERFORMANCE RESULTS



Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

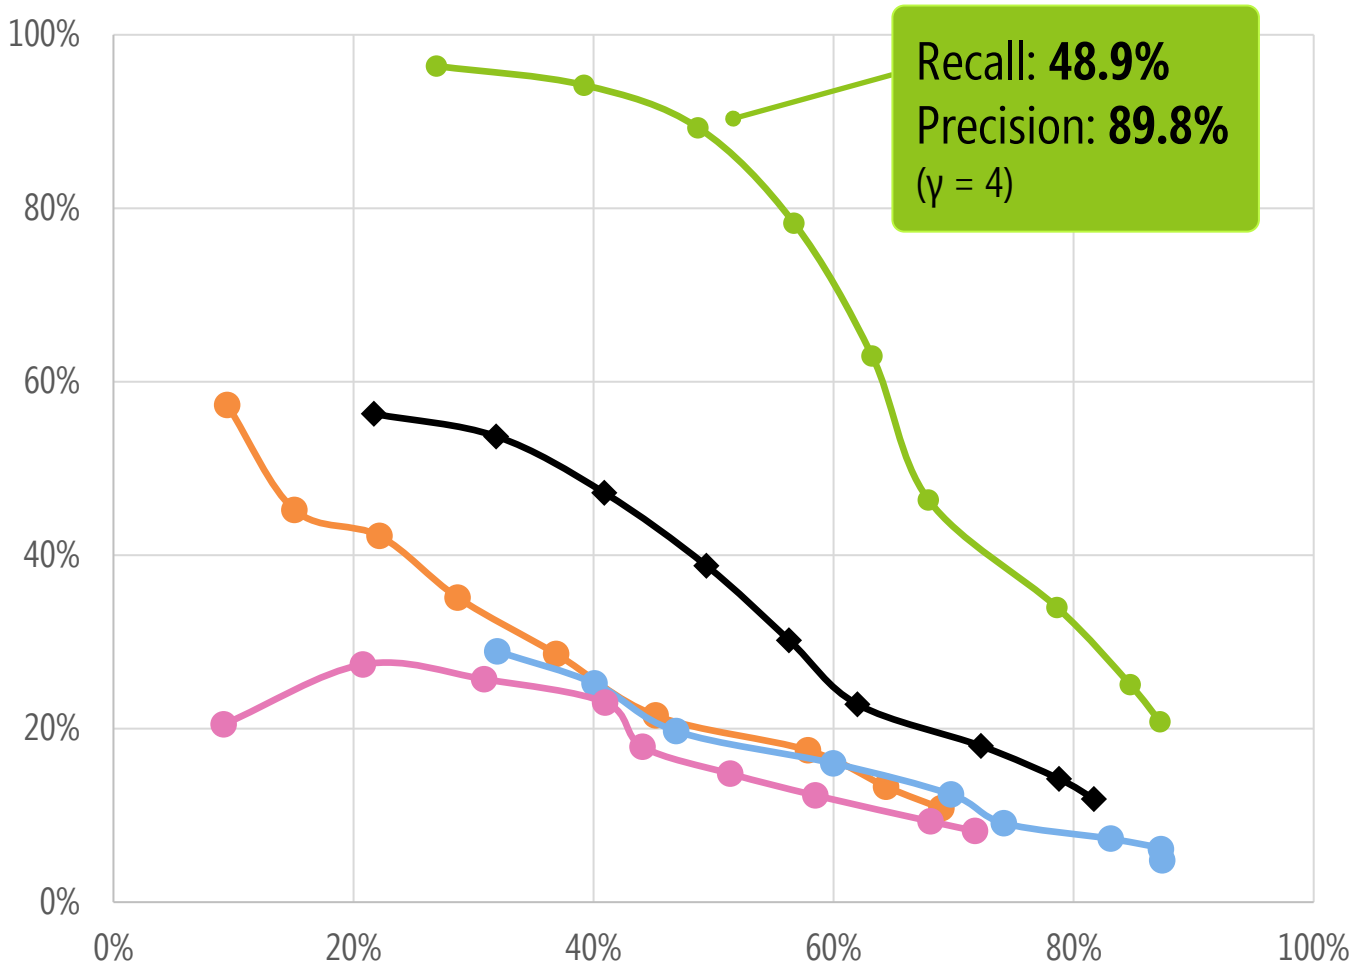
AUTO-LABELING

PERFORMANCE RESULTS BY LABEL TYPE



AUTO-LABELING

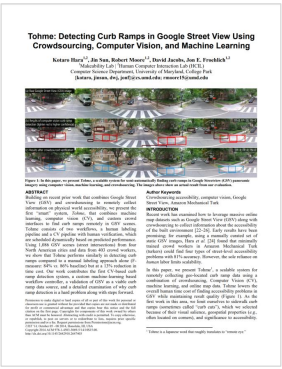
PERFORMANCE RESULTS BY LABEL TYPE



Precision
Of everything we found, was it correct?

Recall
Did we find all sidewalk features / problems?

UIST'14

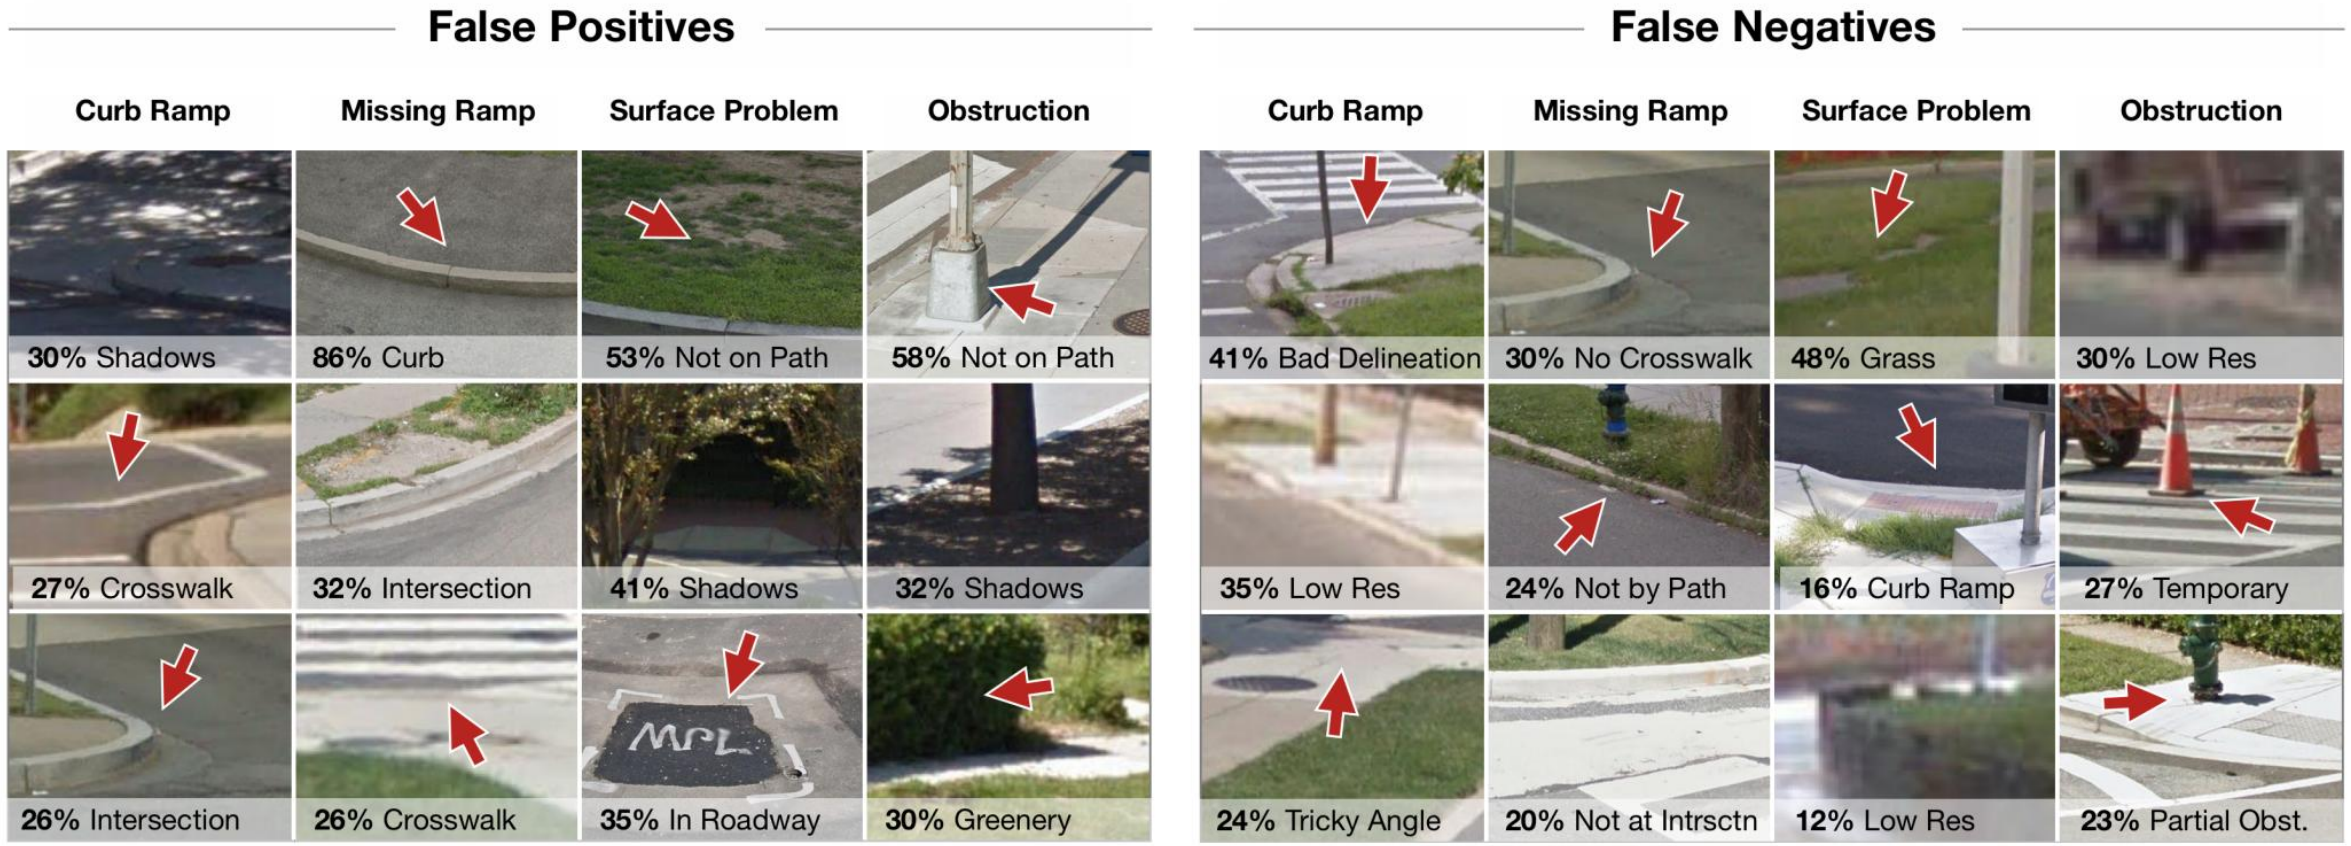

Tohme: Detecting Curb Ramps in Google Street View Using Crowdsourcing, Computer Vision, and Machine Learning

CURB RAMP PERFORMANCE

	Recall	Precision
AI (SVM)	67%	26%
Humans	85%	89%

AUTO-LABELING

COMMON FAILURES



Weld *et al.,* Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery, ASSETS'19

RESEARCH

AUTO-LABELING



GALEN WELD

Best Paper

ASSETS'18

A Feasibility Study of Using Google Street View and Computer Vision to Track the Evolution of Urban Accessibility

Ladan Najafizadeh
University of Maryland, College Park
ladan.n@gmail.com

Jon E. Froehlich
University of Washington
jonf@cs.washington.edu



Figure 1. In this paper, we examine the feasibility of using Google Street View's "time machine" feature [4] and basic computer vision algorithms to track changes in urban accessibility over time. For each location, accessibility problems are manually labeled in the most recent Street View image (blue outline) then are automatically back propagated through time (red outlines) to track and discover potential changes. In the example here, an object in the pedestrian path has persisted over time to the most recent data (2014), while a sidewalk surface problem from 2007 was resolved by 2009.

ABSTRACT

Previous work has explored scalable methods to collect data on the accessibility of the built environment by combining manual labeling, computer vision, and online map imagery. In this poster paper, we explore how to extend these methods to track the evolution of urban accessibility over time. Using Google Street View's "time machine" feature, we introduce a three-stage classification framework: (i) manually labeling accessibility problems in one time period; (ii) classifying the labeled image patch into one of five accessibility categories; (iii) localizing the patch in all previous snapshots. Our preliminary results analyzing 1633 Street View images across 376 locations demonstrate feasibility.

Author Keywords

Urban accessibility; computer vision; Google Street View

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI)

INTRODUCTION

Recent work has explored scalable methods to identify and characterize accessibility features in the built environment using remote crowdsourcing, machine learning, and online map datasets (*e.g., Google Street View* (GSV) [5, 7, 11], satellite photographs [1]). For example, *Tohme* [7] combines computer vision with web-based crowd work to semi-automatically label curb ramps in GSV. While accurately finding and assessing accessibility features in map imagery is still an active research area, in this poster paper, we begin to explore a related but even more data-intensive process—how to semi-automatically track the evolution of urban accessibility over time using historical map data (Figure 1).

Our work builds on decades of past research in urban studies, geography, and ecology, which analyze temporal changes in land use from remote sensors. Typically, however, the focus is on macroscopic trends (*e.g.,* urbanization [8, 14, 18], deforestation [13]), which do not require the detailed sensing of small entities that our work requires (*e.g.,* light poles, curb ramps). In addition, rather than rely on satellite images, we use the historical omnidirectional panoramic imagery found in GSV's "time machine" [4]. With the emergence of large-scale image sets and an interest in vision algorithms to support autonomous vehicles, computer scientists have also begun to develop techniques to detect and model urban change [2, 9, 12]. Our techniques are informed by these approaches but with a distinct focus on tracking accessibility.

Our contributions include: (i) a preliminary examination of using GSV's "time machine" as a data source for tracking (in)accessible pedestrian infrastructure over time; (ii) an initial three-stage classification framework for labeling and categorizing accessibility features through time; (iii) a preliminary study validating our approach.

FEASIBILITY STUDY

To examine the feasibility of our approach, we created a test dataset, implemented a classification framework, and performed initial validation. Based on [6, 11], we track five classes of sidewalk features: *accessible sidewalks* (*i.e.*, no problems), *accessible curb ramps*, *missing curb ramps*, *objects in path*, and *surface problems*.

Dataset

We built our dataset by randomly selecting locations in Washington DC and Maryland, examining the GSV imagery to identify accessibility features, and then using "time machine" to capture historical panoramas. As we are primarily interested in how accessibility features change over time, we iteratively diversified the dataset to include locations where features: (i) changed over time; (ii) persisted over time; or (iii) were occluded in at least one time period (*e.g.*, by a passing car), making it difficult to track temporal changes. For each location, we captured a screenshot of all available images across time and recorded GPS coordinates, Street View URL, capture timestamp, and the camera's yaw, pitch, and field-of-view.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

ASSETS '18, October 22–24, 2018, Galway, Ireland

© 2018 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-5650-3/18/10.
https://doi.org/10.1145/3234695.3240999

ASSETS'19

Deep Learning for Automatically Detecting Sidewalk Accessibility Problems Using Streetscape Imagery

Galen Weld[1], Esther Jang[1], Anthony Li[2], Aileen Zeng[1], Kurtis Heimerl[1], and Jon E. Froehlich[1]
[1]Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, USA
[2]Department of Computer Science, University of Maryland, College Park, USA
{gweld, infrared, aileenz, kheimerl, jonf}@cs.washington.edu, amli@umd.edu

ABSTRACT

Recent work has applied machine learning methods to automatically find and/or assess pedestrian infrastructure in online map imagery (*e.g.,* satellite photos, streetscape panoramas). While promising, these methods have been limited by two interrelated issues: small training sets and the choice of machine learning model. In this paper, aided by the recently released Project Sidewalk dataset of 300,000+ image-based sidewalk accessibility labels, we present the first examination of deep learning to automatically assess sidewalks in Google Street View (GSV) panoramas. Specifically, we investigate two application areas: automatically *validating* crowdsourced labels and automatically *labeling* sidewalk accessibility issues. For both tasks, we introduce and use a residual neural network (ResNet) modified to support both image and non-image (contextual) features (*e.g.,* geography). We present an analysis of performance, the effect of our non-image features and training set size, and cross-city generalizability. Our results significantly improve on prior automated methods and, in some cases, meet or exceed human labeling performance.

Author Keywords

Neural networks, accessibility, sidewalks, computer vision

ACM Classification Keywords

I.2.10. Artificial Intelligence: Vision and Scene Understanding; I.2.6. Artificial Intelligence: Learning

INTRODUCTION

Sidewalks should benefit all of us. They provide a safe, environmentally-friendly conduit for moving about a city. For people with disabilities, sidewalks can have a significant impact on independence [47], quality of life [38], and overall physical activity [17]. While mapping tools like Google and Apple Maps have begun offering pedestrian-focused features, they do not incorporate sidewalk routes or information on sidewalk accessibility [23], which limits their utility and disproportionately affects people with disabilities. A key challenge is data: Where does it come from? How is it collected?

Traditionally, sidewalk audits—which gather data on the presence and quality of sidewalks—are performed via in-person inspections by city transit departments or community volunteers. However, these audits are expensive, labor intensive, and infrequent. [1] Moreover, the resulting data is in disparate formats, is not typically open (*i.e.,* published online), and is not intended for end-user tools [23, 50]. To expand who can collect sidewalk data and to improve data granularity and freshness, researchers have introduced smartphone-based tools [15, 46, 52] as well as instrumented wheelchairs [35, 39, 51, 57], both of which capture sidewalk information *in situ* as it's experienced. However, these tools have been limited by low adoption, small geographic coverage, and high user burden (*e.g.,* requiring users to take out their phones, load an app, take a picture, annotate it, and upload it) [20, 23].

To partially address these scalability issues, researchers have begun developing automated methods for sidewalk assessment using machine learning and online imagery (*e.g.,* satellite photos [10, 8], panoramic streetscape imagery [31, 32, 59]). While still early, these complementary approaches promise to dramatically decrease manual labor and cost. However, they have been limited by two interrelated issues: small training sets and the choice in machine learning model—both of which negatively impact performance. In this paper, we attempt to address both of these issues.

We present the first examination of deep learning methods to automatically assess sidewalk accessibility in terms of *curb ramps*, *missing curb ramps*, *surface problems*, and *sidewalk obstructions* from widely available streetscape imagery. Our work is enabled by the recently released Project Sidewalk open dataset, which contains a corpus of 300,000+ image-based sidewalk accessibility labels collected via remote crowdsourcing in Google Street View (GSV) [55] (Figure 1). Specifically, we investigate two application tasks using GSV panoramas: automatically *validating* crowdsourced labels and automatically *labeling* sidewalk accessibility issues.

Our research questions include:

- **R1:** How well does our machine learning approach perform across our two tasks (validation and labeling)?
- **R2:** What is the impact of additional, non-image related training features on performance?
- **R3:** How does classification accuracy change as a function of training data amount?
- **R4:** How well does our model generalize across cities?

To address these questions, we trained two sets of deep convolutional neural networks using ResNet-18 [33]—one set for

[1]As one example, the Seattle Department of Transportation completed their first ever sidewalk assessment in 2016, which took 14 interns nearly a year to complete. [1]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASSETS '19, October 28–30, 2019, Pittsburgh, PA, USA

© 2019 ACM. ISBN 978-1-4503-6676-2/19/10…$15.00

DOI: 10.1145/3308561.3353798

ASSETS'21

Experimental Crowd+AI Approaches to Track Accessibility Features in Sidewalk Intersections Over Time

ATHER SHARIF, Allen School of Computer Science, University of Washington, asharif@cs.uw.edu

PAARI GOPAL, Allen School of Computer Science, University of Washington, paarig@cs.uw.edu

MICHAEL SAUGSTAD, Allen School of Computer Science, University of Washington, saugstad@cs.uw.edu

SHIVEN BHATT, Allen School of Computer Science, University of Washington, shivbat@cs.uw.edu

RAYMOND FOK, Allen School of Computer Science, University of Washington, rayfok@cs.uw.edu

GALEN WELD, Allen School of Computer Science, University of Washington, gweld@cs.uw.edu

KAVI DEY, Seattle Academy High School

JON E. FROEHLICH, Allen School of Computer Science, University of Washington, jonf@cs.uw.edu

How do sidewalks change over time? Are there geographic or socioeconomic patterns to this change? These questions are important but difficult to address with current GIS tools and techniques. In this demo paper, we introduce three preliminary crowd+AI (Artificial Intelligence) prototypes to track changes in street intersection accessibility over time—specifically, curb ramps—and report on results from a pilot usability study.

CCS CONCEPTS • Human-centered computing → Accessibility → Accessibility systems and tools

Additional Keywords and Phrases: Mobility, disability, sidewalks, crowdsourcing, machine learning, change tracking

ACM Reference Format:

Ather Sharif, Paari Gopal, Michael Saugstad, Shiven Bhatt, Raymond Fok, Galen Weld, Kavi Dey, Jon E. Froehlich. 2021. Experimental Crowd+AI Approaches to Track Accessibility Features in Sidewalk Intersections Over Time. In The 23rd International ACM SIGACCESS Conference on Computers and Accessibility, October 18–22, 2021, Virtual Event. ACM, New York, NY, USA. ACM ISBN 978-1-4503-8306-6/21/10. https://doi.org/10.1145/3441852.3476549

1 Introduction

In 1990, the US passed the *Americans with Disabilities Act* (ADA) [1] requiring that public infrastructure—including sidewalks and street crossings—be accessible. Yet, more than 30 years later, cities struggle to meet accessibility requirements, often only pursuing large-scale sidewalk renovations in response to civil rights litigation such as in New York [12], Seattle [9], and Los Angeles [18]. Observing these challenges and to help stimulate and structure ADA renovations and city planning, in 2015, the *US Federal Highway Administration* requested that local governments develop sidewalk ADA transition plans, including an inventory of accessibility barriers and a description of accessible renovations [25]. In a recent study of 401 municipalities, however, only 54 (13%) had published plans and only seven met the minimum ADA criteria [7].

Such findings reflect the challenge in making infrastructure accessible. Viable solutions require substantial political, economic, and technical investment—training, resources, community involvement, specialized tools, and the work and coordination of multiple governmental agencies [19]. And there is a lack of open tools, techniques, and datasets to

RESEARCH

AUTO-LABELING



JOHN O'MEARA

22

k Accessibility Assessments: Initial Experiments
nd Cross-city Training on Performance

omputer Science, University of Washington, USA
Computer Science, University of Washington, USA
of Computer Science, University of Washington, USA
f Computer Science, University of Washington, USA
ol of Computer Science, University of Washington, USA
Imaging and Data Analysis (VIDA) Center, New York University, USA
of Computer Science, University of Washington, USA

echniques are being used to semi-automatically analyze the accessibility of built envi-
bout how to effectively combine the two. We present two experiments examining the
classifying sidewalk accessibility features in streetscape images. In Experiment 1, we
has been voted correct by the crowd but is more expensive to collect—compared with a
ment 2, we examine whether crowdsourced labeled data gathered in one city can be used
r, these experiments contribute to the growing literature in Crowd+AI approaches for
identify pertinent challenges.

din, Johnson Kuang, Michael Saugstad, Maryam Hosseini, and Jon E. Froehlich. 2022.
ssessments: Initial Experiments Examining Label Quality and Cross-city Training on
pages. https://doi.org/XXXXXXX.XXXXXXX

ncreasing emphasis on the design of inclusive cities [13], urban streets and
The problem is not just a lack of accessibility but a lack of data on sidewalk
mentally impacts policy making, urban planning, and the design of accessibility
aps [5, 15]. Traditionally, sidewalk data is gathered via in-person inspections
wever, this approach is laborious, expensive, and infrequent [2, 5, 12]. Thus,
methods using remote crowdsourcing [7, 16], computer vision [1, 9, 10, 19], or

scalable, they require copious amounts of training data. Large-scale, standard-
3] or *Mapillary Vistas* [14] exist and have dramatically accelerated research in
labeled datasets do not include sidewalk accessibility features. Furthermore,
rface condition assessment [10] have benefited from including streetscape
ive, researcher-labeled datasets for training, which limits experimental scope
ds or cities). To achieve larger training sets, others have explored crowdsourced

1

ASSETS'23

BusStopCV: A Real-time AI Assistant for Labeling Bus Stop Accessibility Features in Streetscape Imagery

Minchu Kulkarni[1,2] Chu Li[1] Jaye Ahn[1,2] Katrina Ma[3]
Zhihan Zhang[1] Michael Saugstad[1] Yochai Eisenberg[4] Valerie Novack[5]
Brent Chamberlain[5] Jon E. Froehlich[1]
[1]Allen School of Computer Science, University of Washington [2]MHCI+D Program [3]Human-Centered Design and Engineering, University of Washington [4]Disability and Human Development, University of Illinois Chicago
[5]Landscape Architecture & Environmental Planning, Utah State University



Figure 1: The BusStopCV human+AI data collection workflow: a real-time YOLOv8 CV model automatically detects bus stop features such as shelters and benches in streetscape imagery (left). Users can verify detections via lightweight click interactions (middle) or manually label features not detected by the model (right). Verified bounding boxes turn from dashed to solid lines. See demonstration video in supplementary materials.

ABSTRACT

Public transportation provides vital connectivity to people with disabilities, facilitating access to work, education, and health services. While modern navigation applications provide a suite of information about transit options—including real-time updates about bus or train arrivals—they lack data about the accessibility of the transit stops themselves. Bus stop features such as seatings, shelters, and landing areas are critical, but few cities provide this information. In this demo paper, we introduce BusStopCV, a Human+AI web prototype for scalably collecting data on bus stop features using real-time computer vision and human labeling. We describe BusStopCV's design, custom training with the YOLOv8 model, and an evaluation of 100 randomly selected bus stops in Seattle, WA. Our findings demonstrate the potential of BusStopCV and highlight opportunities for future work.

KEYWORDS

urban accessibility, bus stops, accessible transit systems, computer vision, crowdsourcing

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ASSETS '23, October 22–25, 2023, New York, NY, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0220-4/23/10.
https://doi.org/10.1145/3597638.3614481

ACM Reference Format:
Minchu Kulkarni, Chu Li, Jaye Ahn, Katrina Ma, Zhihan Zhang, Michael Saugstad, Yochai Eisenberg, Valerie Novack, Brent Chamberlain, Jon E. Froehlich. 2023. BusStopCV: A Real-time AI Assistant for Labeling Bus Stop Accessibility Features in Streetscape Imagery. In *The 25th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '23), October 22–25, 2023, New York, NY, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3597638.3614481

1 INTRODUCTION

Public transportation systems provide many benefits, such as mitigating traffic congestion [9], reducing pollution [9], enhancing economic growth [6], and improving health and well-being [16]; however, significant barriers continue to limit how people with disabilities use public transportation [26]. While a city's entire transit system requires careful assessment—from getting on a train or bus to finding accessible and safe seating—the transit stop itself is often overlooked, perceived as a mere waiting point instead of an integral part of mobility [1].

Modern navigation tools such as Google Maps [8] and Apple Maps [4] offer real-time bus information but fall short in providing data on essential bus stop accessibility features like seating availability, shelter provisioning, conditions of the landing area, and sidewalk connectivity [5]. Some cities publish open data about bus stop features; however, this practice is bespoke and limited by data collection costs and a lack of data standards [5]. In this demo paper, we introduce a new human+AI approach [3] to semi-automatically gather data on bus stops using real-time computer vision (CV) and

ASSETS'24

Towards Fine-Grained Sidewalk Accessibility Assessment with Deep Learning: Initial Benchmarks and an Open Dataset

Xinlei Liu*
Kevin Wu*
xllegion@cs.uw.edu
kwu2004@cs.uw.edu
Allen School of Computer Science
University of Washington, USA

Minchu Kulkarni
Allen School of Computer Science
University of Washington, USA
minchu@cs.uw.edu

Michael Saugstad
Allen School of Computer Science
University of Washington, USA
saugstad@cs.uw.edu

Peyton Rapo
Allen School of Computer Science
University of Washington, USA
peytor01@cs.uw.edu

Jeremy Freiburger
Allen School of Computer Science
University of Washington, USA
jeremyfr@uw.edu

Maryam Hosseini
City and Regional Planning
University of California, Berkeley
maryamh@berkeley.edu

Chu Li
Allen School of Computer Science
University of Washington, USA
chuchuli@cs.uw.edu

Jon E. Froehlich
Allen School of Computer Science
University of Washington, USA
jonf@cs.uw.edu



Figure 1: We examine whether deep learning models can classify sidewalk accessibility conditions from pre-cropped 640x640 streetscape images—*e.g.*, whether a curb ramp is *too steep, too narrow,* or *missing a tactile indicator* or if a sidewalk panel is *uneven, bumpy,* or composed of *brick/cobblestone.* The grid above showcases all 33 conditions we attempt to infer.

Abstract

We examine the feasibility of using deep learning to infer 33 classes of sidewalk accessibility conditions in pre-cropped streetscape images, including *bumpy, brick/cobblestone, cracks, height difference (uplifts), narrow, uneven/slanted, pole,* and *sign.* We present two experiments: first, a comparison between two state-of-the-art computer vision models, Meta's *DINOv2* and OpenAI's *CLIP-ViT*, on a cleaned dataset of ~24k images; second, an examination of a larger but noisier crowdsourced dataset (~87k images) on the best performing model from Experiment 1. Though preliminary, Experiment 1 shows that certain sidewalk conditions can be identified with high precision and recall, such as missing tactile warnings on curb ramps and grass grown on sidewalks, while Experiment 2 demonstrates that larger but noisier training data can have a detrimental effect on performance. We contribute an open dataset and classification benchmarks to advance this important area.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ASSETS '24, October 27–30, 2024, St. John's, NL, Canada
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0677-6/24/10
https://doi.org/10.1145/3663548.3688531

CV4A11y'25

RampNet: A Two-Stage Pipeline for Bootstrapping Curb Ramp Detection in Streetscape Images from Open Government Metadata

John S. O'Meara[1,2] Jared Hwang[2] Zeyu Wang[2]
Michael Saugstad[2] Jon E. Froehlich[2]
[1]Issaquah High School [2]University of Washington

https://github.com/ProjectSidewalk/RampNet

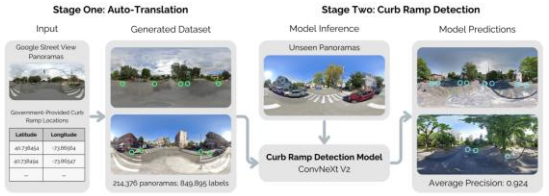


Figure 1. We introduce RampNet, a custom two-stage pipeline for bootstrapping curb ramp detection models in streetscape images from open government data. In Stage 1, we auto-generate a labeled streetview curb ramp dataset by translating government-provided curb ramp location data (*i.e.*, <lat, long> lists) into pixel labels on Google Street View panoramas. Stage 2 then uses this generated dataset to train a detection model that predicts curb ramp points in unseen panoramas.

Abstract

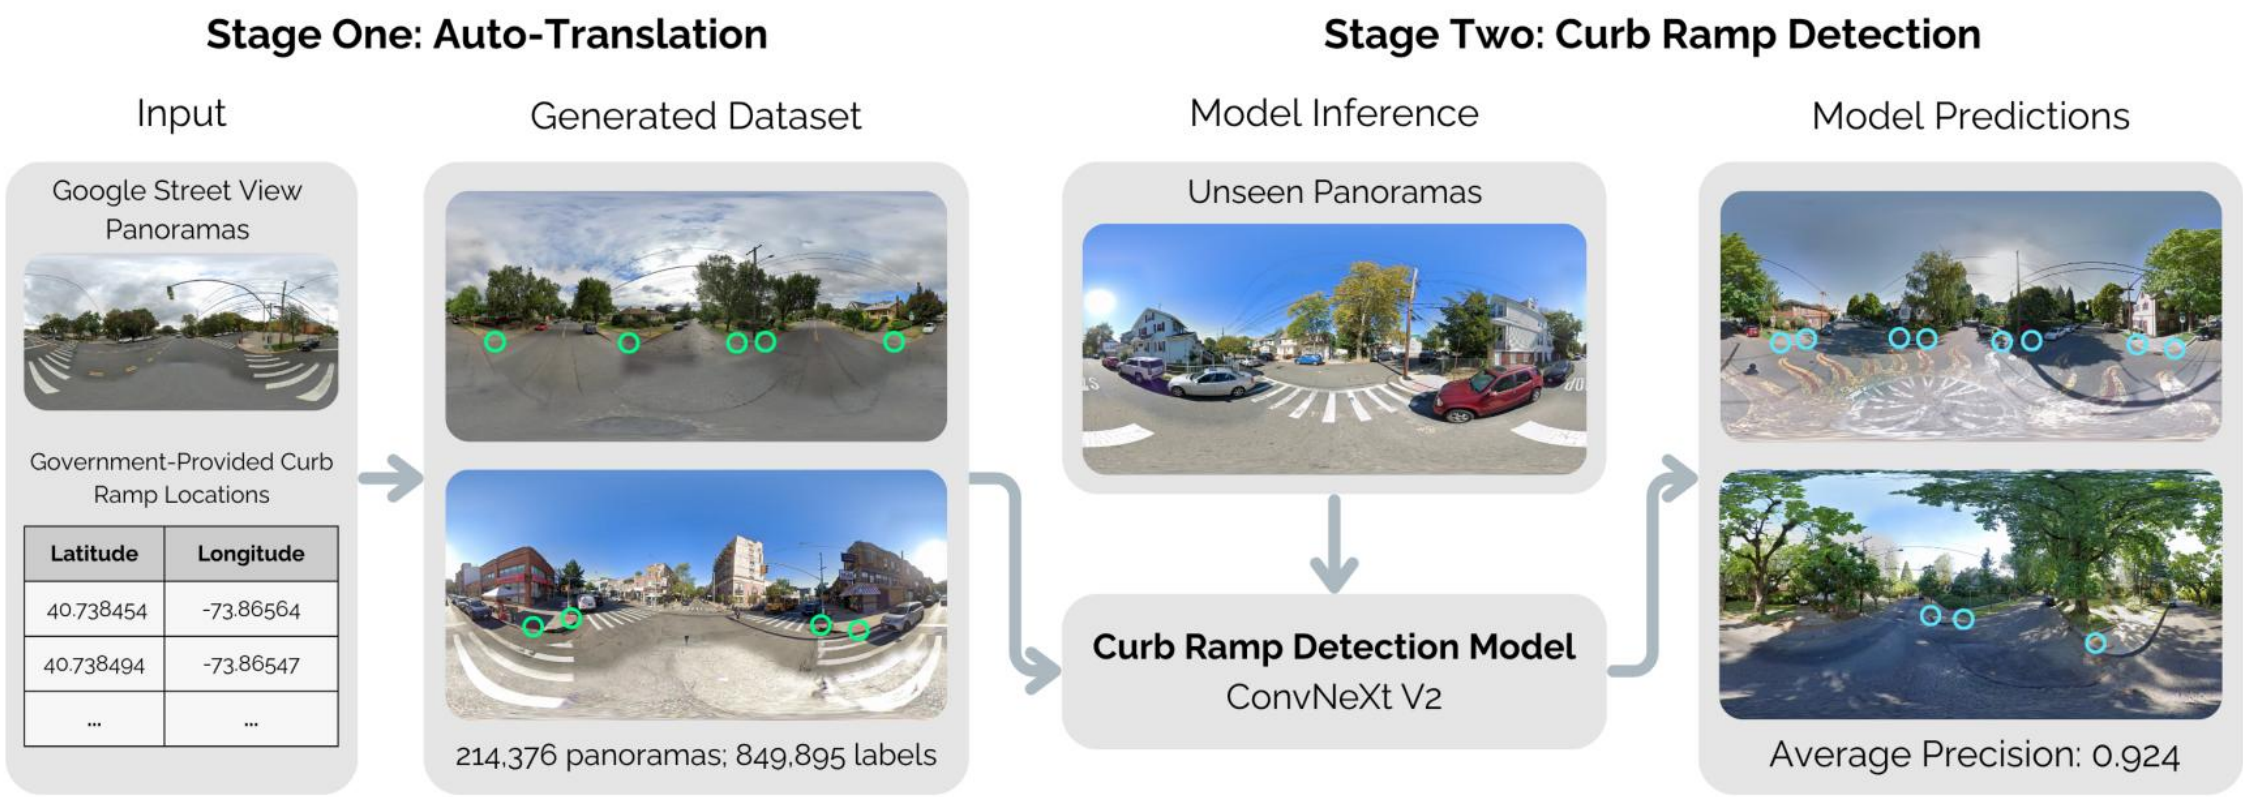
Curb ramps are critical for urban accessibility, but robustly detecting them in images remains an open problem due to the lack of large-scale, high-quality datasets. While prior work has attempted to improve data availability with crowdsourced or manually labeled data, these efforts often fall short in either quality or scale. In this paper, we introduce and evaluate a two-stage pipeline to scale curb ramp detection datasets and improve model performance. In Stage 1, we generate a dataset of more than 210,000 annotated Google Street View (GSV) panoramas by auto-translating government-provided curb ramp location data to pixel coordinates in panoramic images. In Stage 2, we train a curb ramp detection model (modified ConvNeXt V2) from the generated dataset, achieving state-of-the-art performance. To evaluate both stages of our pipeline, we compare to manually labeled panoramas. Our generated dataset achieves 94.0% precision and 92.5% recall, and our detection model reaches 0.9236 AP—far exceeding prior work. Our work contributes the first large-scale, high-quality curb ramp detection dataset, benchmark, and model.

1. Introduction

Curb ramps are critical to urban accessibility [24, 33], impacting both independence [26] and wellbeing [39] for people with mobility disabilities. However, information regarding curb ramp location and condition is limited, making it difficult to plan and maintain accessible infrastruc-

AUTO-LABELING

RAMPNET

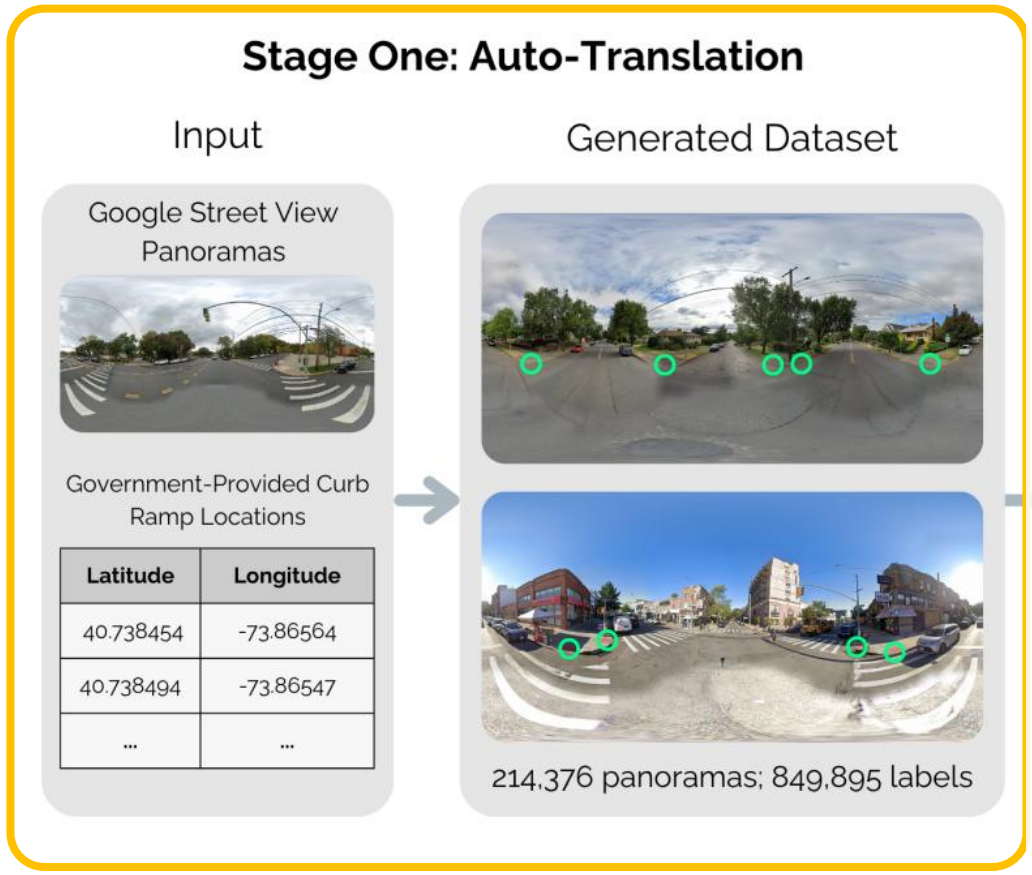


O'Meara, J., *et al.* RampNet: A Two-Stage Pipeline for Bootstrapping Curb Ramp Detection in Streetscape Images from Open Government Metadata. ICCV CV4A11y 2025



AUTO-LABELING

LARGE TRAINING DATASET



3 US cities: NYC, Portland, Bend

214k labeled panoramas

850k curb ramp labels

O'Meara, J., *et al.* RampNet: A Two-Stage Pipeline for Bootstrapping Curb Ramp Detection in Streetscape Images from Open Government Metadata. ICCV CV4A11y 2025

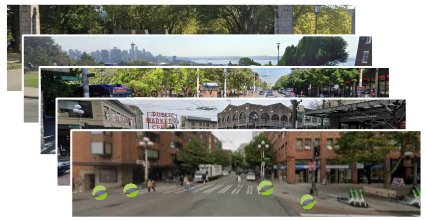
AUTO-LABELING

RAMPNET VS. PRIOR WORK

TEST DATASET

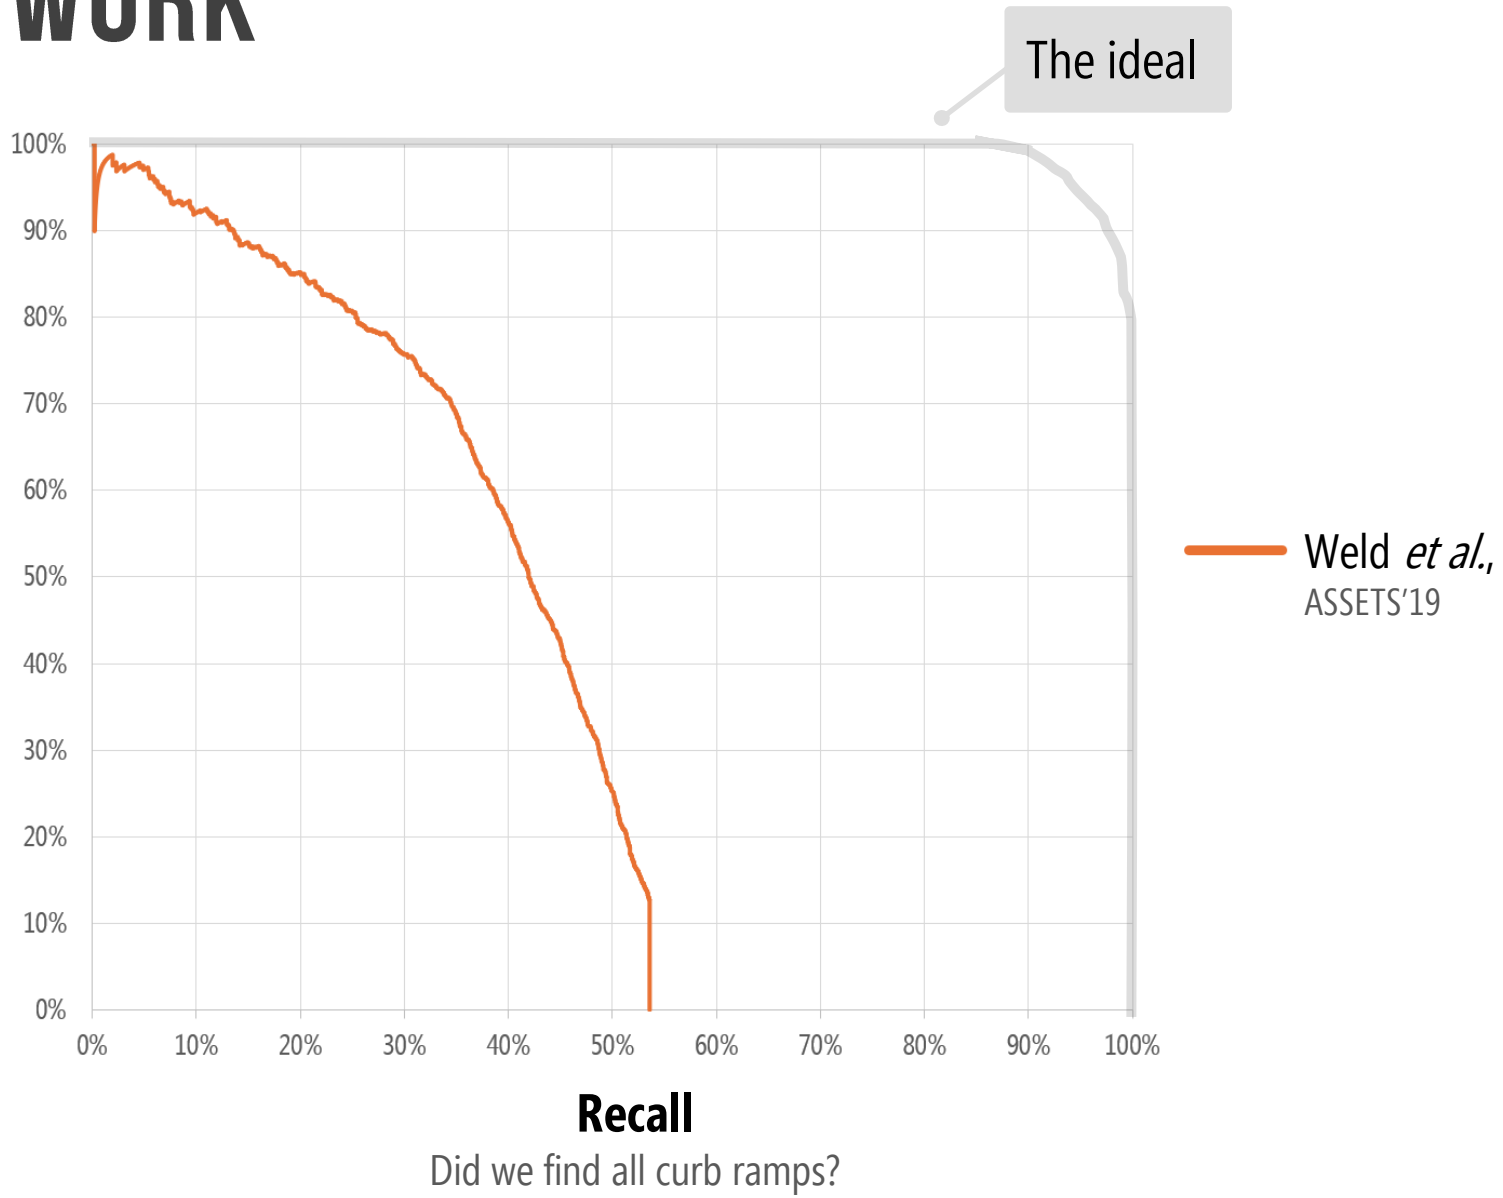
1,000 Panoramas

3,919 Curb ramp labels



Precision

Of everything found, were they all correctly curb ramps?



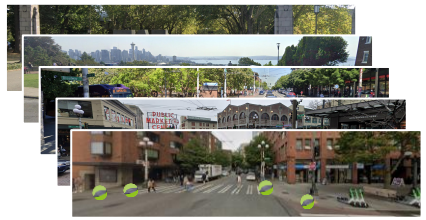
AUTO-LABELING

RAMPNET VS. PRIOR WORK

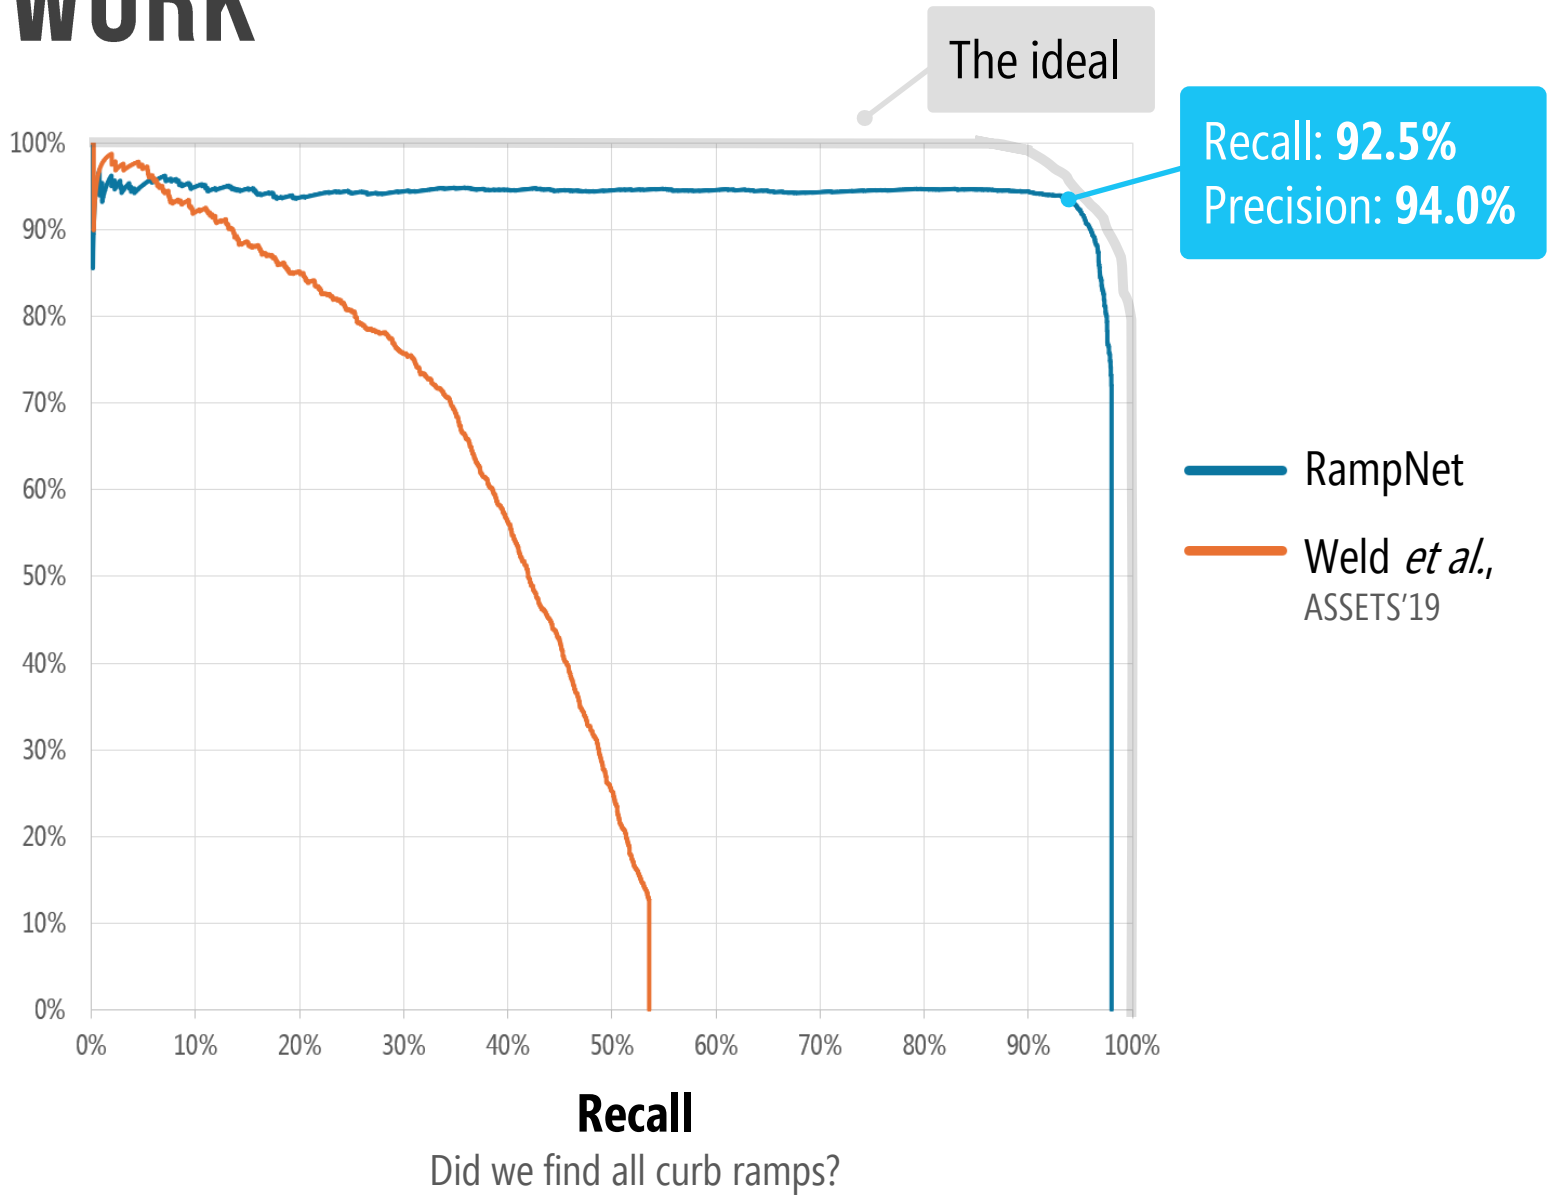
TEST DATASET

1,000 Panoramas

3,919 Curb ramp labels



Precision
Of everything found, were they all correctly curb ramps?



Recall
Did we find all curb ramps?

AUTO-LABELING

INTEGRATING INTO PROJECT SIDEWALK

Stage One: Auto-Translation

Input

Google Street View Panoramas



Government-Provided Curb Ramp Locations

Latitude	Longitude
40.738454	-73.86564
40.738494	-73.86547
...	...

Generated Dataset



214,376 panoramas; 849,895 labels

Stage Two: Curb Ramp Detection

Model Inference

Unseen Panoramas



Curb Ramp Detection Model
ConvNeXt V2

Model Predictions



Average Precision: 0.924

Stage Three: Integration
Project Sidewalk

Stage Four: Test Deployment
Vancouver, WA



O'Meara, J., *et al.* RampNet: A Two-Stage Pipeline for Bootstrapping Curb Ramp Detection in Streetscape Images from Open Government Metadata. ICCV CV4A11y 2025

VANCOUVER, WA

http://vancouver-wa.projectsidewalk.org/

SINGLE
AI Labeler

64.8K
Curb Ramp Labels














VANCOUVER, WA
http://vancouver-wa.projectsidewalk.org/

VANCOUVER, WA

http://vancouver-wa.projectsidewalk.org/



Three Experts



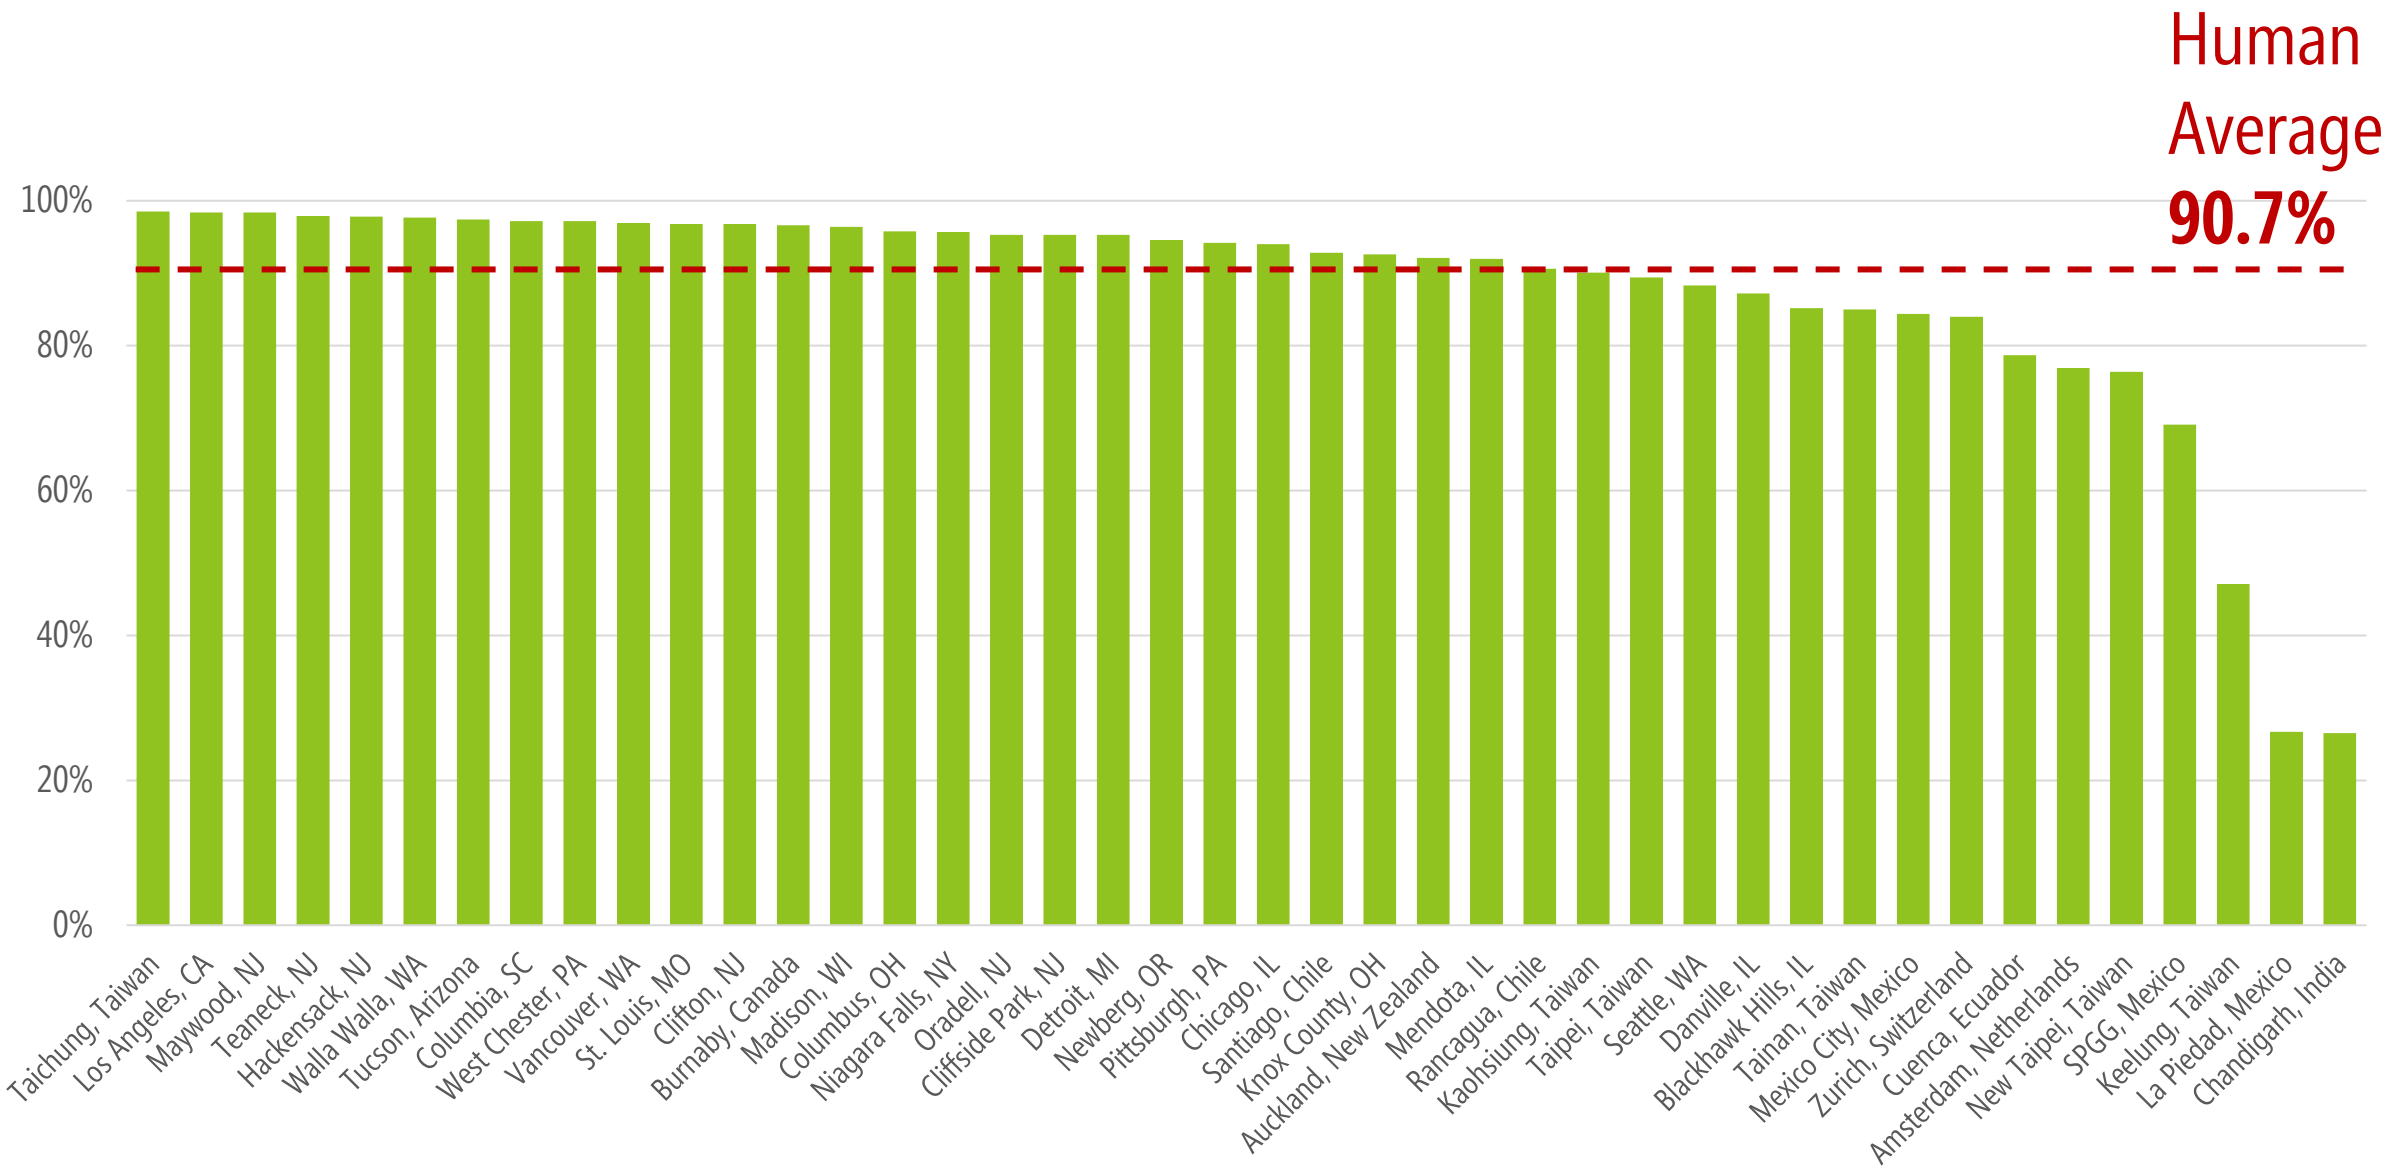
Reviewed 2550 RampNet Labels



Precision 96.9%

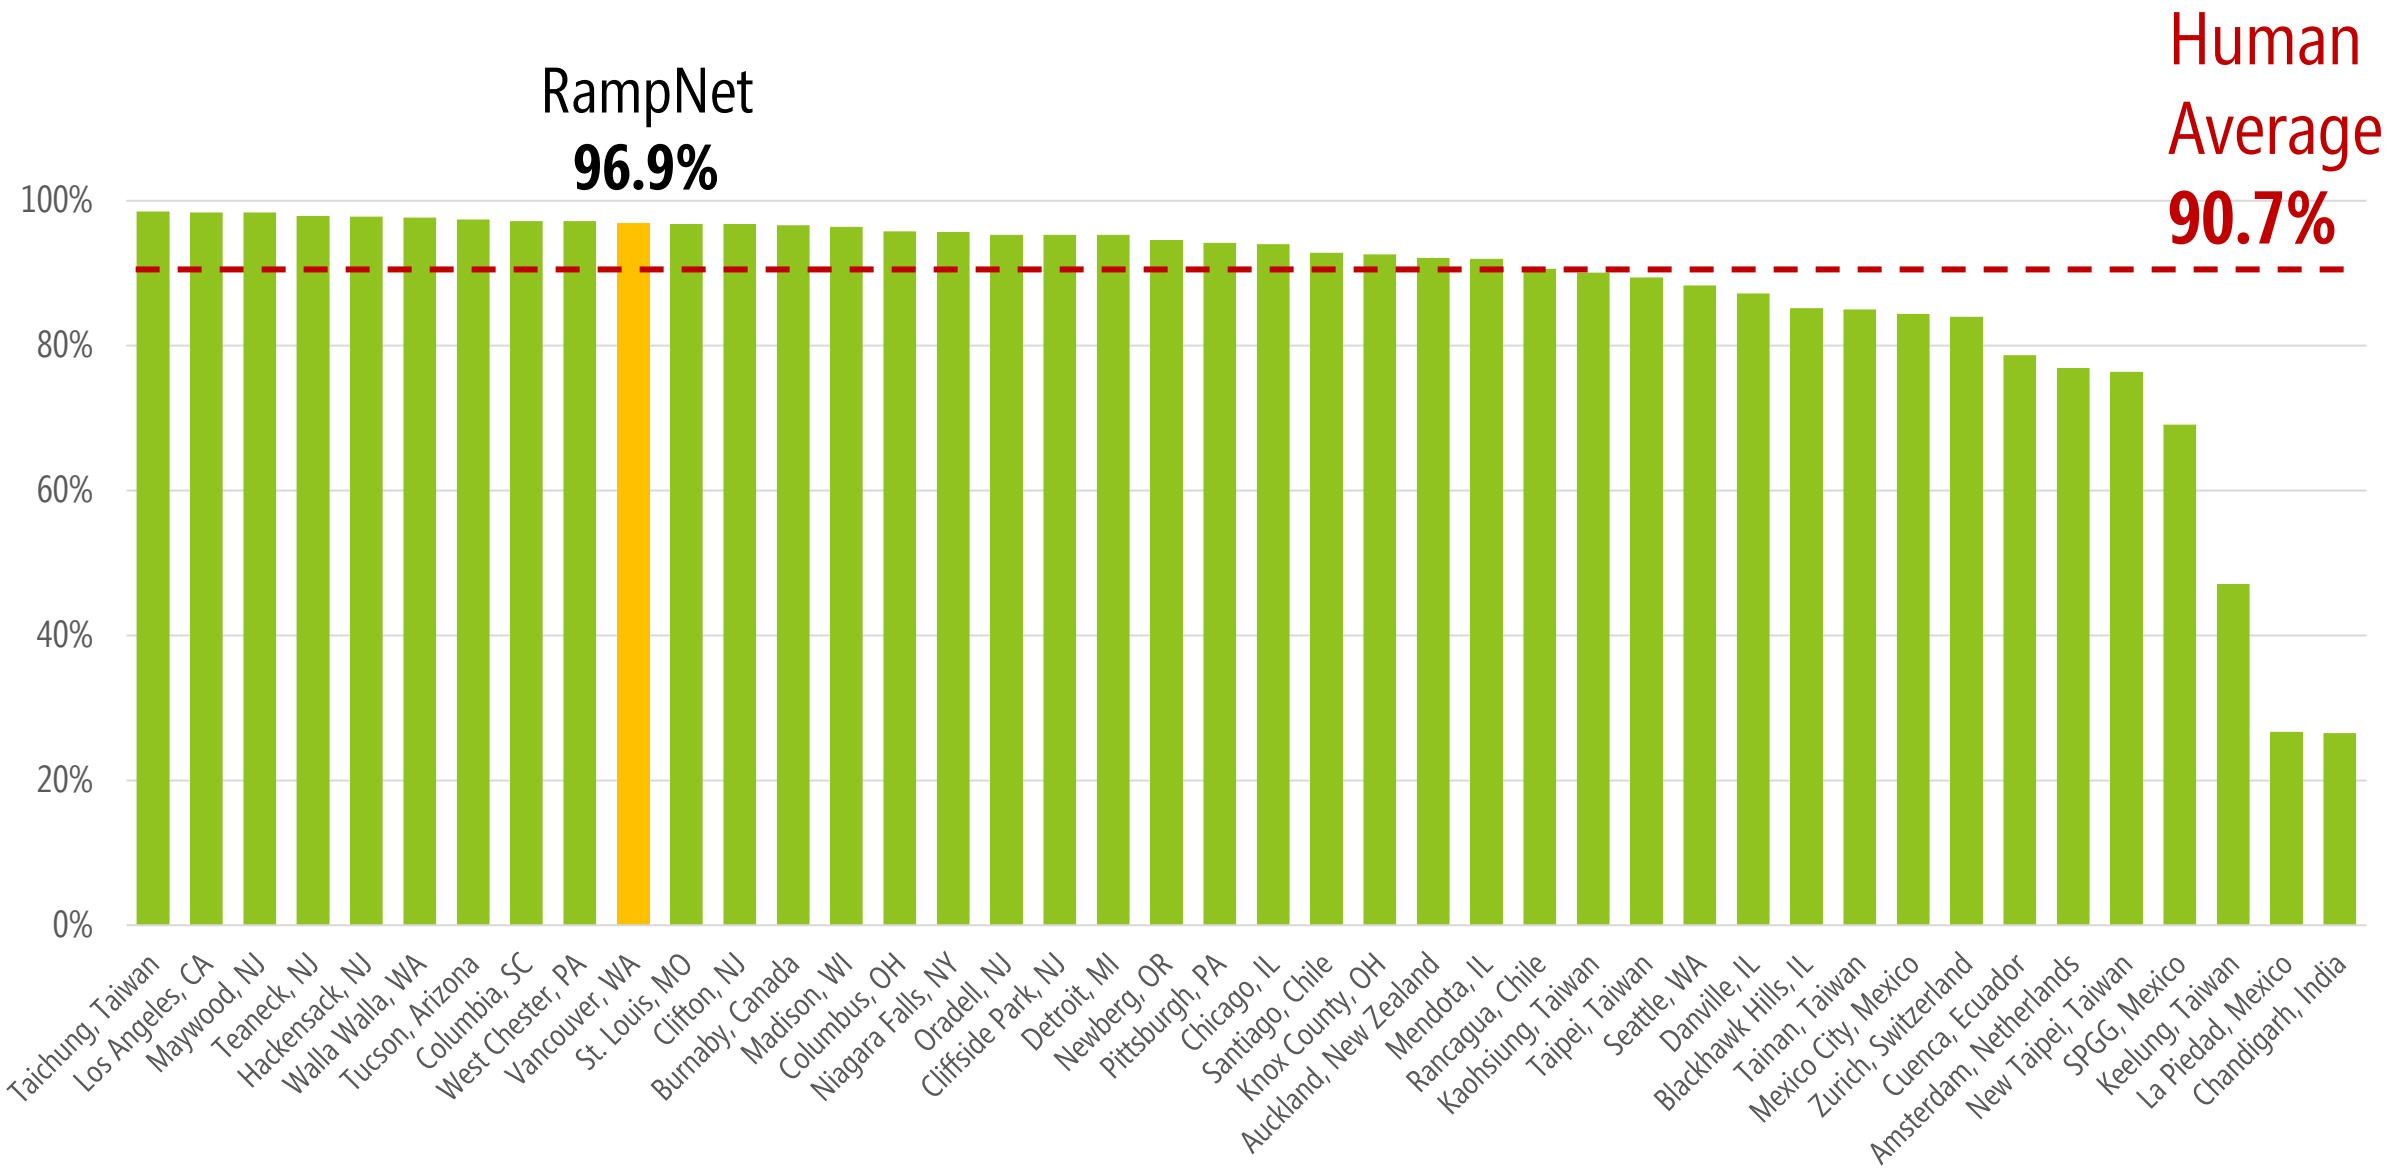
AUTO-LABELING

CONTEXTUALIZING RAMPNET PERFORMANCE



AUTO-LABELING

CONTEXTUALIZING RAMPNET PERFORMANCE



Bend, OR, US

Google Street View

6 curb ramps, 6 detected

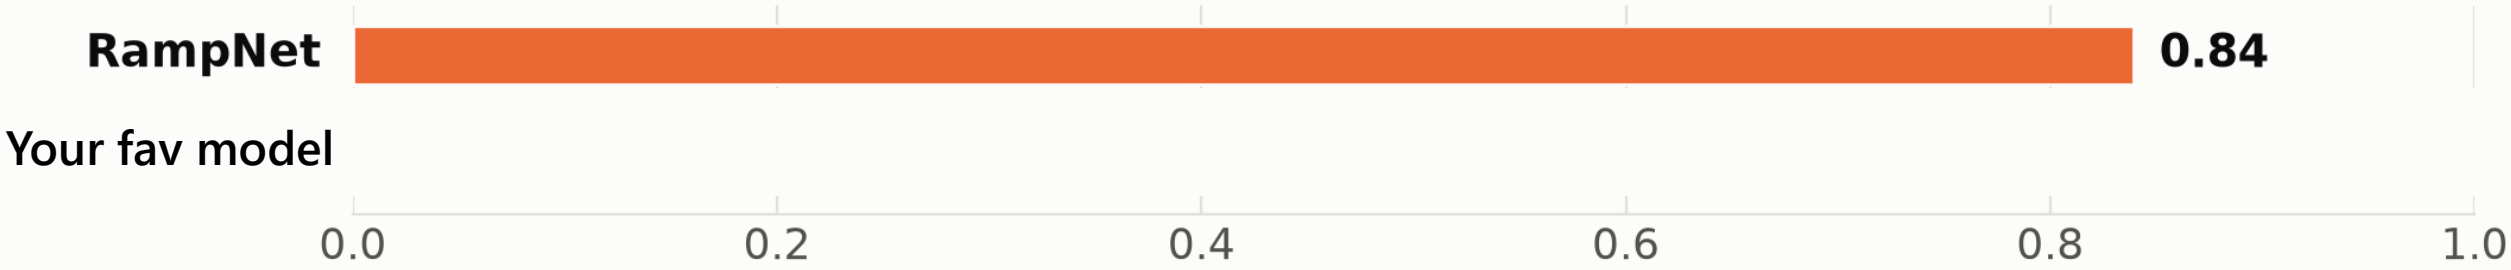
1 / 22



What about **frontier or open-weight models**?

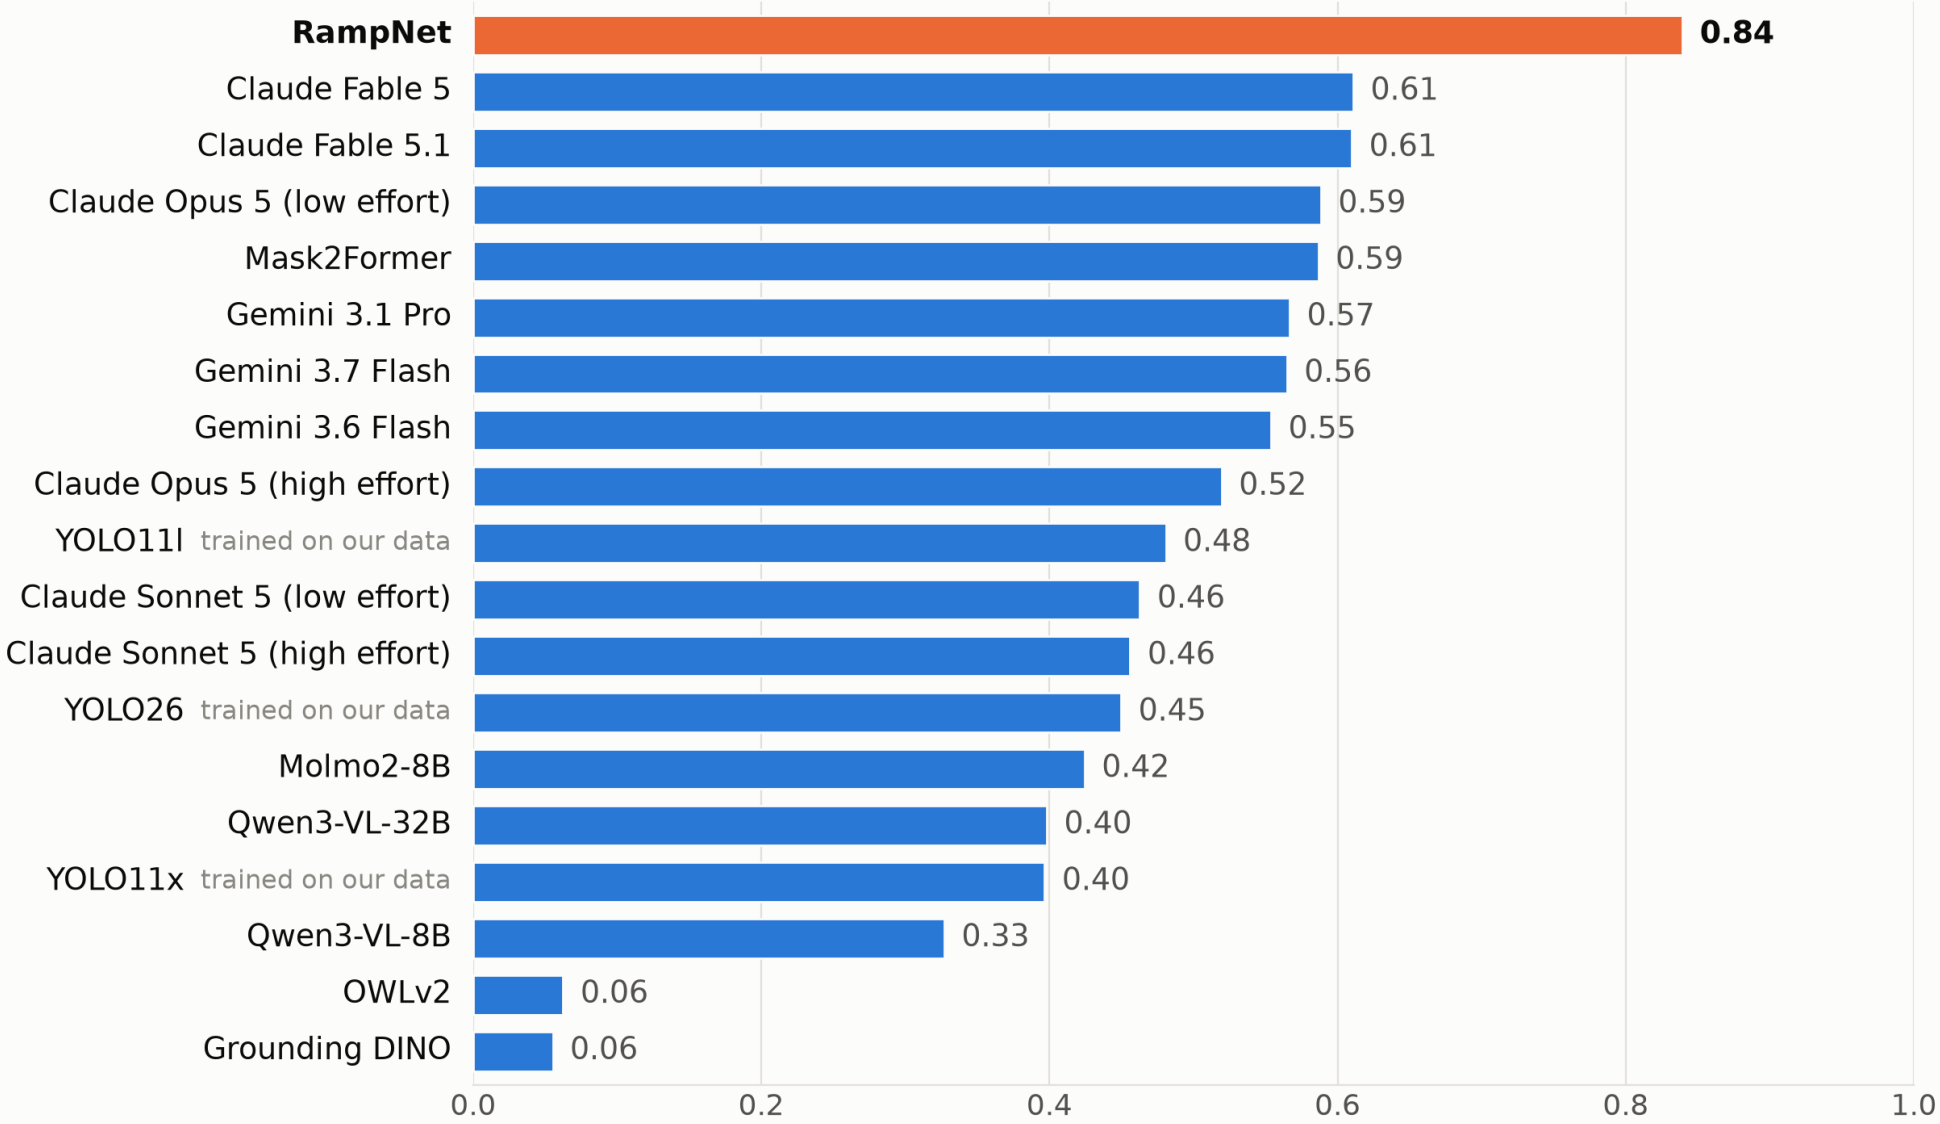
AUTO-LABELING

COMPARING RAMPNET TO FRONTIER MODELS: F1 SCORES



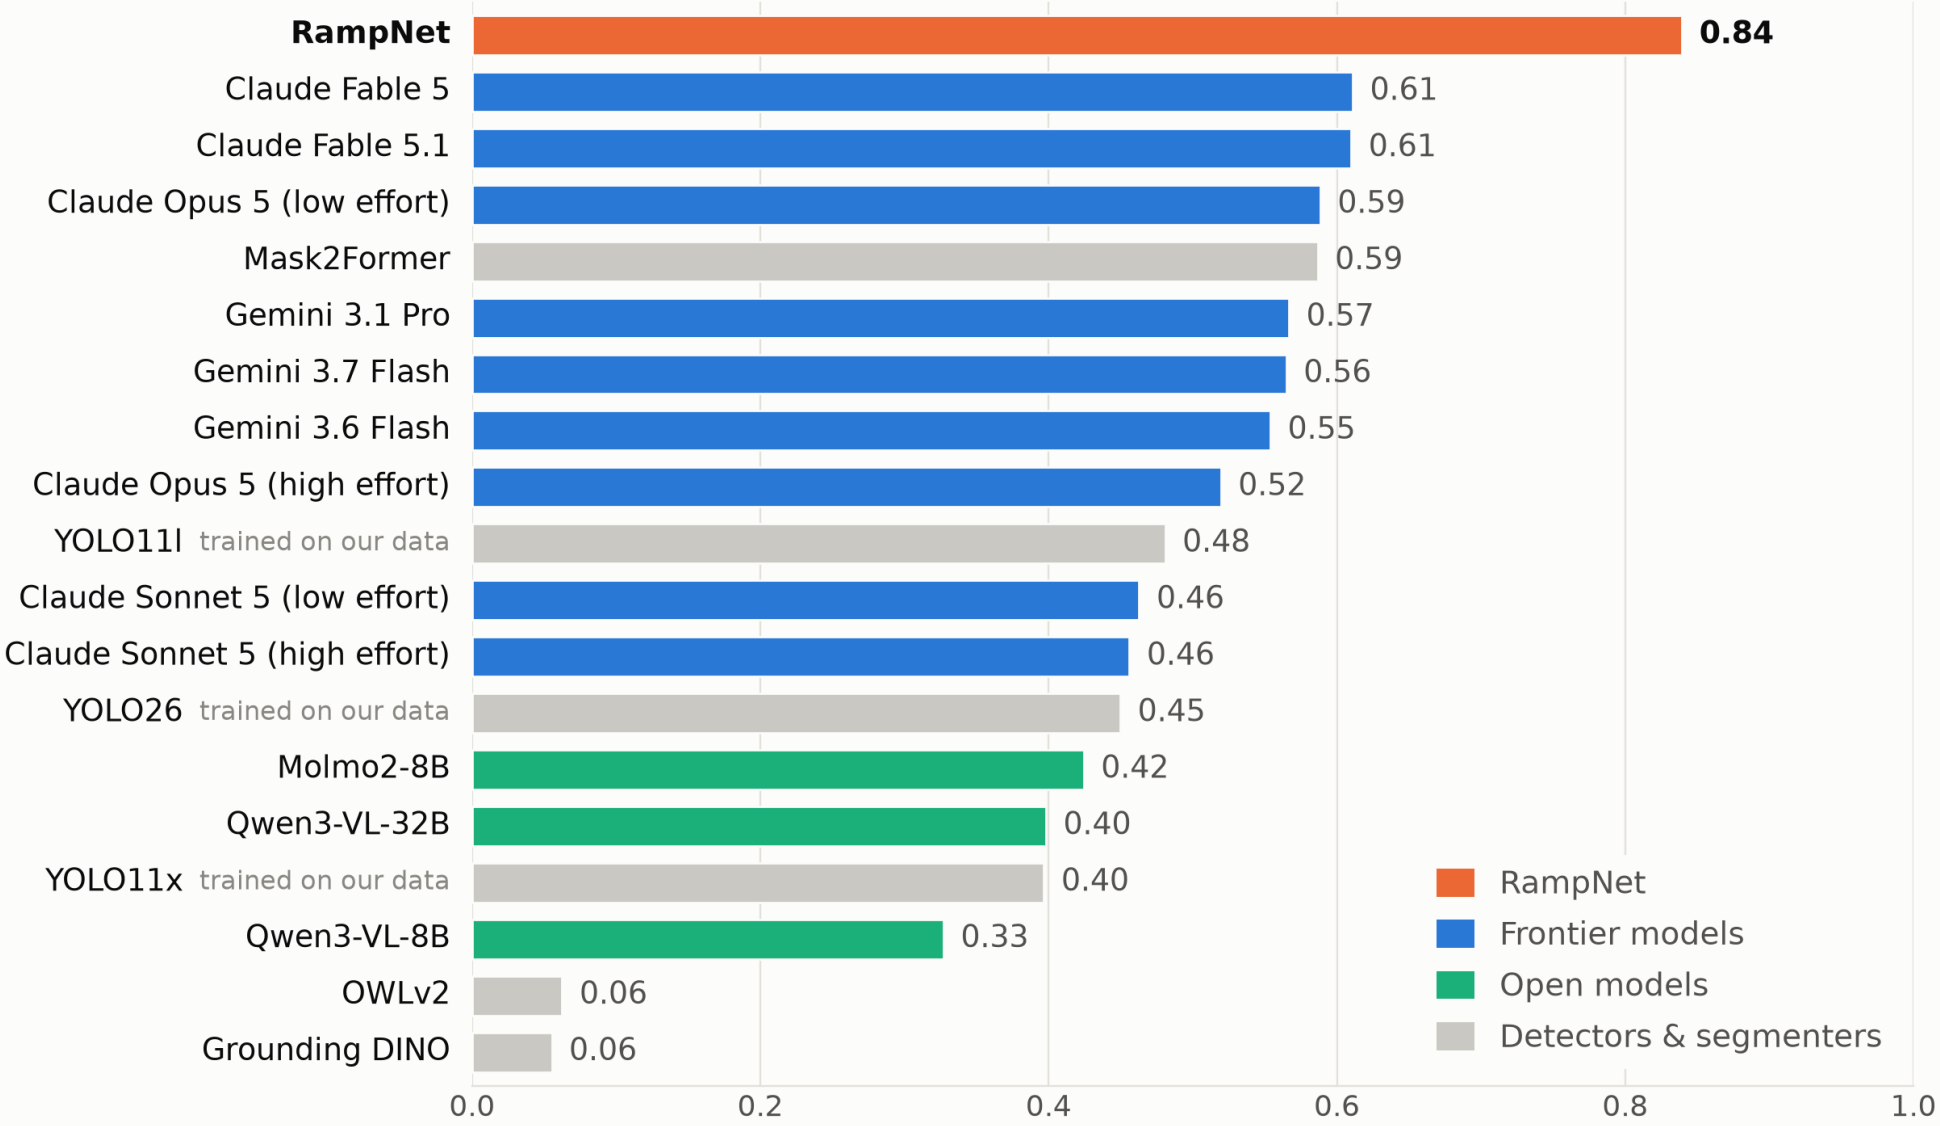
AUTO-LABELING

COMPARING RAMPNET TO FRONTIER MODELS: F1 SCORES



AUTO-LABELING

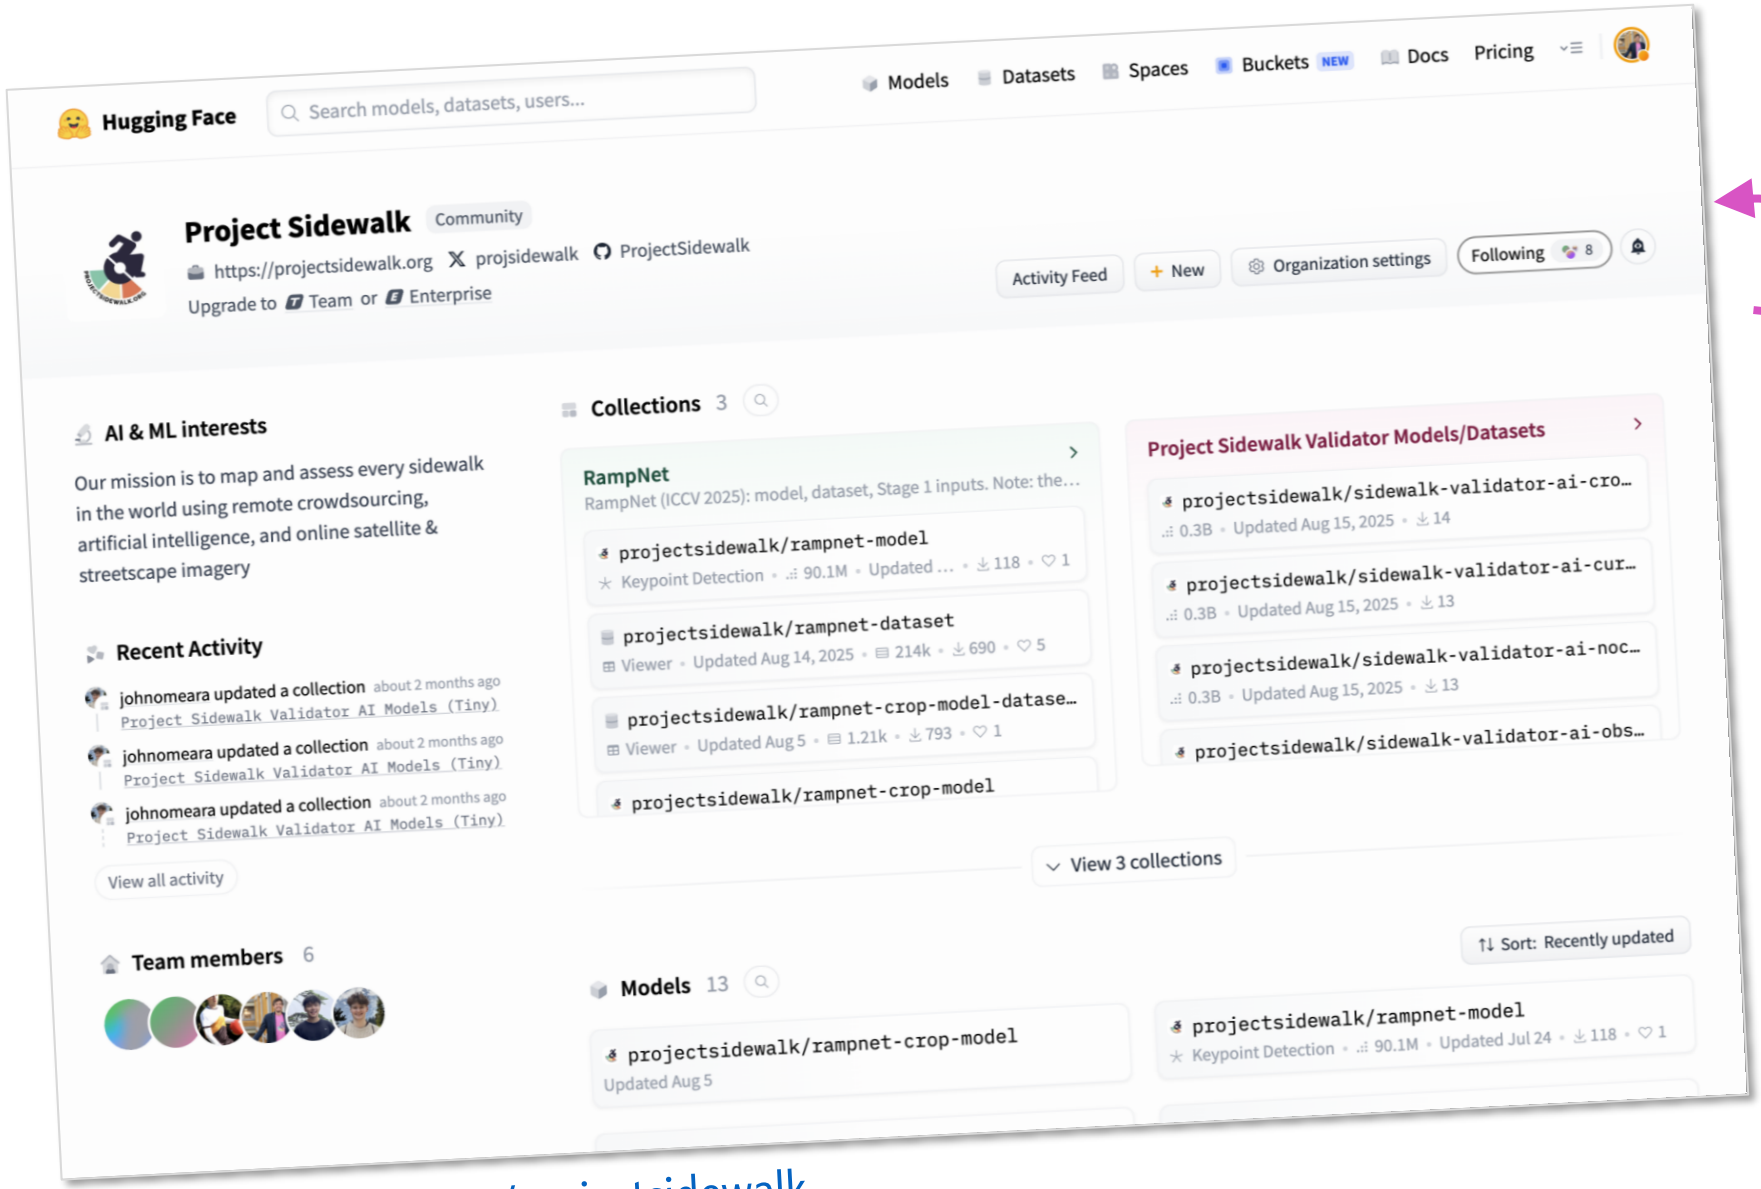
COMPARING RAMPNET TO FRONTIER MODELS: F1 SCORES




LABEL ENTIRE CITIES IN MINUTES
WEST MINNEHAHA
MINNEHAHA
CARTER PARK
FOURTH PLAIN VILLAGE
BAGLEY DOWNS
MEADOW HOMES
CENTRAL PARK
VANCOUVER
HARNEY HEIGHTS
NORTHCREST
SOUTH CLIFF
EDGEWOOD PARK
HUDSON'S BAY
DUBOIS PARK
EVERGREEN SHORES
EVERGREEN HIGHLANDS
Columbia River
North Portland Harbor

More work needed for additional **sidewalk features** and **problems**





https://huggingface.co/projectsidewalk

Try to **beat us**!

(32k downloads on Hugging Face)

2 LABEL SIDEWALKS

Collect geo-located sidewalk assessments

Labeling Mission Routes

Streetscape Imagery APIs

Google Street View

Meta Mapillary

Panoramax

iNovitas infra3D

Community Partners

Labeling Missions

AI Labeler

Custom Trained ConvNeXt V2

CURB RAMPS

MISSING RAMP

SURFACE PROB

OBSTACLE

NO SIDEWALK

CROSSWALK

PED SIGNAL

3 VALIDATE LABELS

Human-AI quality control

Human Validation UI

Is this a **Surface Problem**?

Agree | Disagree | Not sure

AI-Assisted Human Validation

Custom Trained DINOv2

AI Validation

Agree

Disagree

4 POSTPROCESSING

Prepare data for apps

Raw Labels

Remove poor-quality labelers

Geo-cluster proximal labels

Label Clusters

2 LABEL SIDEWALKS

Collect geo-located sidewalk assessments

Labeling Mission Routes

Streetscape Imagery APIs

Google Street View

Meta Mapillary

Panoramax

iNovitas infra3D

Community Partners

Labeling Missions

AI Labeler

Custom Trained ConvNeXt V2

CURB RAMPS

MISSING RAMP

SURFACE PROB

OBSTACLE

NO SIDEWALK

CROSSWALK

PED SIGNAL

3 VALIDATE LABELS

Human-AI quality control

Human Validation UI

AI-Assisted Human Validation

Custom Trained DINOv2

AI Validation

Agree

Disagree

4 POSTPROCESSING

Prepare data for apps

Raw Labels

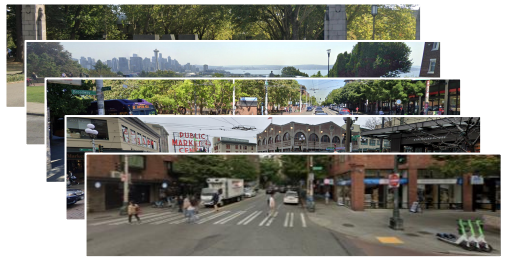
Remove poor-quality labelers

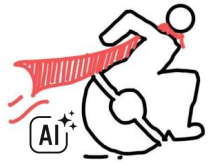
Geo-cluster proximal labels

Label Clusters

AI LABELER

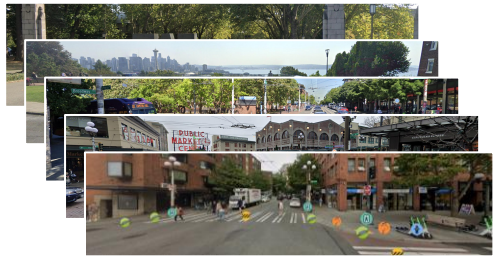
Model Input
Unlabeled panos





Custom Trained
AI Labeler

Model Output
Fully labeled panos



Easier **task**

AI VALIDATOR

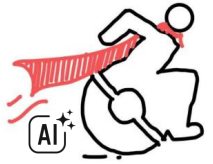
Model Input
Labeled crops











Custom Trained
AI Validator

Model Output
Label correctness?
Missing tag?
...

AUTO-VALIDATOR

AI VOTES ON LABEL CORRECTNESS

	Is this a **curb ramp?**	... a **missing curb ramp?**	... a **surface problem?**	... an **obstacle?**
Agree				
Disagree				

AUTO-VALIDATOR

SYSTEM DIAGRAM

Labeled crop
Image + label







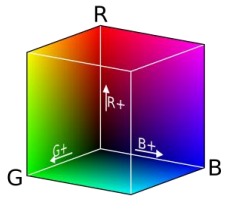


Resize Crop
256x256px

640x640

256x
256

Image Tensor
3x256x256



DINOv2-Large
ViT-L/14 Distilled
Backbone

Feature
Extractor

Classification Head
Binary classifier

Output
Vote

Agree

Disagree



Marked Crosswalk

Severity ●○○ AI 1 0

AI voted **yes** on this label



Marked Crosswalk

Severity ●○○ AI 1 0



Missing Curb Ramp

Severity ●○○ AI 1 0 Tags alternate route present



Curb Ramp

Severity ●○○ AI 1 0

AI VALIDATOR

AUTOMATIC VALIDATION PERFORMANCE

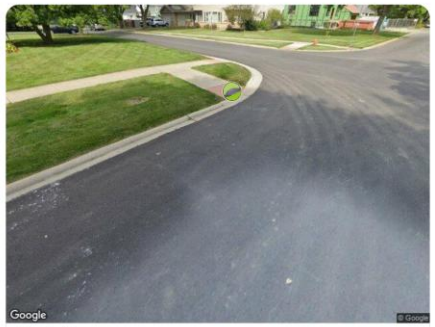


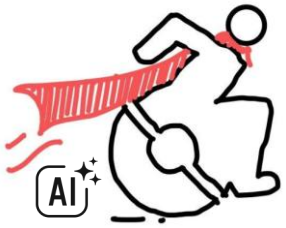












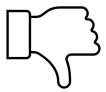
66,132 Labels

AI Votes

Human Majority Vote



96.3% Agreement



60.2% Agreement


WHAT'S THE IMPACT?
LABELS
100
1k
10k
100k
PROJECTSIDEWALK.ORG
mapbox

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

Service Learning

Research Engine

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



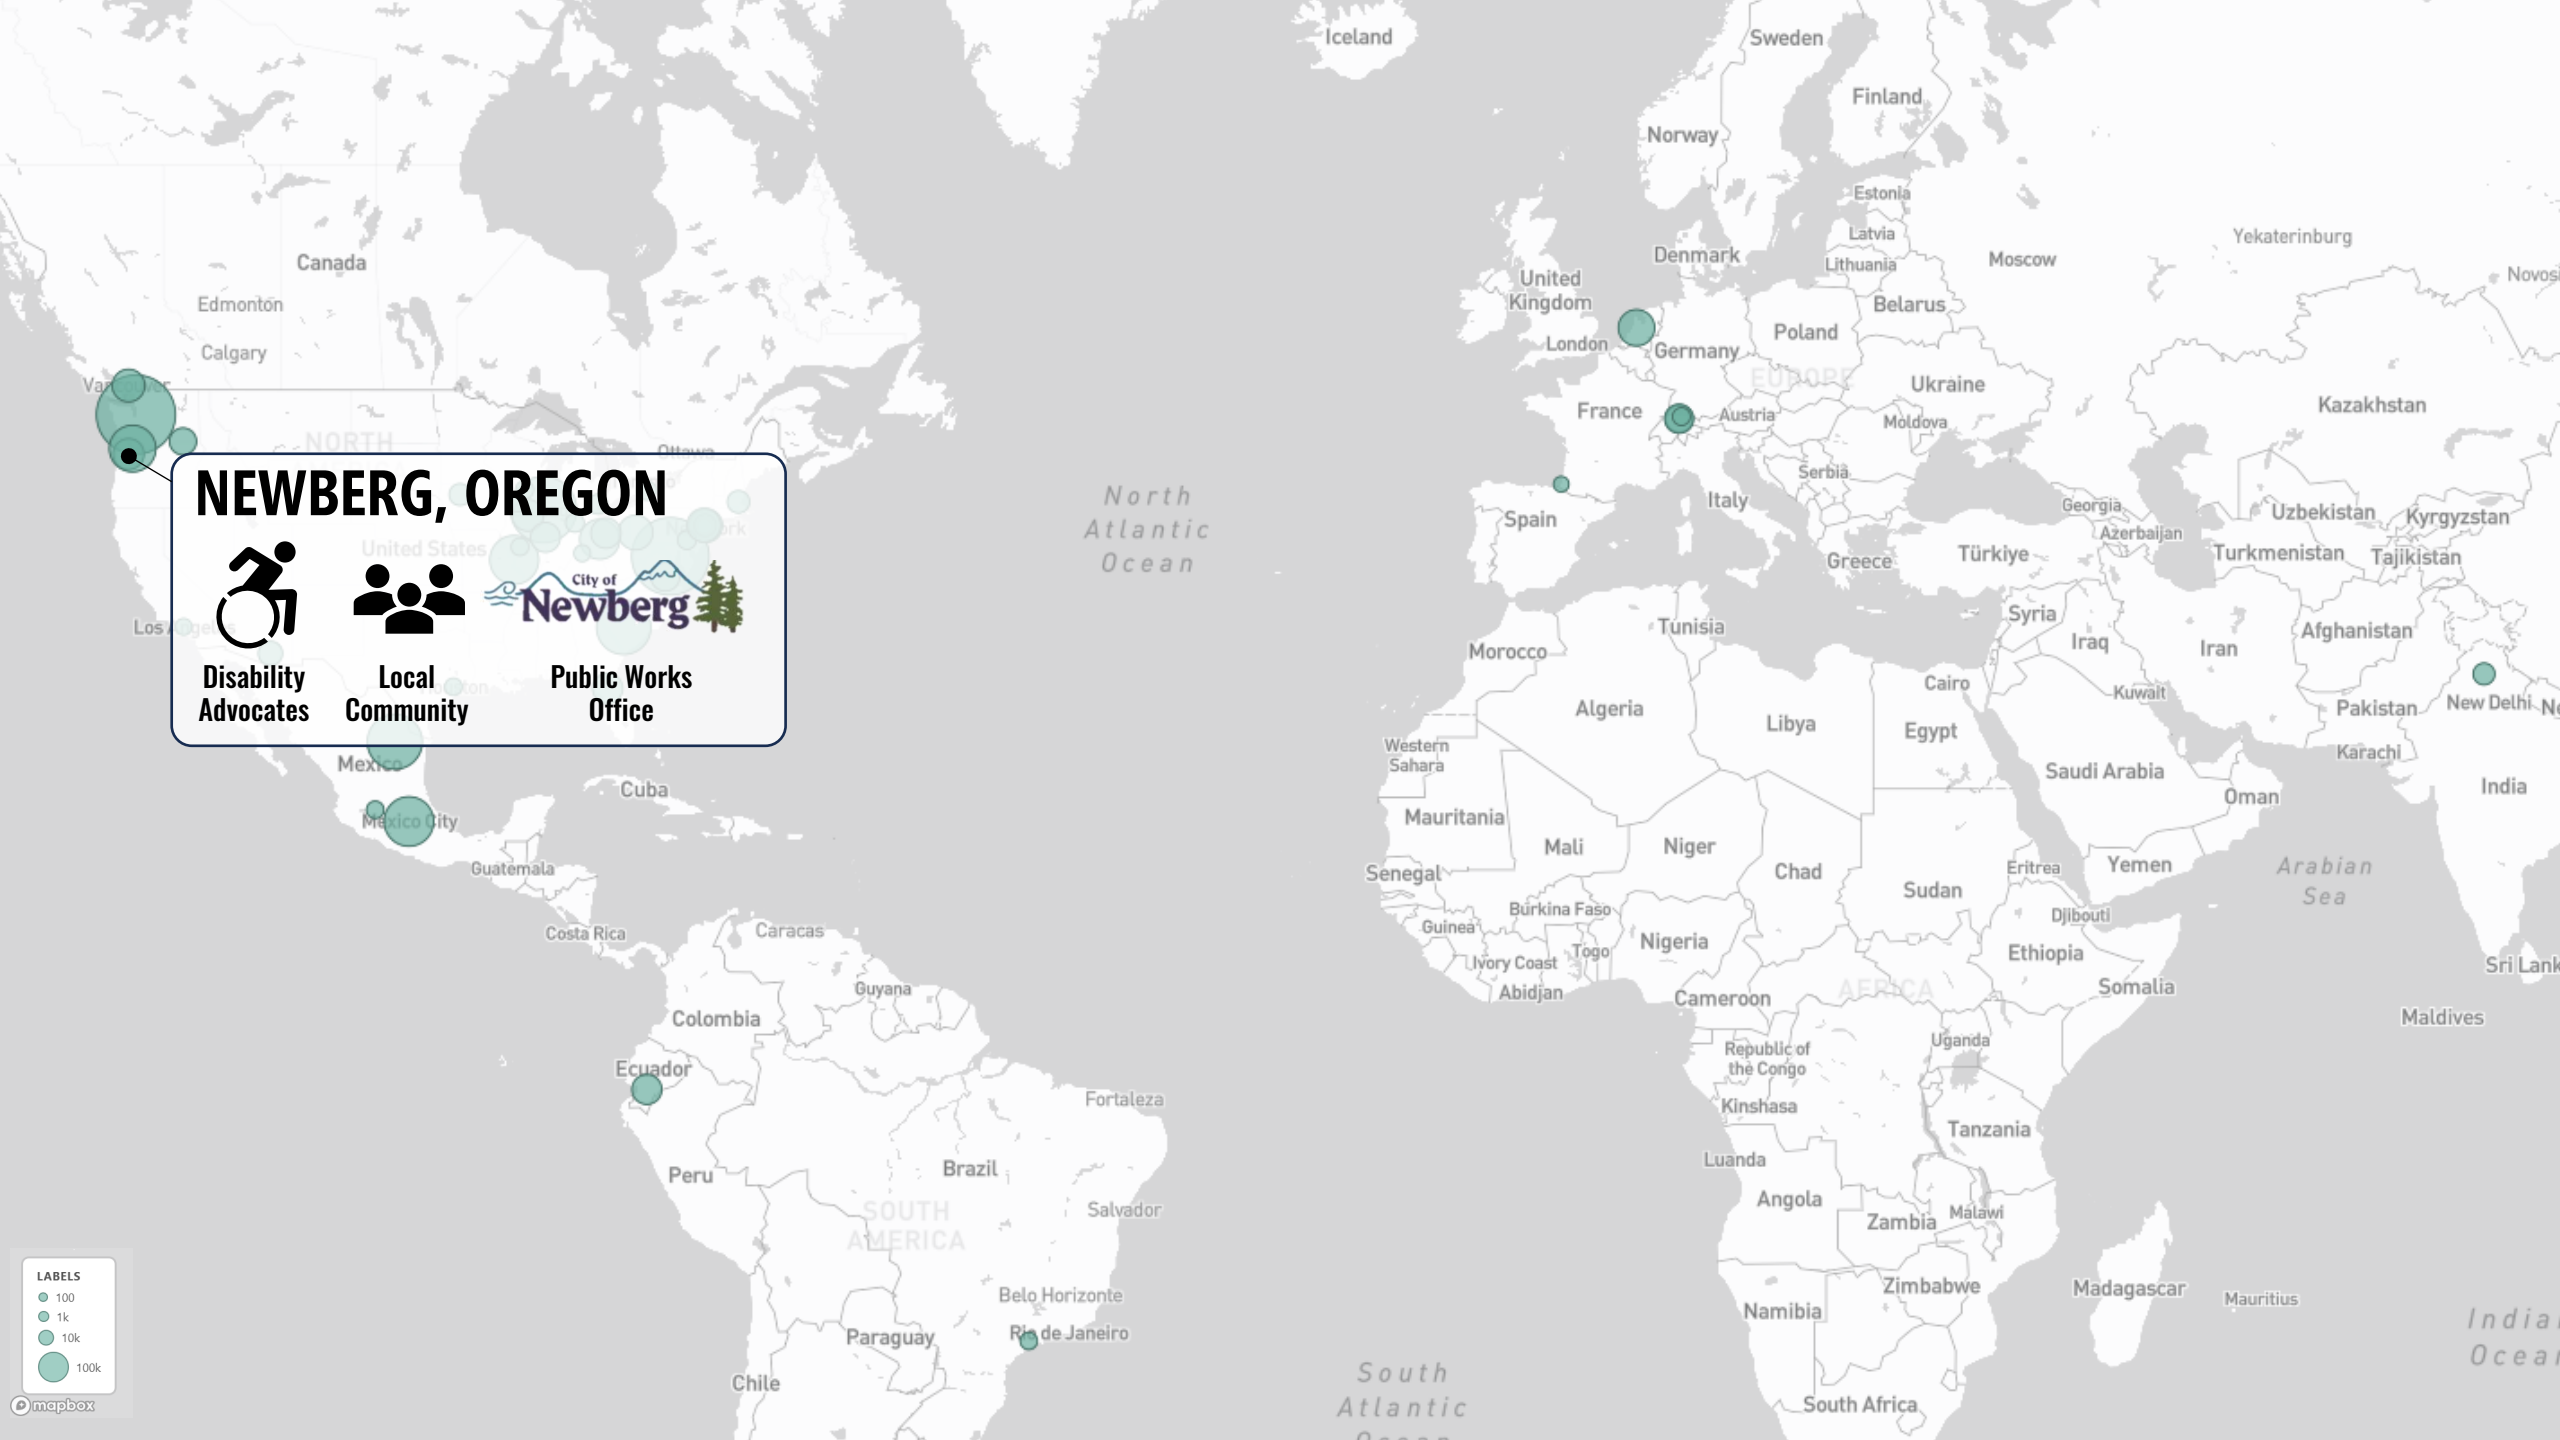
PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

Service Learning

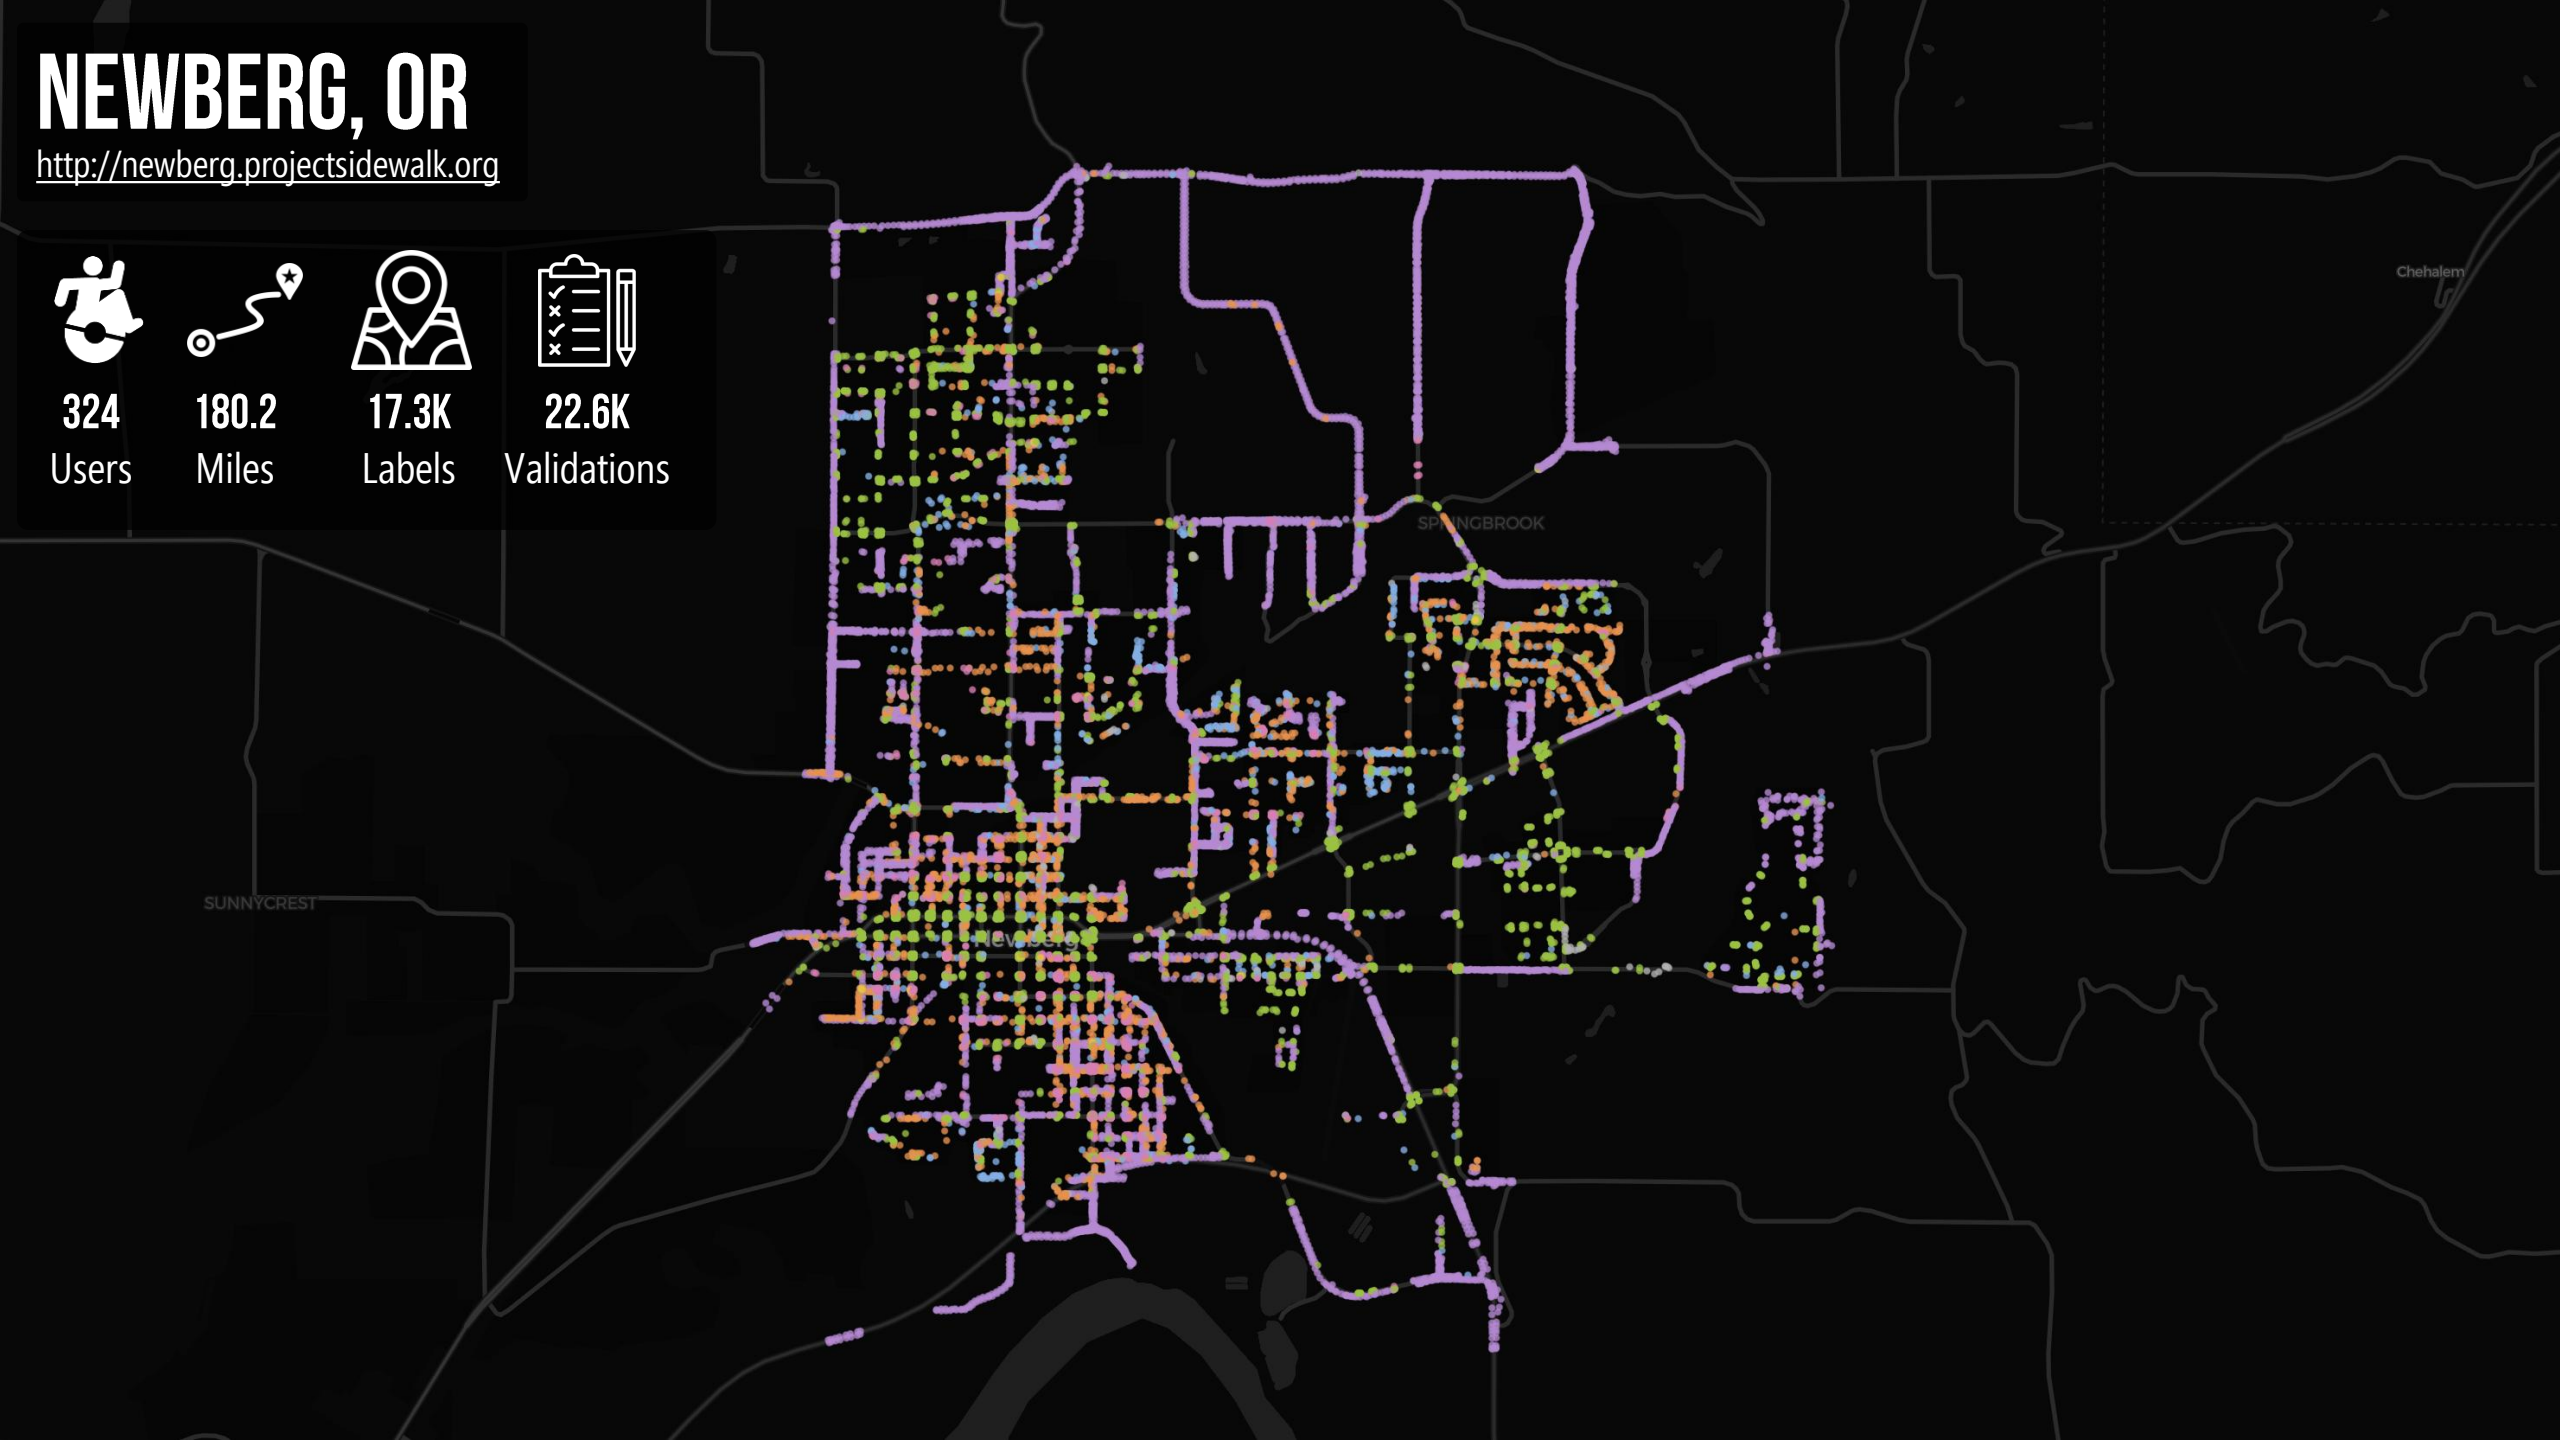
Research Engine


NEWBERG, OREGON
Disability Advocates
Local Community
Public Works Office
City of Newberg
Canada
Edmonton
Calgary
United States
Mexico
Mexico City
Cuba
Guatemala
Costa Rica
Caracas
Guyana
Colombia
Ecuador
Peru
Brazil
Fortaleza
Salvador
Belo Horizonte
Rio de Janeiro
Paraguay
Chile
North Atlantic Ocean
South Atlantic Ocean
Iceland
Sweden
Finland
Norway
Estonia
Latvia
Lithuania
Denmark
United Kingdom
London
Belarus
Poland
Germany
Ukraine
France
Austria
Moldova
Serbia
Spain
Italy
Greece
Türkiye
Georgia
Azerbaijan
Kazakhstan
Uzbekistan
Kyrgyzstan
Turkmenistan
Tajikistan
Afghanistan
Syria
Iraq
Iran
Kuwait
Pakistan
New Delhi
Karachi
India
Arabian Sea
Oman
Saudi Arabia
Yemen
Tunisia
Morocco
Algeria
Libya
Egypt
Cairo
Western Sahara
Mauritania
Mali
Niger
Chad
Sudan
Eritrea
Djibouti
Ethiopia
Somalia
Senegal
Burkina Faso
Guinea
Ivory Coast
Togo
Abidjan
Nigeria
Cameroon
Republic of the Congo
Kinshasa
Uganda
Tanzania
Luanda
Angola
Zambia
Malawi
Zimbabwe
Namibia
Madagascar
Mauritius
South Africa
Maldives
Sri Lanka
Yekaterinburg
Moscow
LABELS
100
1k
10k
100k
mapbox

NEWBERG, OR

http://newberg.projectsidewalk.org

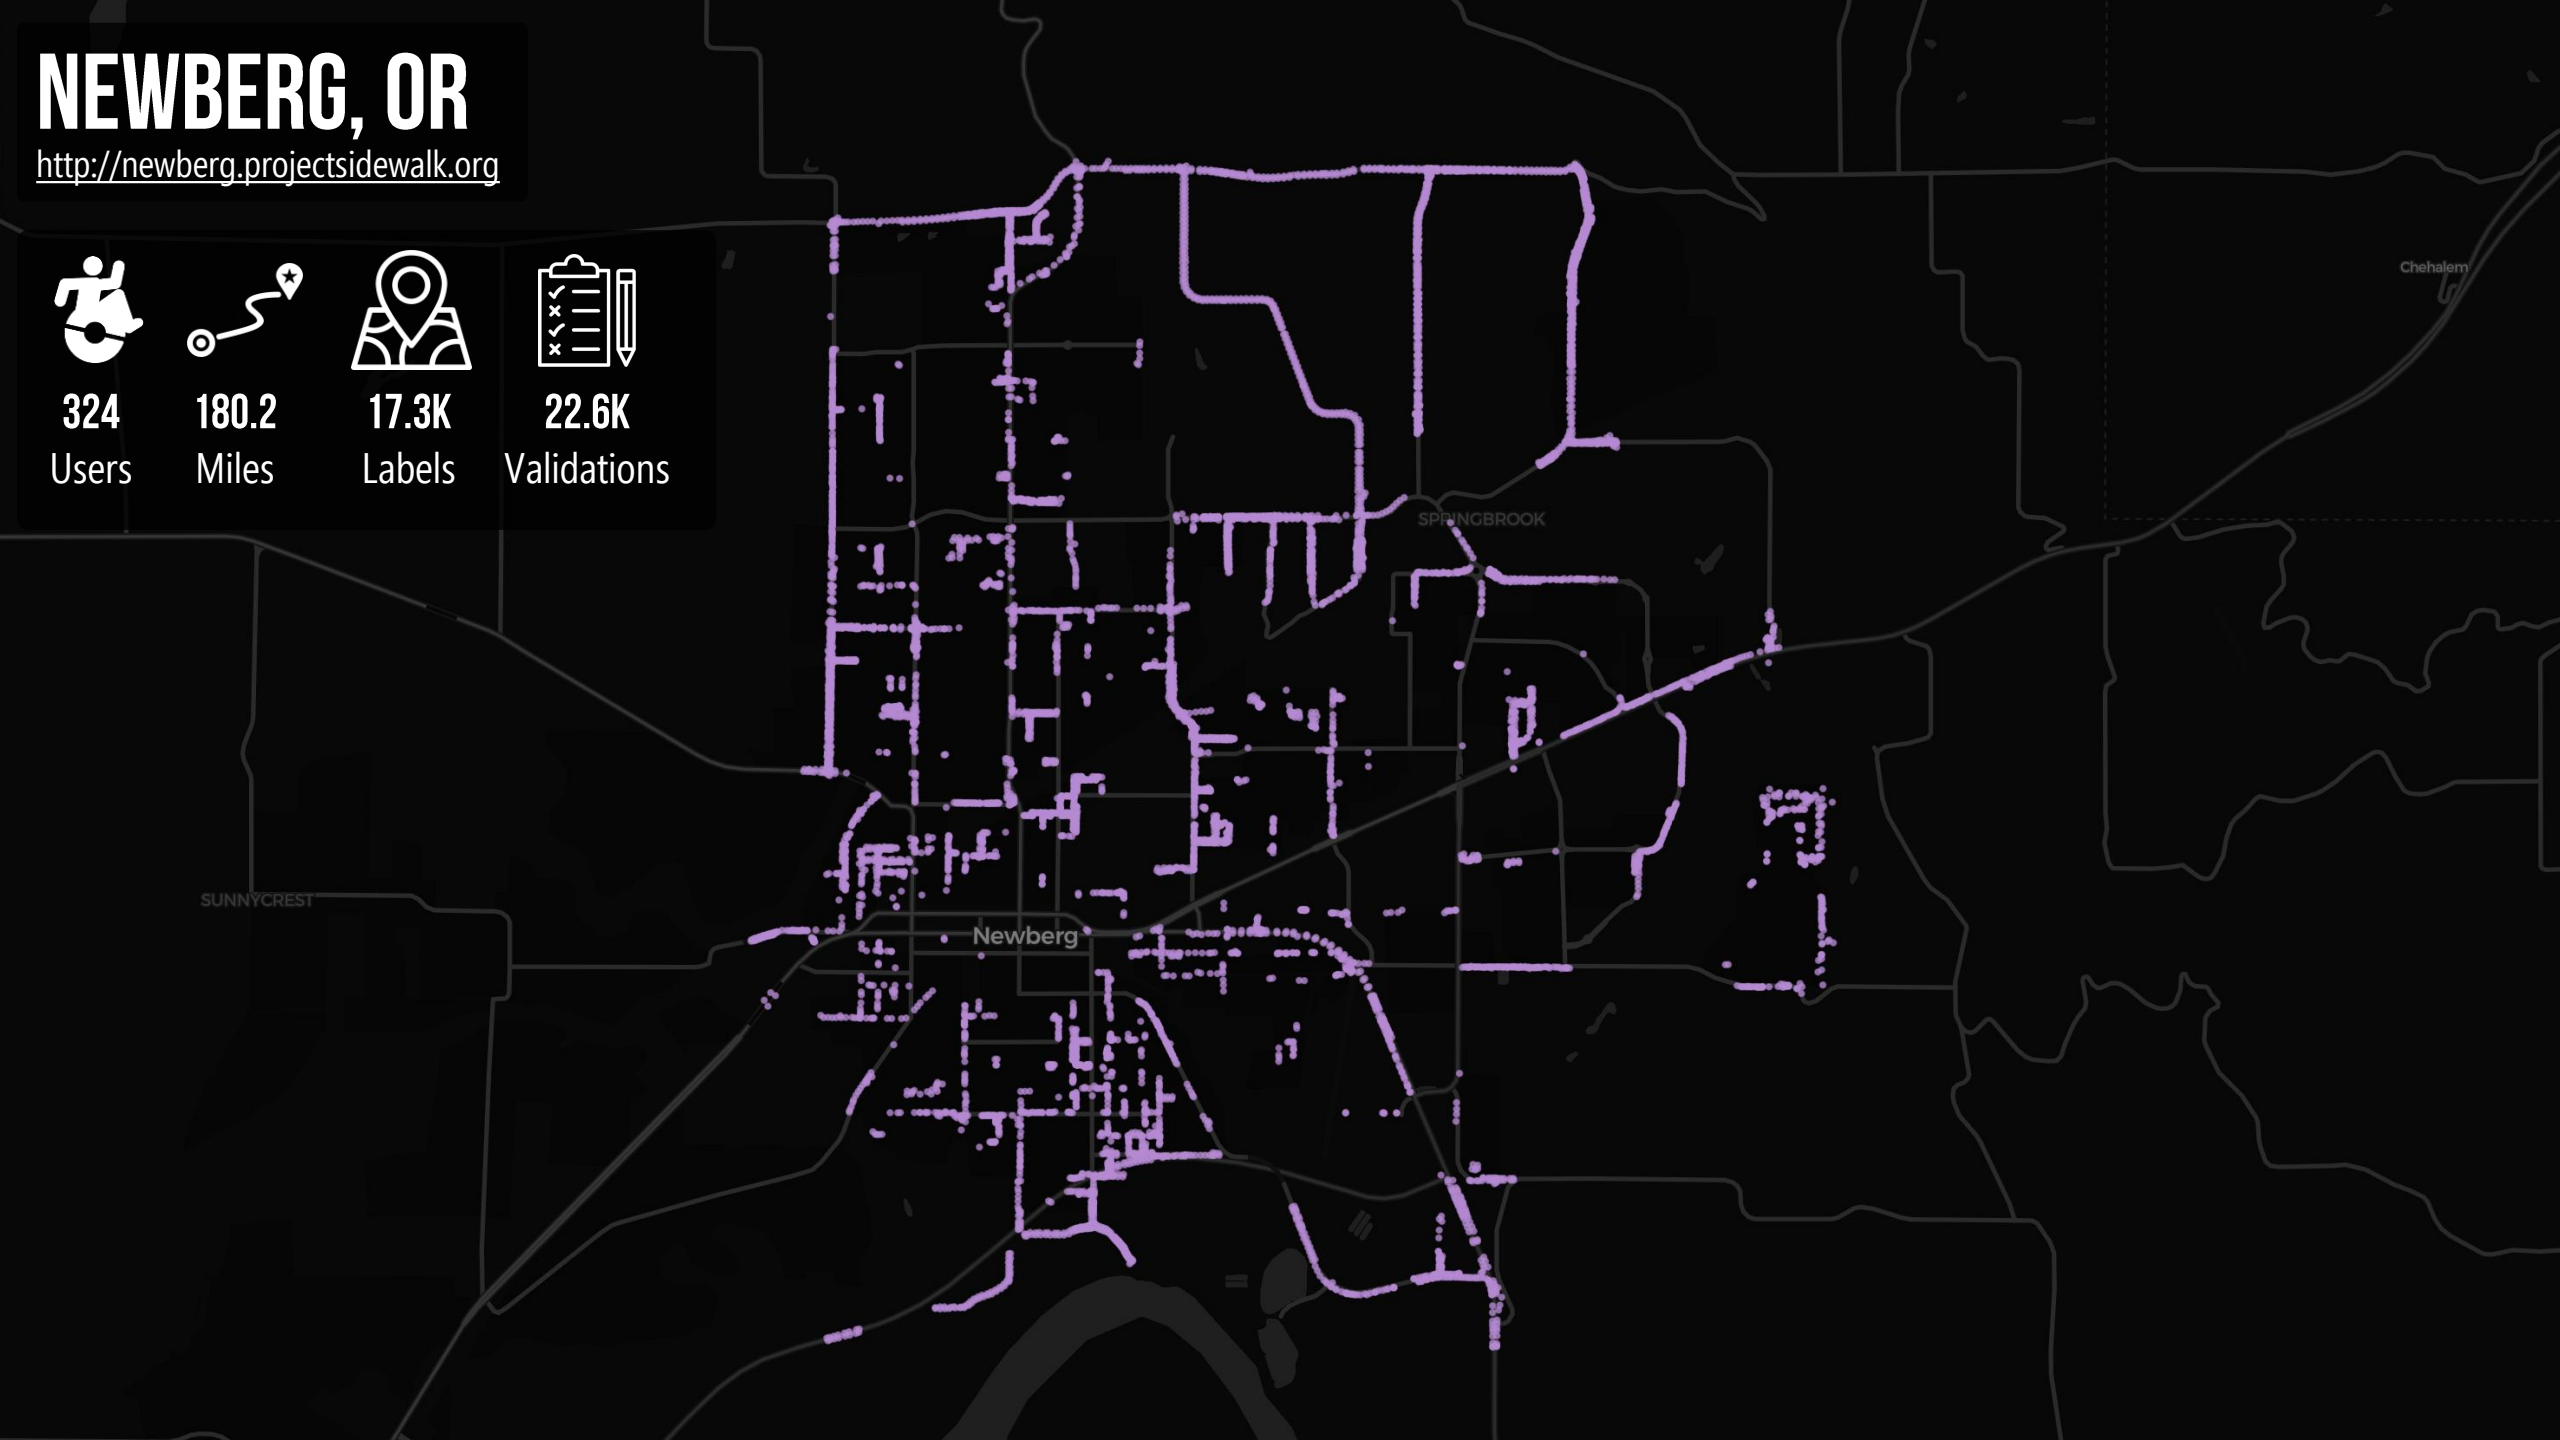
324	180.2	17.3K	22.6K
Users	Miles	Labels	Validations



NEWBERG, OR

http://newberg.projectsidewalk.org

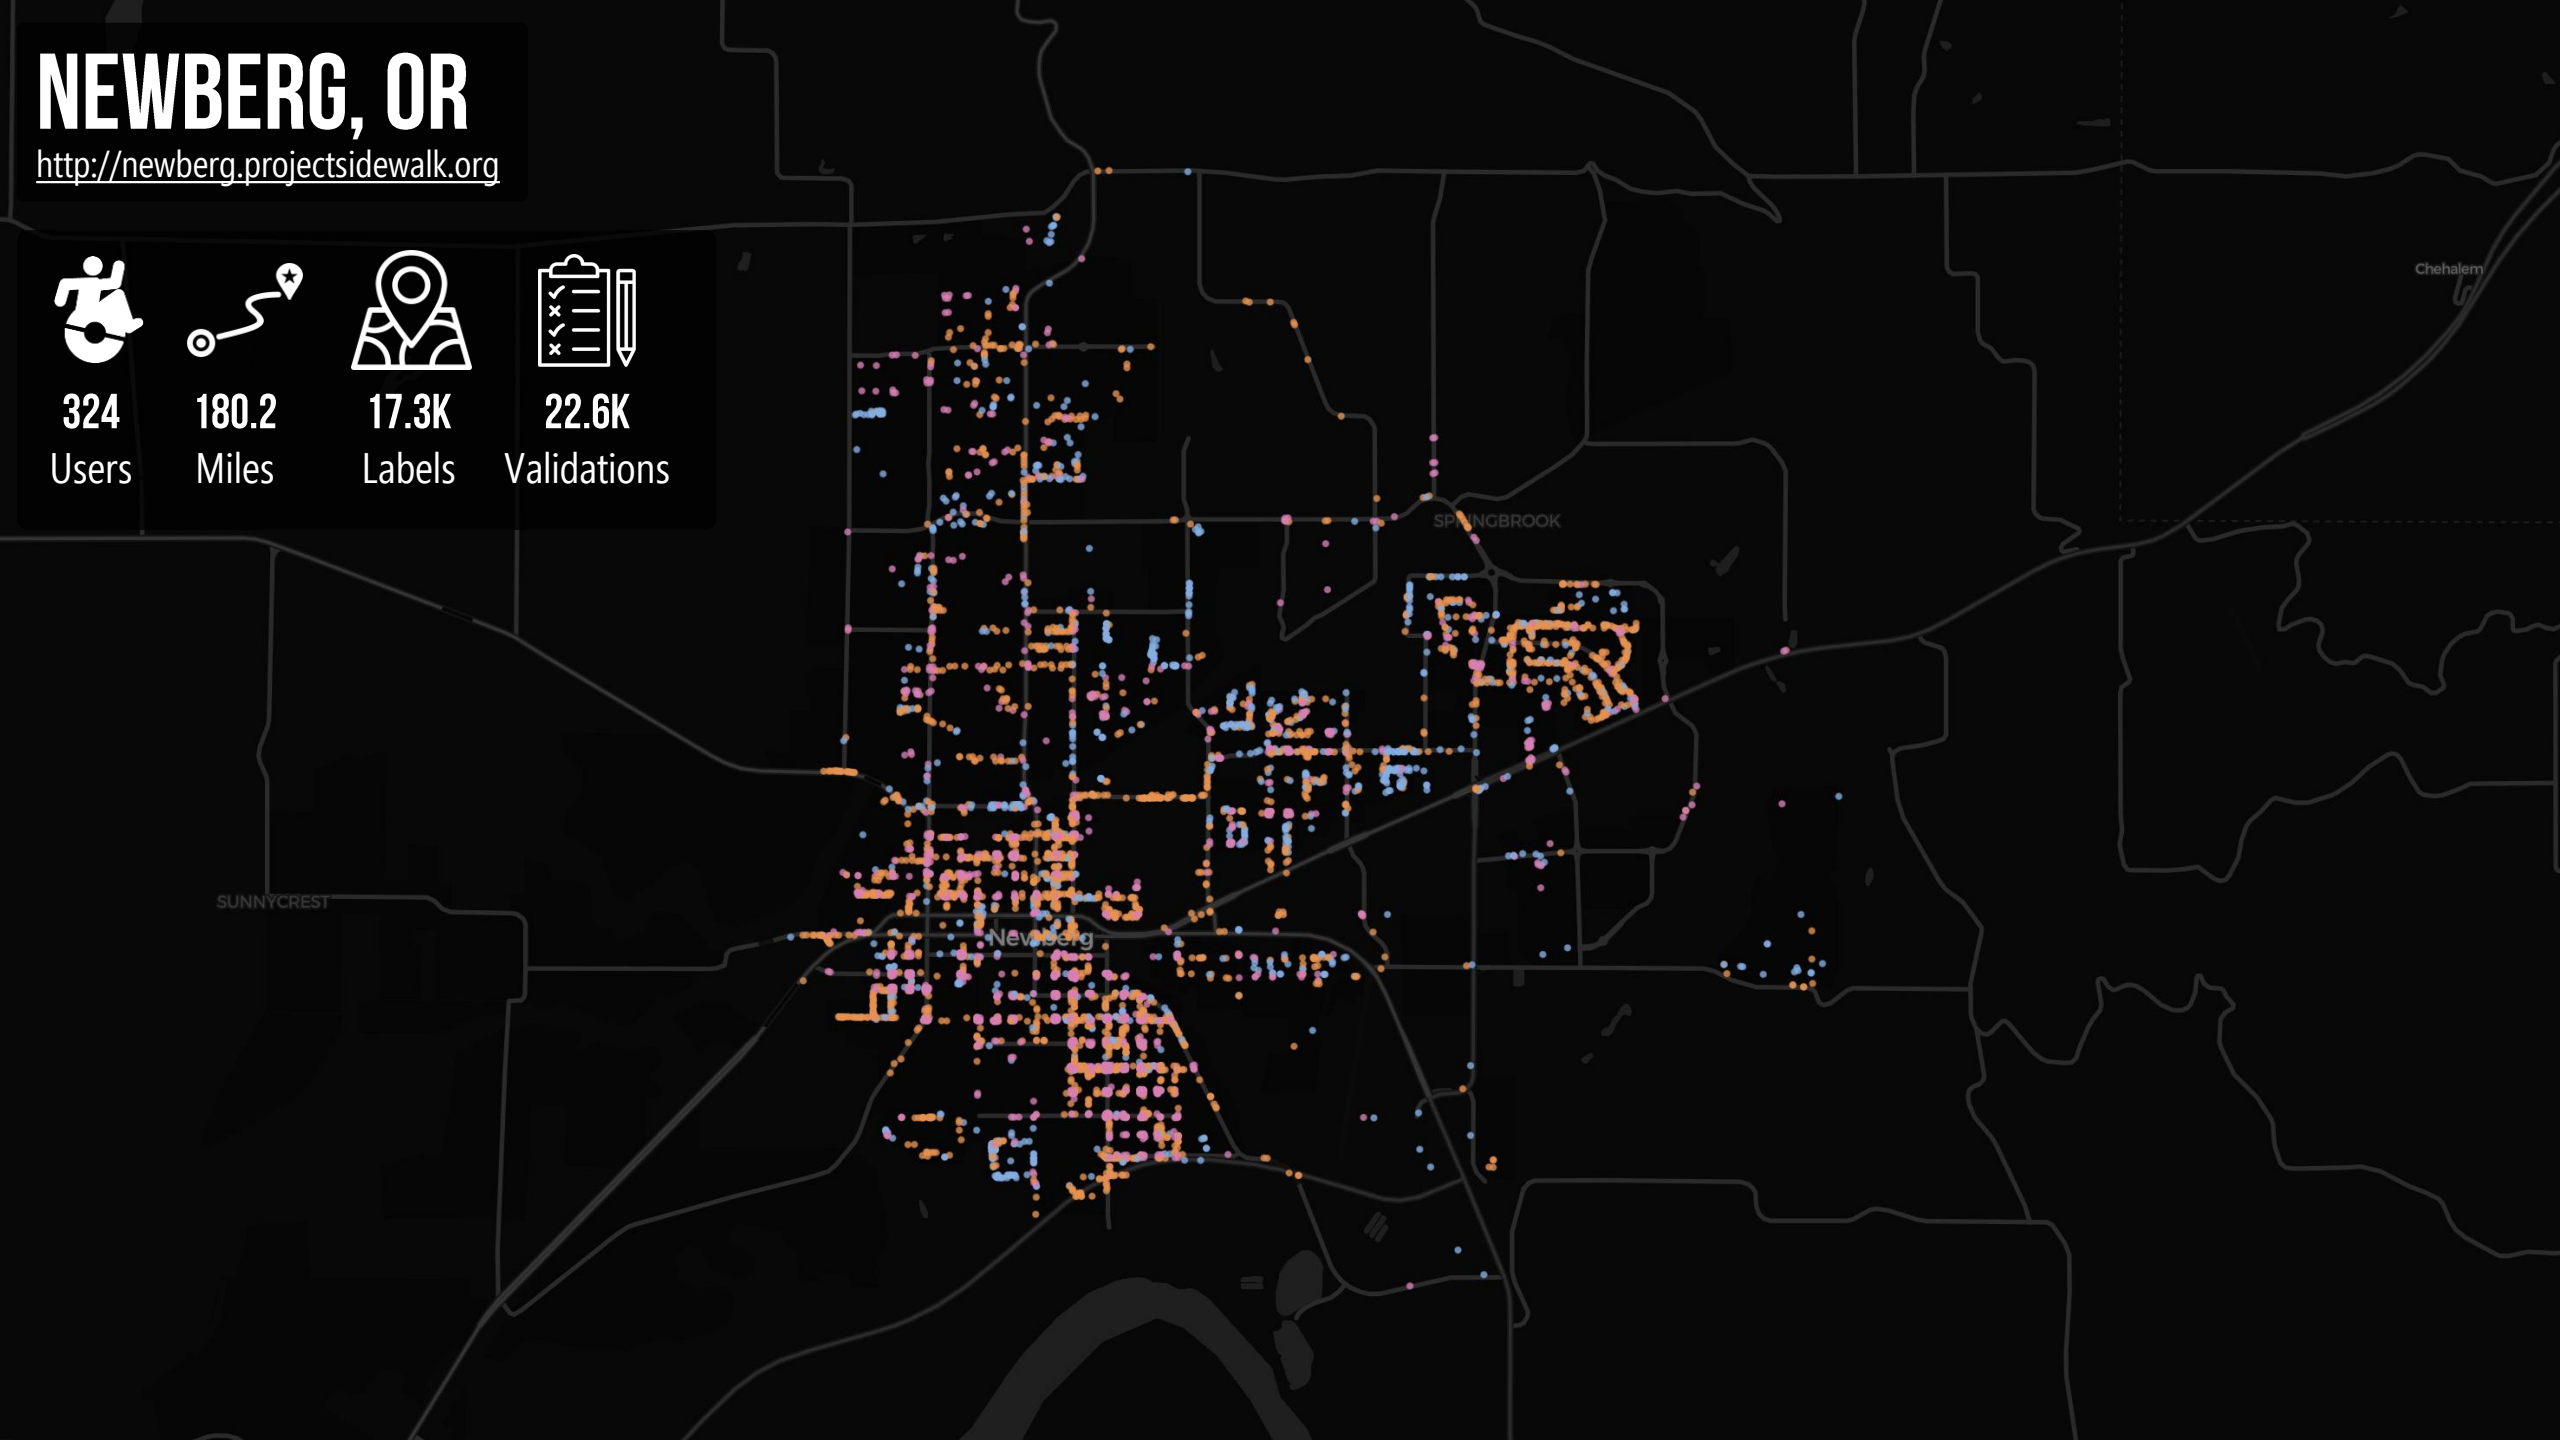
324	180.2	17.3K	22.6K
Users	Miles	Labels	Validations



NEWBERG, OR

http://newberg.projectsidewalk.org

324	180.2	17.3K	22.6K
Users	Miles	Labels	Validations





NEWBERG, ORG

NEW SIDEWALK PROGRAMS

City of Newberg

SIDEWALK REPAIR PROGRAM

RESIDENTIAL GRANT AND LOAN APPLICATION PACKET

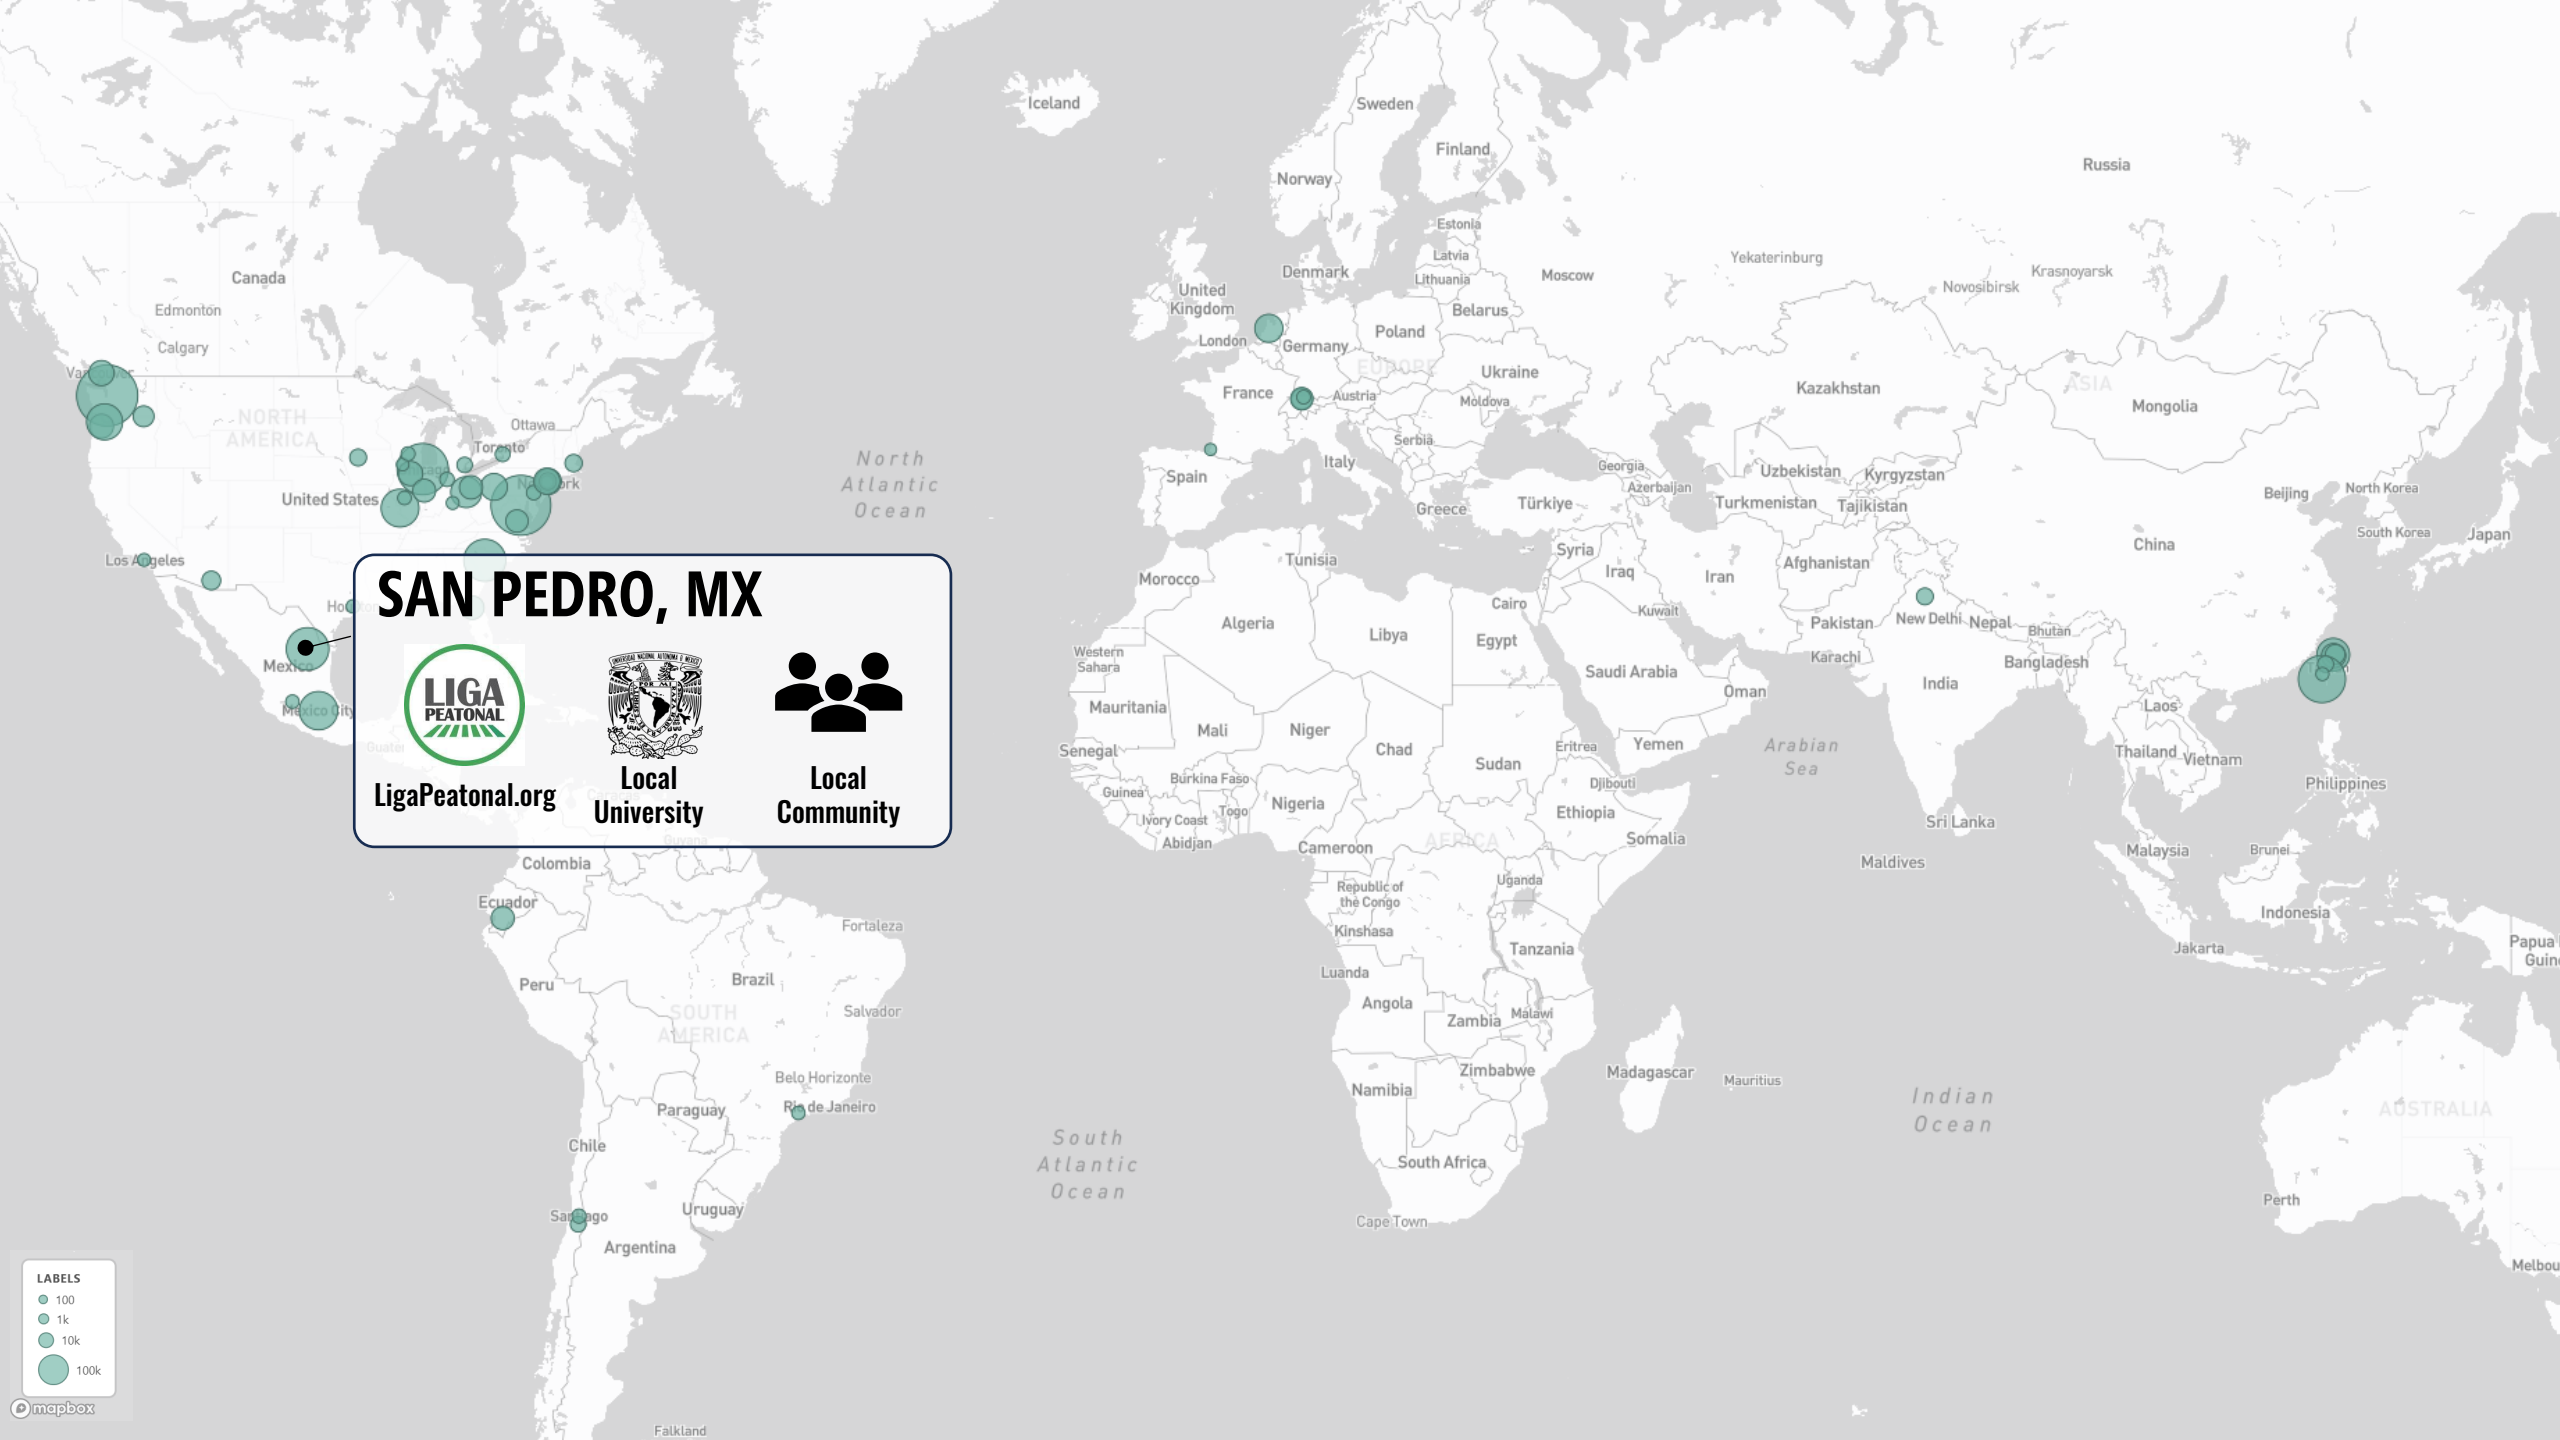


Sidewalks are for Everybody

Sidewalks make the entire community accessible,
are required under the Federal ADA regulations,
and enhance your home's value and curb appeal.

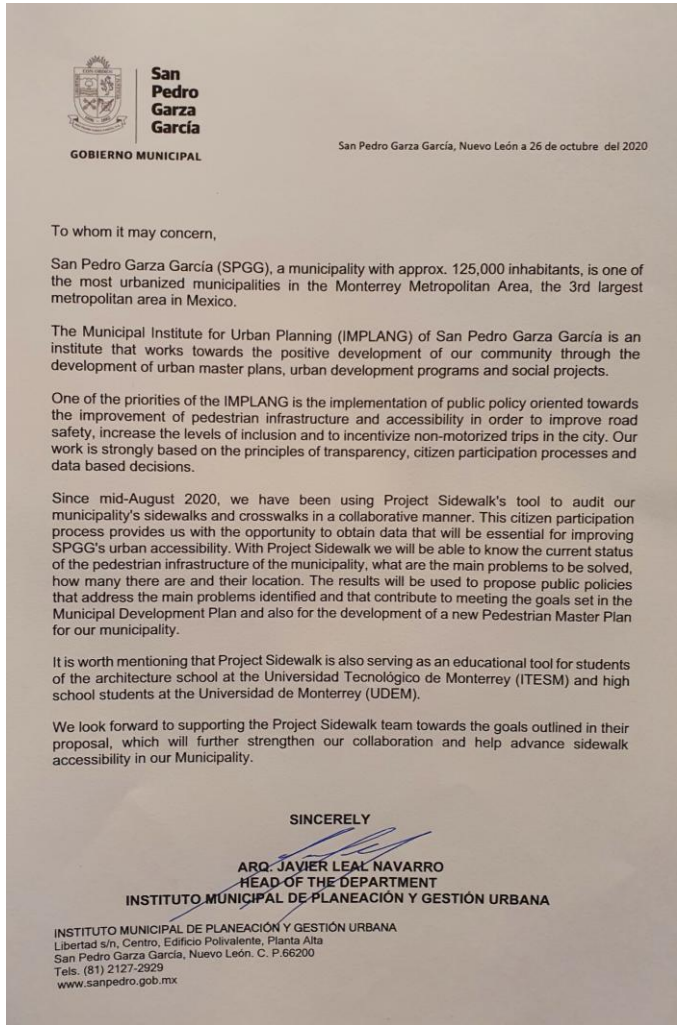
City Engineer's Office: Email: sidewalkrepairprogram@newbergoregon.gov | 503.537.1273 | PO Box 970

Newberg City Hall | 414 E. First Street | Newberg, OR 97132 | (503) 538-9421 | M-F 8:30-4:30PM


SAN PEDRO, MX
LIGA PEATONAL
LigaPeatonal.org
Local University
Local Community
LABELS
100
1k
10k
100k
mapbox

PROJECT SIDEWALK MEXICO

SAN PEDRO, MX



San Pedro Garza García
GOBIERNO MUNICIPAL

San Pedro Garza García, Nuevo León a 26 de octubre del 2020

To whom it may concern,

San Pedro Garza García (SPGG), a municipality with approx. 125,000 inhabitants, is one of the most urbanized municipalities in the Monterrey Metropolitan Area, the 3rd largest metropolitan area in Mexico.

The Municipal Institute for Urban Planning (IMPLANG) of San Pedro Garza García is an institute that works towards the positive development of our community through the development of urban master plans, urban development programs and social projects.

One of the priorities of the IMPLANG is the implementation of public policy oriented towards the improvement of pedestrian infrastructure and accessibility in order to improve road safety, increase the levels of inclusion and to incentivize non-motorized trips in the city. Our work is strongly based on the principles of transparency, citizen participation processes and data based decisions.

Since mid-August 2020, we have been using Project Sidewalk's tool to audit our municipality's sidewalks and crosswalks in a collaborative manner. This citizen participation process provides us with the opportunity to obtain data that will be essential for improving SPGG's urban accessibility. With Project Sidewalk we will be able to know the current status of the pedestrian infrastructure of the municipality, what are the main problems to be solved, how many there are and their location. The results will be used to propose public policies that address the main problems identified and that contribute to meeting the goals set in the Municipal Development Plan and also for the development of a new Pedestrian Master Plan for our municipality.

It is worth mentioning that Project Sidewalk is also serving as an educational tool for students of the architecture school at the Universidad Tecnológico de Monterrey (ITESM) and high school students at the Universidad de Monterrey (UDEM).

We look forward to supporting the Project Sidewalk team towards the goals outlined in their proposal, which will further strengthen our collaboration and help advance sidewalk accessibility in our Municipality.

SINCERELY

ARQ. JAVIER LEAL NAVARRO
HEAD OF THE DEPARTMENT
INSTITUTO MUNICIPAL DE PLANEACIÓN Y GESTIÓN URBANA

INSTITUTO MUNICIPAL DE PLANEACIÓN Y GESTIÓN URBANA
Libertad s/n, Centro, Edificio Polivalente, Planta Alta
San Pedro Garza García, Nuevo León. C. P.66200
Tels. (81) 2127-2929
www.sanpedro.gob.mx

Project Sidewalk provides us with data that is **essential to improving San Pedro's urban accessibility**. With Project Sidewalk, we **know the main problems** to be solved, how many problems there are, and their location... The results will be used to inform a **new Pedestrian Master Plan** for our municipality.

San Pedro Garza García
GOBIERNO MUNICIPAL

ETIQUETAS DE RAMPA PEATONAL | ETIQUETAS DE BANQUETA | OTRAS ETIQUETAS DE ACCESIBILIDAD

Explorar | Rampa peatonal (C) | Rampa peatonal ausente (M) | Obstáculo en la banqueta | Problema en superficie | No hay banqueta | Cruce peatonal marcado (W) | Semáforo peatonal | Otro

Acercar (Z)

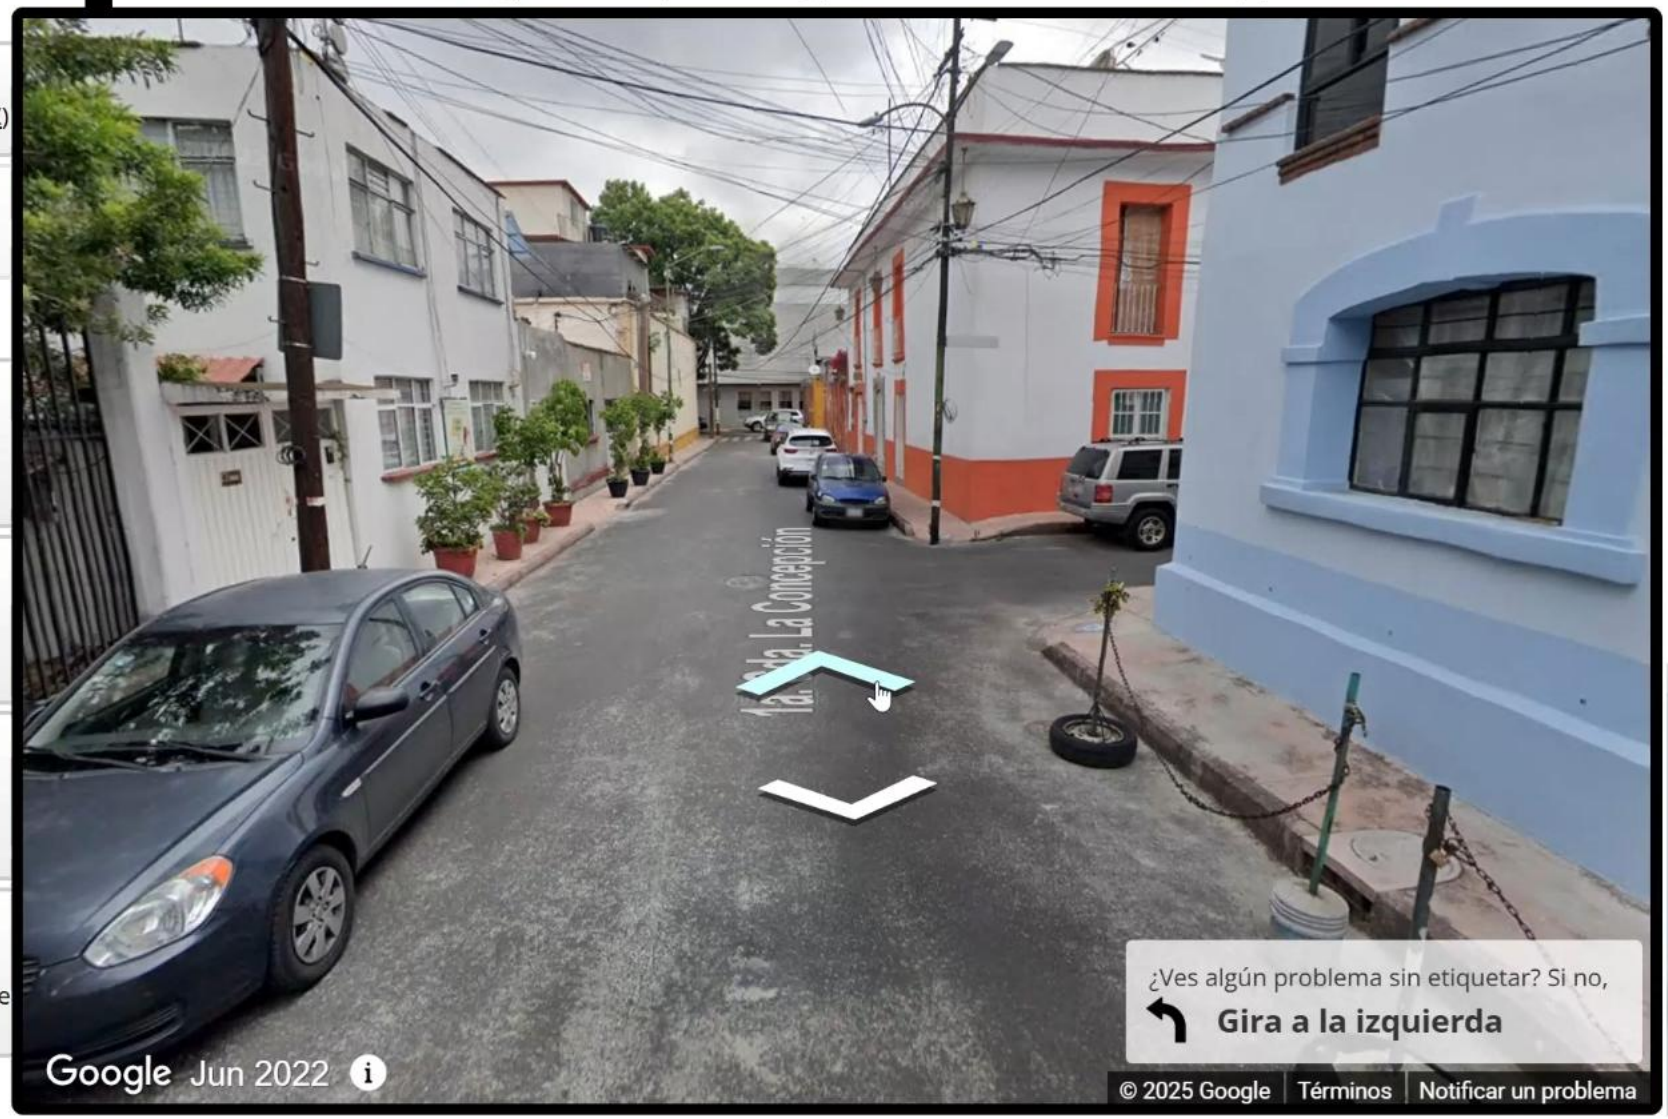
Alejar

Sonido

Saltar

Atorado

Retroalimentación



Estadísticas generales

2,11 kilómetros

233 etiquetas

92,1% de precisión

Colonia actual

Centro de Azcapotzalco, CDMX

0,01 kilómetros

10 etiquetas

Misión actual

Explora 75 m y encuentra todos los atributos de accesibilidad de la banqueta

25% completado

3 rampas peatonales

2 problemas en superficie

3 rampas peatonales ausentes

0 no hay banquetas

2 obstáculos

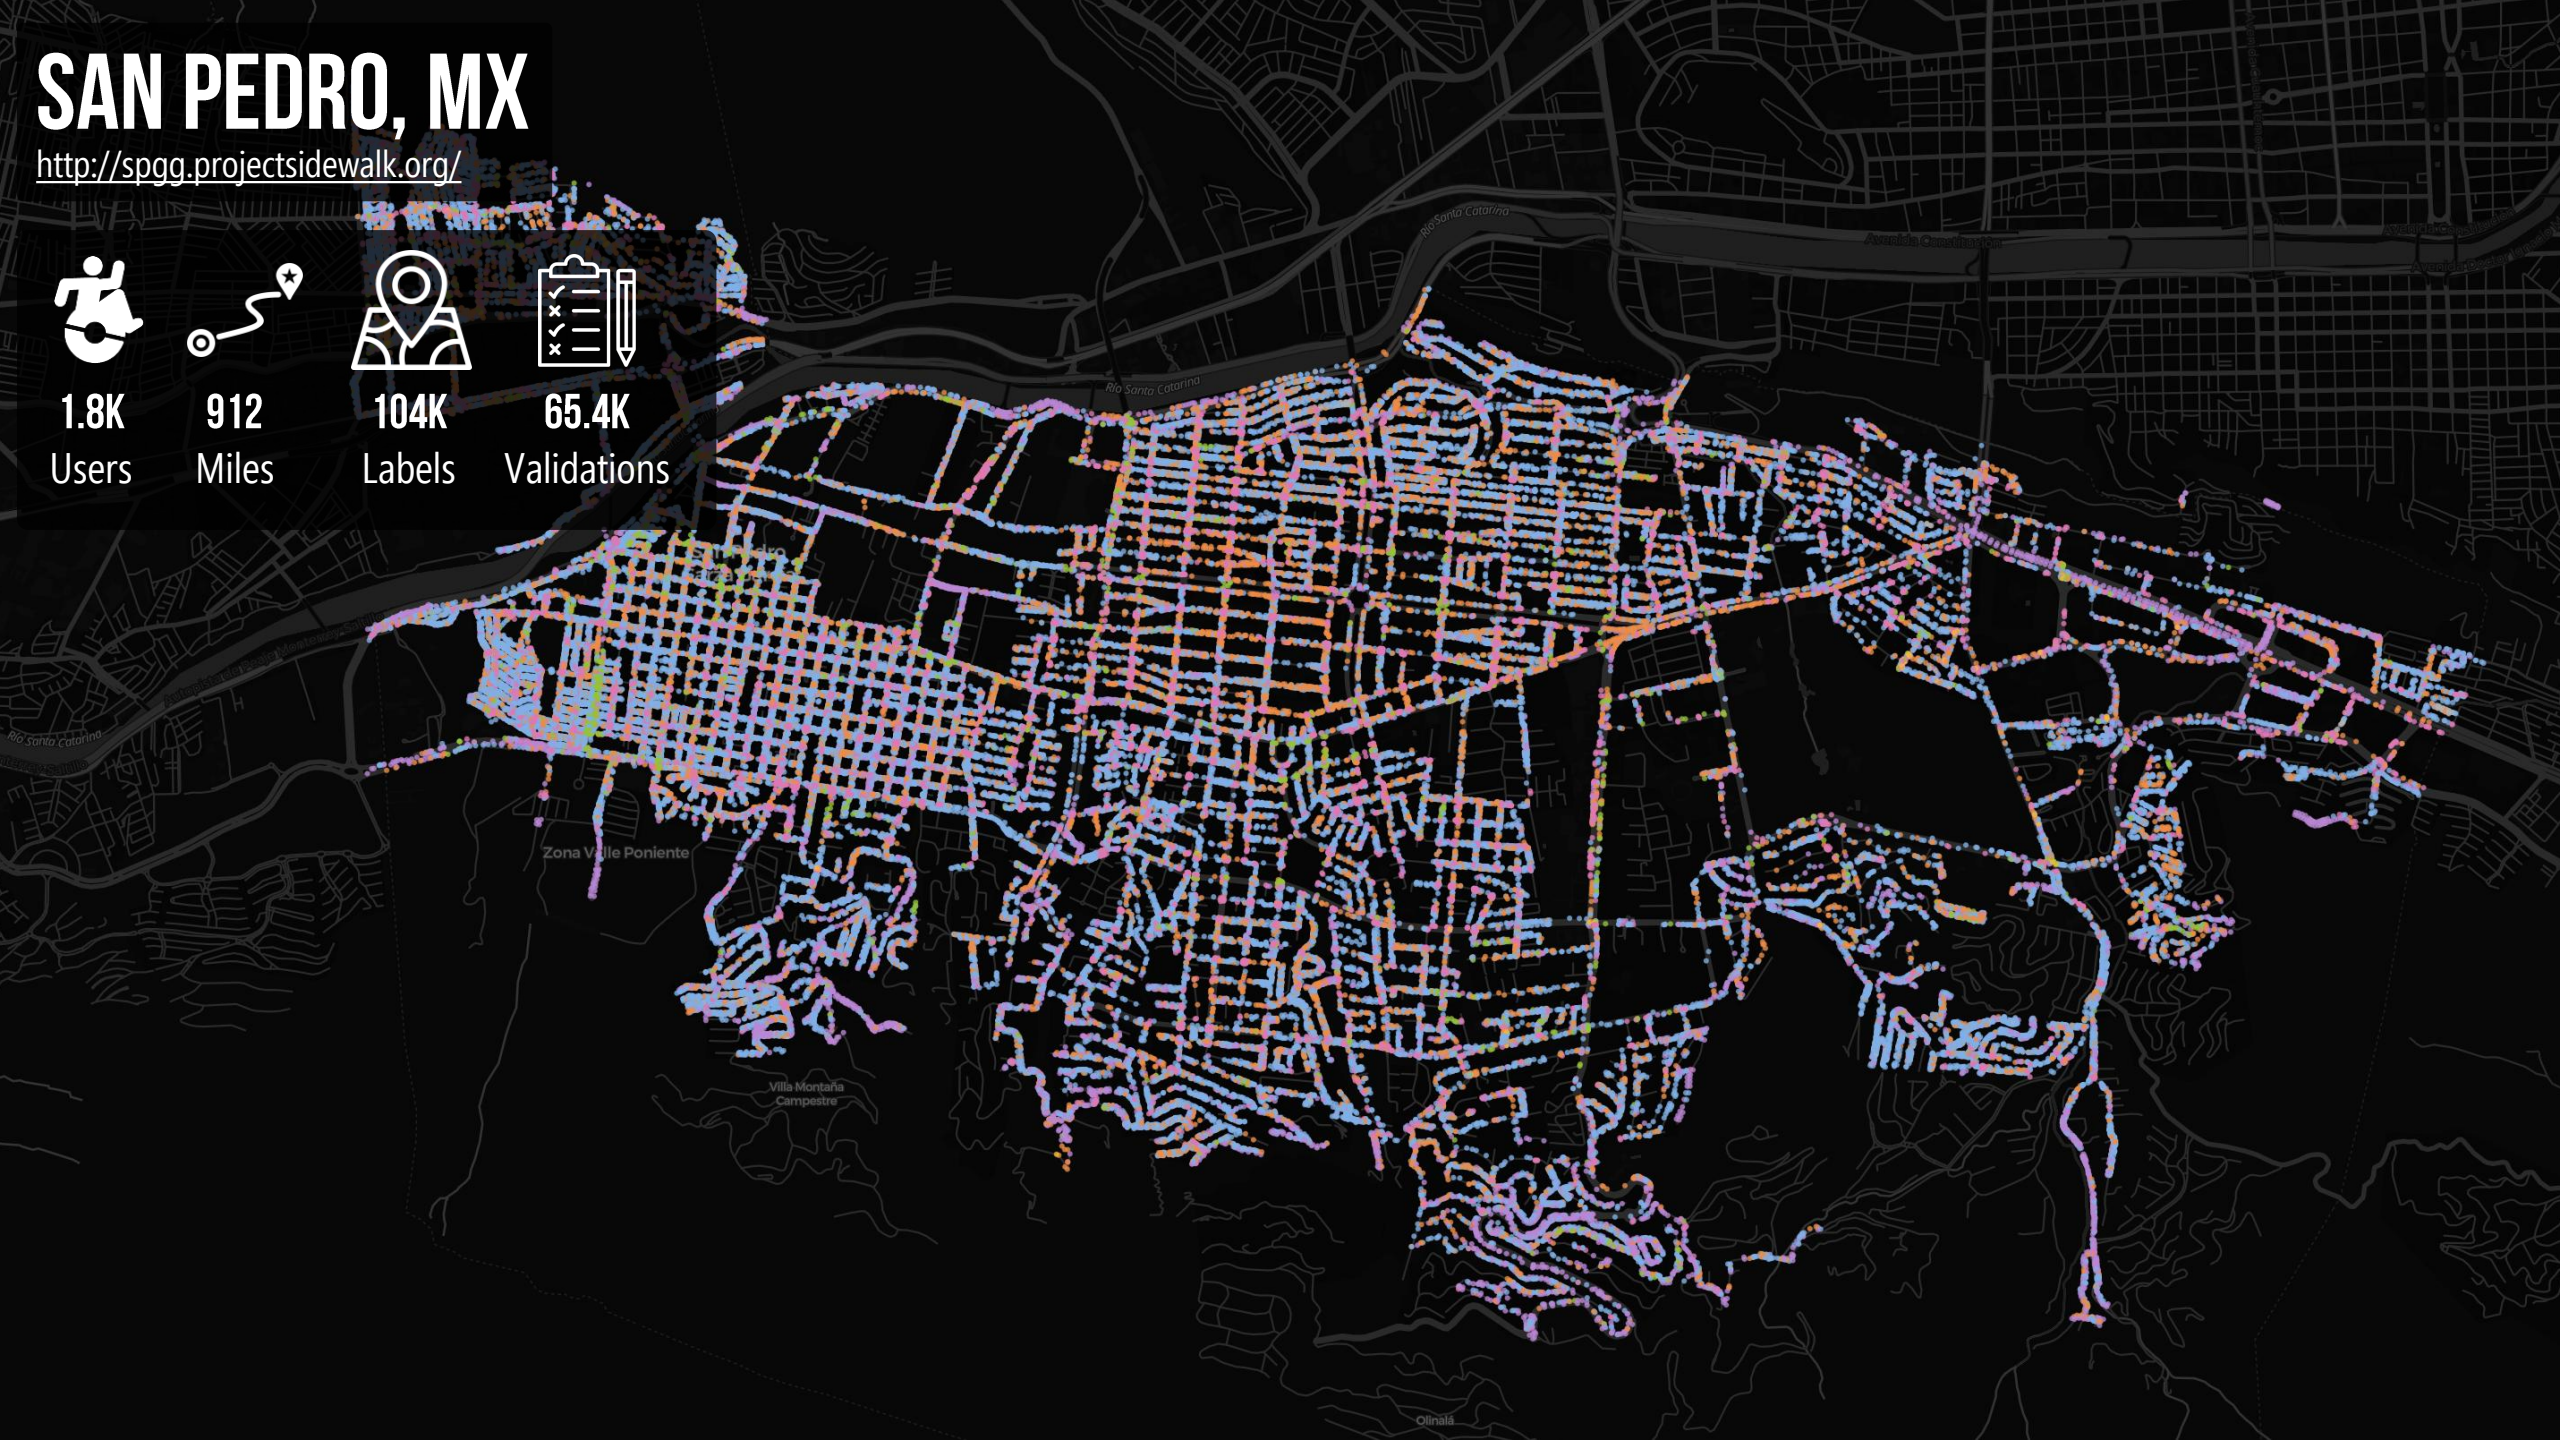
0 otros



SAN PEDRO, MX

http://spgg.projectsidewalk.org/

1.8K	912	104K	65.4K
Users	Miles	Labels	Validations



SAN PEDRO, MX

http://spgg.projectsidewalk.org/

1.8K	912	104K	65.4K
Users	Miles	Labels	Validations



CURB RAMPS

High Severity

Not passable


Narrow + obstacle


Not enough landing space


Points into traffic


Narrow


No friction/tactile strip


Not level with street


Steep + obstacle


Steep + narrow


Poor design



CURB RAMPS

High Severity

Not passable


Narrow + obstacle


Not enough landing space


Points into traffic


Narrow


No friction/tactile strip


Not level with street


Steep + obstacle


Steep + narrow


Poor design


Passable
1
2
3
4
5
Not Passable
Can only be entered from one direction. Has a stee
Add Tags:
narrow
points into traffic
missing friction strip
steep
not enough landing space
Temporary (e.g., construction)
OK
Do you see any unlabeled problems? If not,
Turn slightly left
Google
© 2020 Google
Terms of Use
Report a problem



ONU HABITAT
POR UN MEJOR FUTURO URBANO

ENFOQUE ACTIVISTA

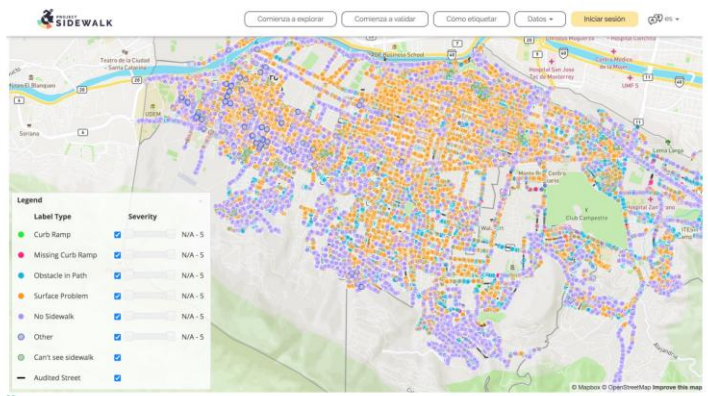
E

EvaluANDO: del activismo peatonal a la colaboración comunitaria para el registro de obstáculos en las banquetas

Escrito por
Claudina de Gyves y Ana Rodríguez

Ubicación
San Pedro Garza García, México

Palabras clave
activismo peatonal, movilidad sostenible, infraestructura peatonal, participación remota

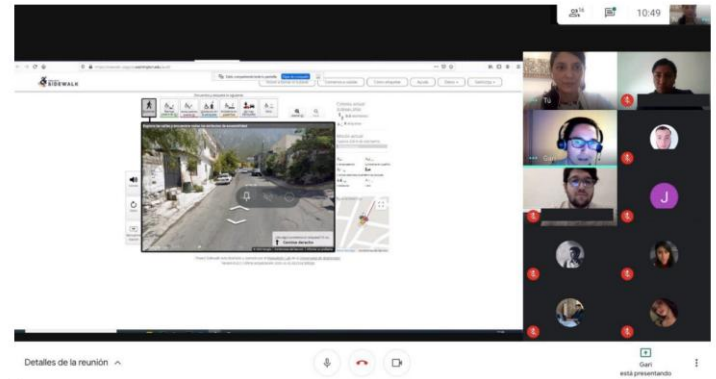


>>
Fotografía 2. Mapa de etiquetas en Project Sidewalk
Fuente: Liga Peatonal (2021).

La vinculación fue posible gracias a que ya había un historial de activismo peatonal en la metrópoli y el acercamiento no fue solo con la Liga Peatonal como ONG, sino de la mano de Makeability Lab, un actor técnico-académico que mostró disposición a contextualizar su plataforma a las necesidades de las calles mexicanas. Aunado a este proceso, la situación por la COVID-19 detonó una serie de intervenciones en el espacio público por parte del municipio de San Pedro Garza García, enfocadas en promover la movilidad sostenible, destacando las ciclovías emergentes y la aceleración de otros proyectos en el espacio público que estaban en puerta. Todo esto generó un escenario adecuado para la colaboración de EvaluANDO SPGG, en la que todos los actores involucrados estaban conscientes de la importancia de contar con información precisa sobre las condiciones de las calles en el municipio.

Recientemente, en mayo de 2021, tras 9 meses de trabajo y con la participación de 1099 personas se lograron cubrir los 570.2 km de vialidades que tiene el municipio de SPGG y se generaron 105 177 etiquetas (Makeability Lab, 2021) en un ejercicio inédito a nivel nacional de participación ciudadana para ubicar los obstáculos de movilidad peatonal.

El caso de EvaluANDO SPGG destaca no solo por haber completado el mapeo del municipio y ser resultado de una colaboración multisectorial entre gobierno local, sociedad civil y academia, sino porque los resultados son ahora insumos valiosos del municipio para la creación de nuevos planes y proyectos. Los planes en proceso de elaboración, tanto de movilidad activa como de seguridad vial, con los resultados de EvaluANDO, ayudarán a identificar estrategias aterrizadas a la realidad y fomentar una mayor participación ciudadana, al involucrar a la población desde su diagnóstico y permitir la descarga de los datos generados en formato editable.



>>
Fotografía 4. Mapatón San Pedro Garza García
Fuente: Liga Peatonal (2020).

En el proceso de levantamiento de información, Liga Peatonal trató de complementar el trabajo asincrónico e individual con cuatro eventos donde varias personas se conectaban de manera simultánea a probar la herramienta y resolver dudas sobre su funcionamiento. Se convocó a dos sesiones dirigidas a las personas ciudadanas del municipio, con el nombre de Mapatones, y otras dos orientadas a estudiantes universitarios, en formato de talleres en los que se les introducía al tema de movilidad peatonal y donde se generaron propuestas para atender los problemas principales.

Si bien este proceso ha permitido el involucramiento de adolescentes y jóvenes en el análisis crítico de su entorno urbano, todavía presenta oportunidades de mejora en la inclusión de personas que no tienen acceso a dispositivos de internet. Ante esta situación, sería conveniente explorar el ejercicio analógico del mapeo en sitio con herramientas impresas y más con el fin de fortalecer la convivencia vecinal y promover la organización, que con la precisión del levantamiento. En estos escenarios de atención a la población en condición de vulnerabilidad, tal vez destaquen otros elementos a mapear no tan relacionados con ser un obstáculo en las banquetas, sino ya más encaminados a una ausencia de infraestructura como la misma pavimentación de las calles, las banquetas o la falta de conectividad con otros sectores.

Además, en las reflexiones en torno al uso de la herramienta y las necesidades para el diagnóstico urbano incluyente, se identificó como área de oportunidad un mapeo con perspectiva de género, que pudiera incluir no solamente obstáculos en los trayectos identificados, sino también situaciones y elementos propios de la infraestructura que provocan una sensación de inseguridad, pero que no representan como tal un obstáculo, como sí lo hacen los muros ciegos, la falta de luminarias, los recovecos o terrenos baldíos.

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

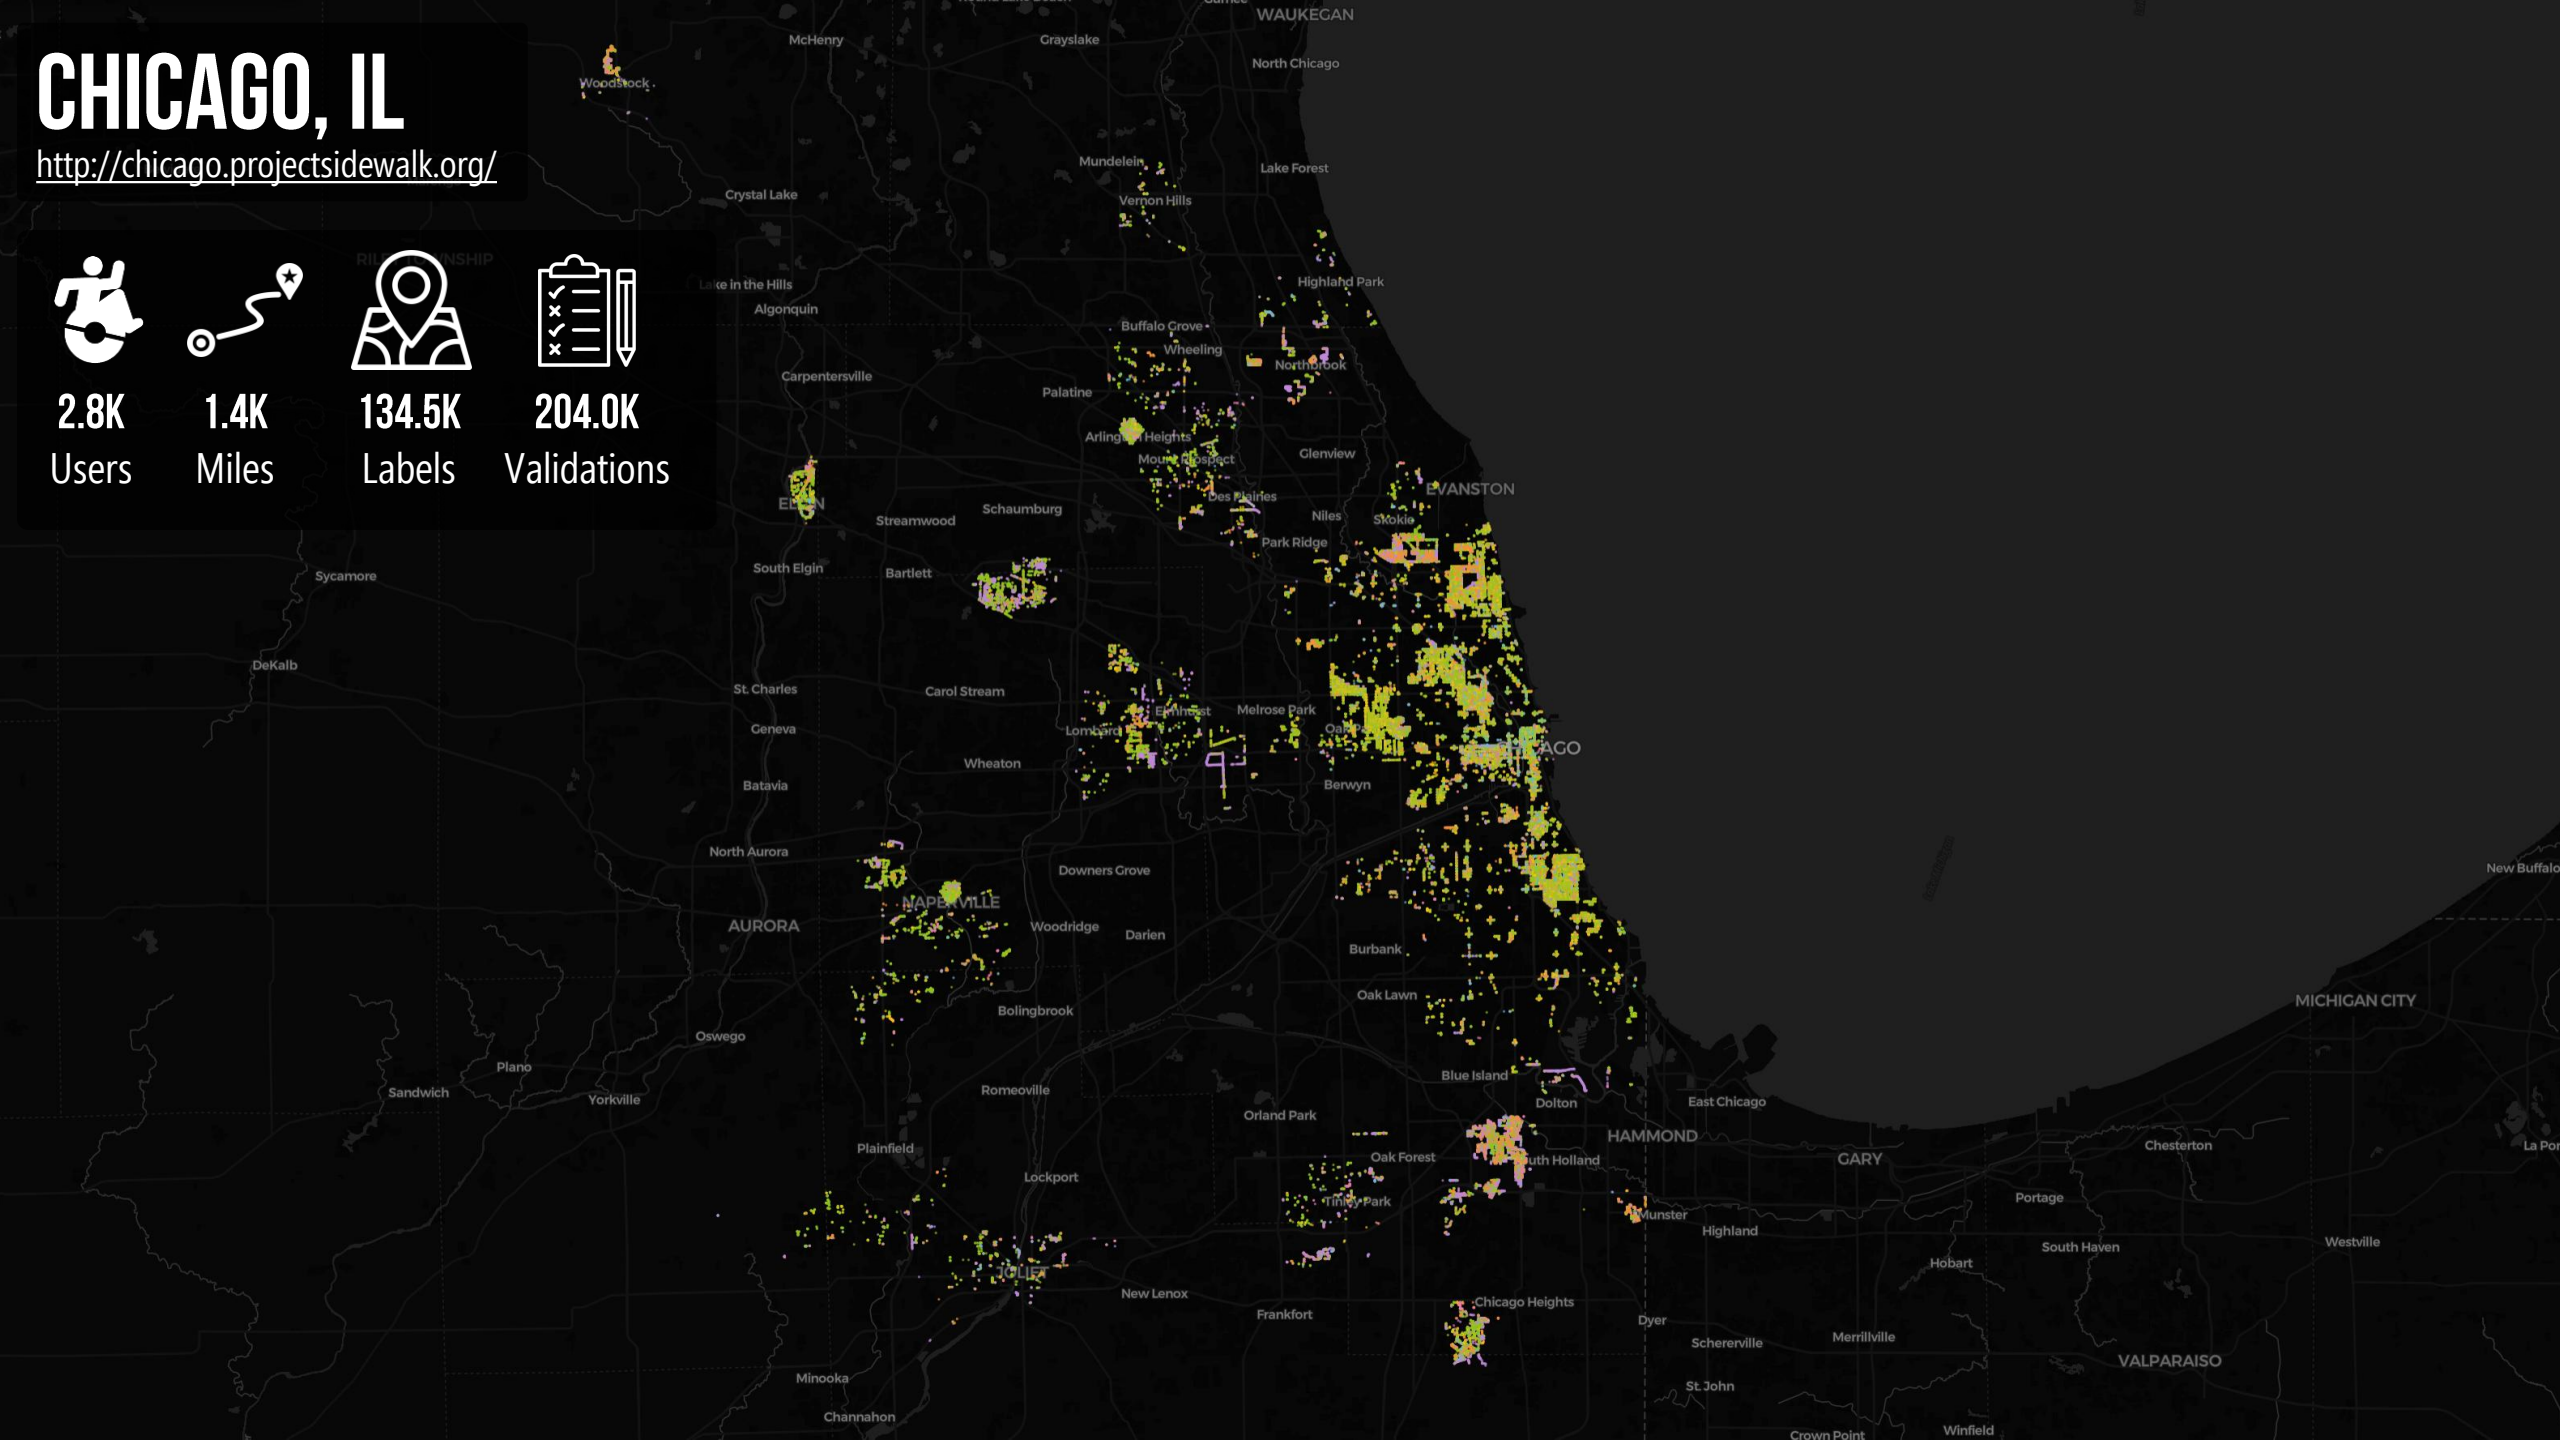
Service Learning

Research Engine

CHICAGO, IL

http://chicago.projectsidewalk.org/

2.8K	1.4K	134.5K	204.0K
Users	Miles	Labels	Validations



Project Sidewalk helped us **make better decisions** in how to **prioritize our $1.5m infrastructure spending** as well as alert appropriate **city departments** to **high-priority accessibility issues.**

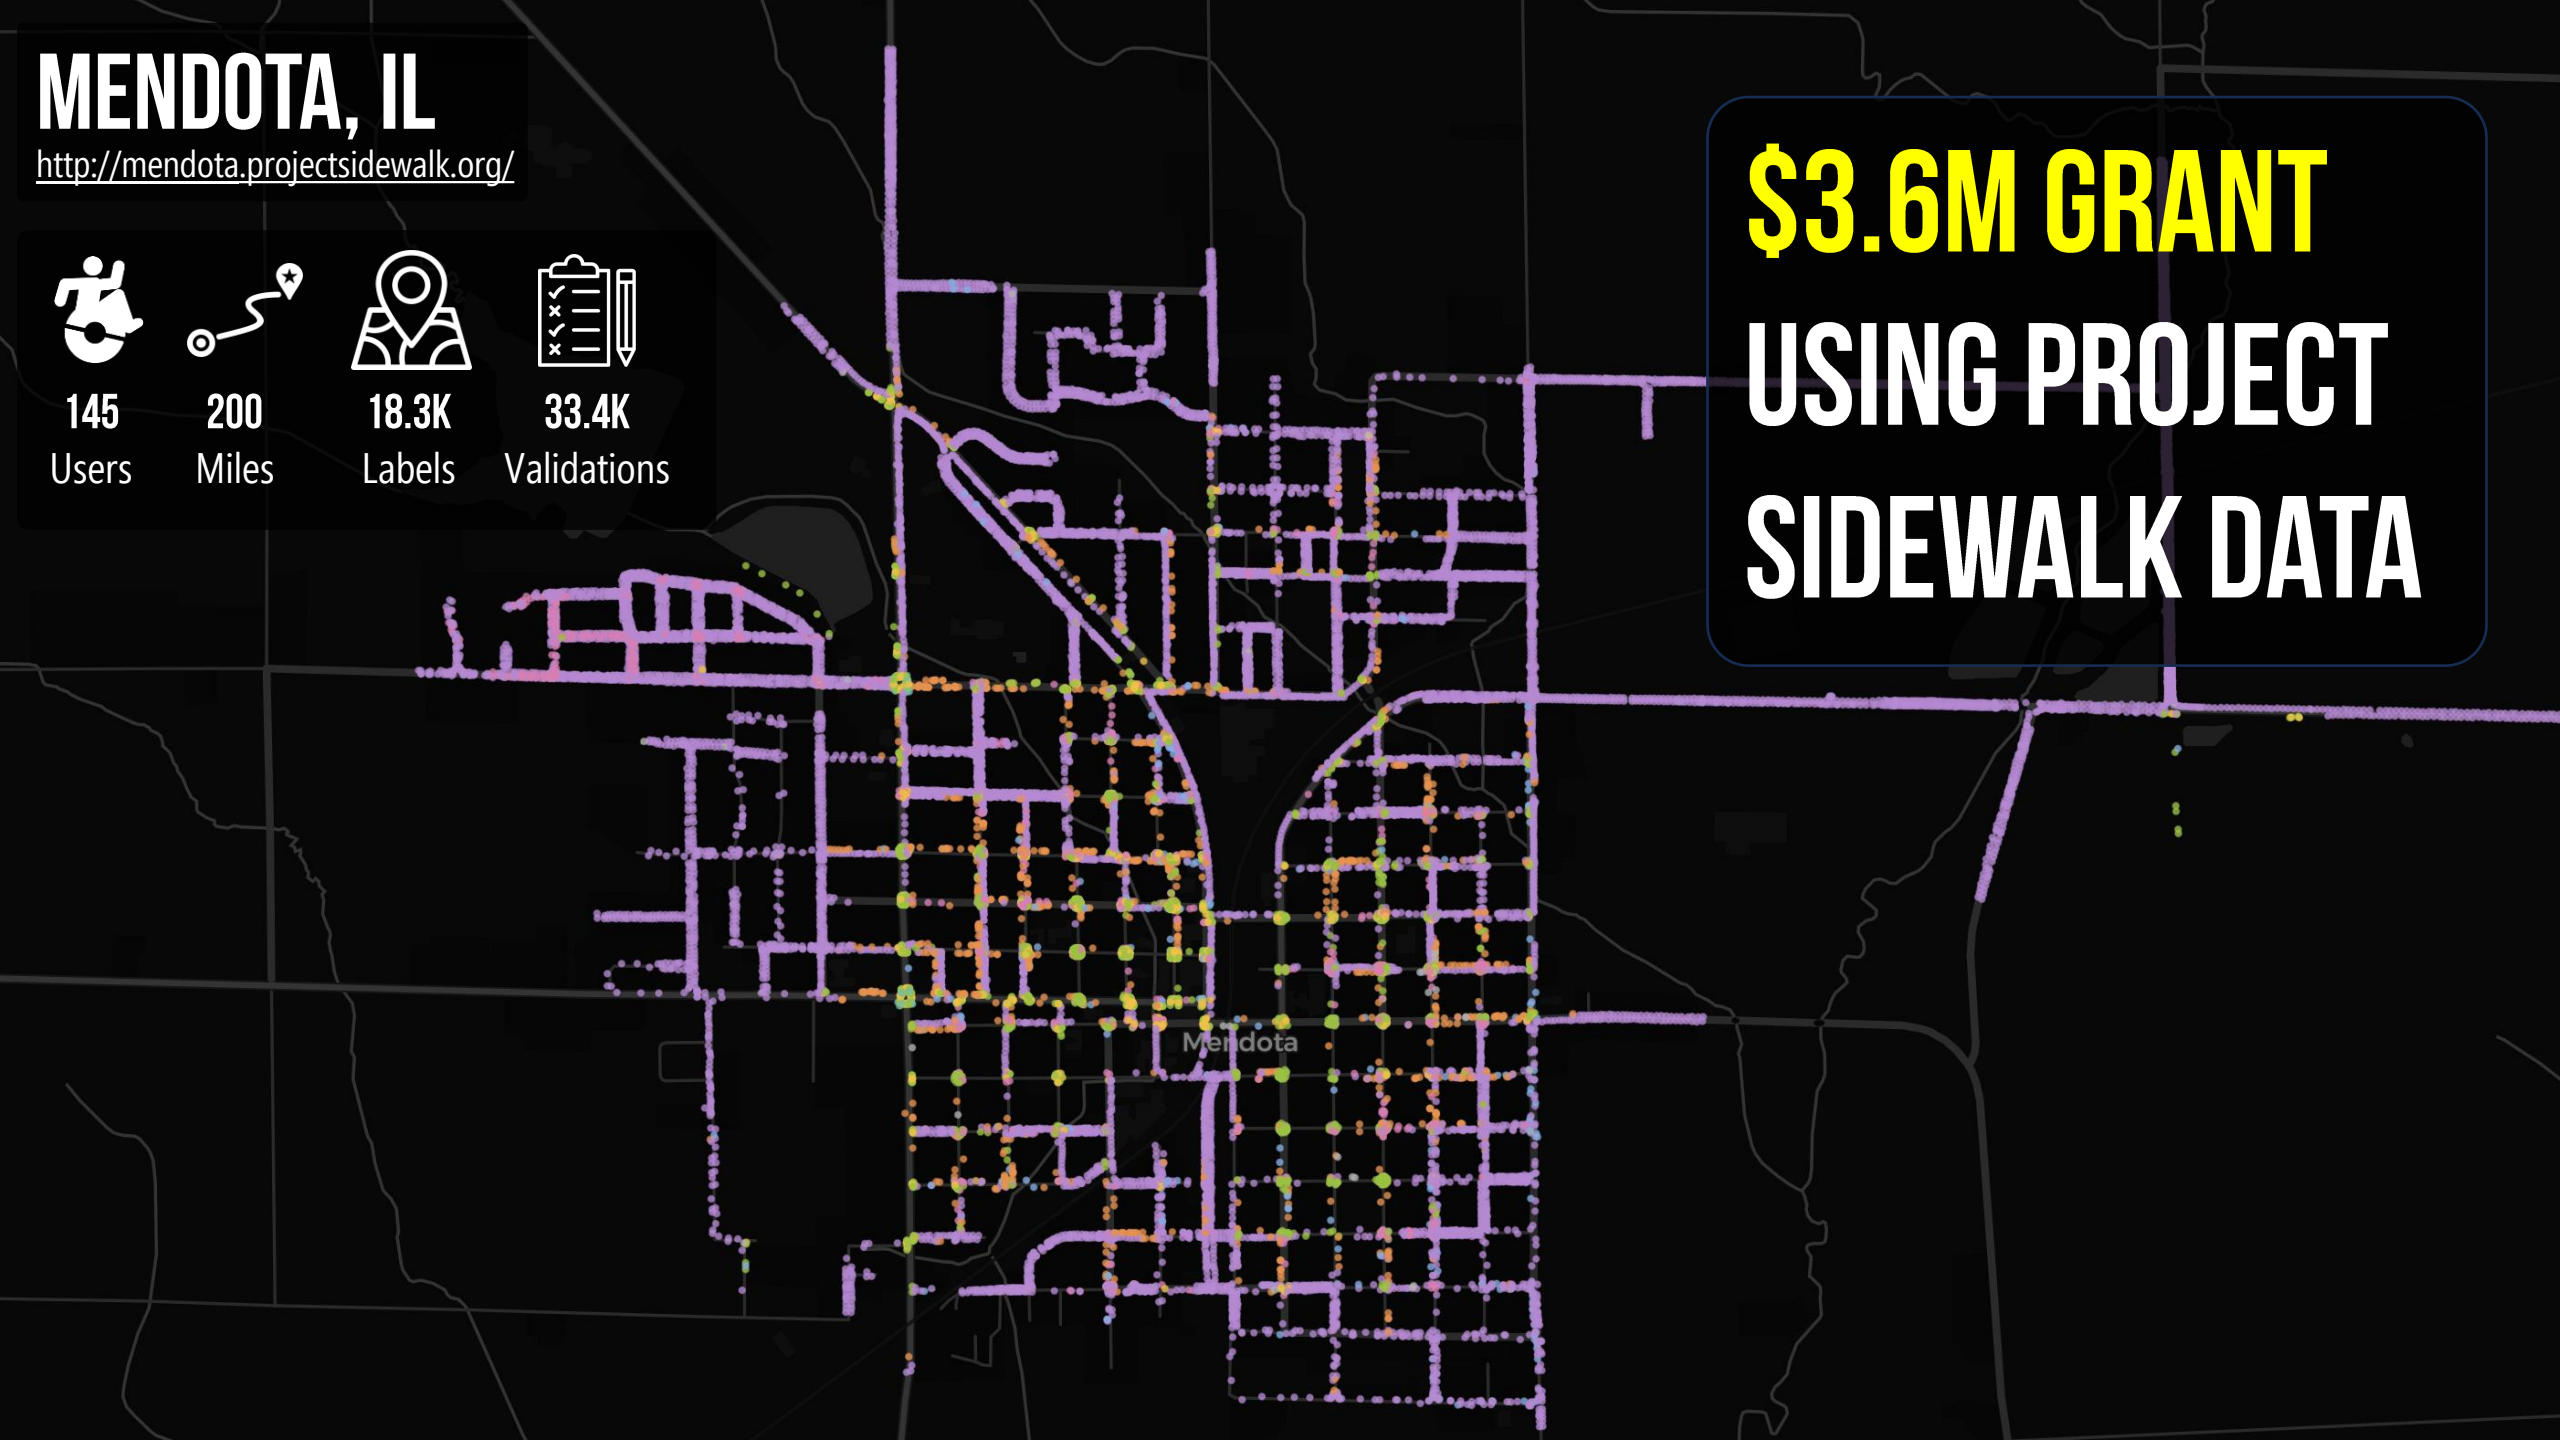
- ALISON MURPHY

Director of Communications & Development
40th Ward of Chicago

MENDOTA, IL

http://mendota.projectsidewalk.org/

145	200	18.3K	33.4K
Users	Miles	Labels	Validations



$3.6M GRANT USING PROJECT SIDEWALK DATA

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

Service Learning

Research Engine

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

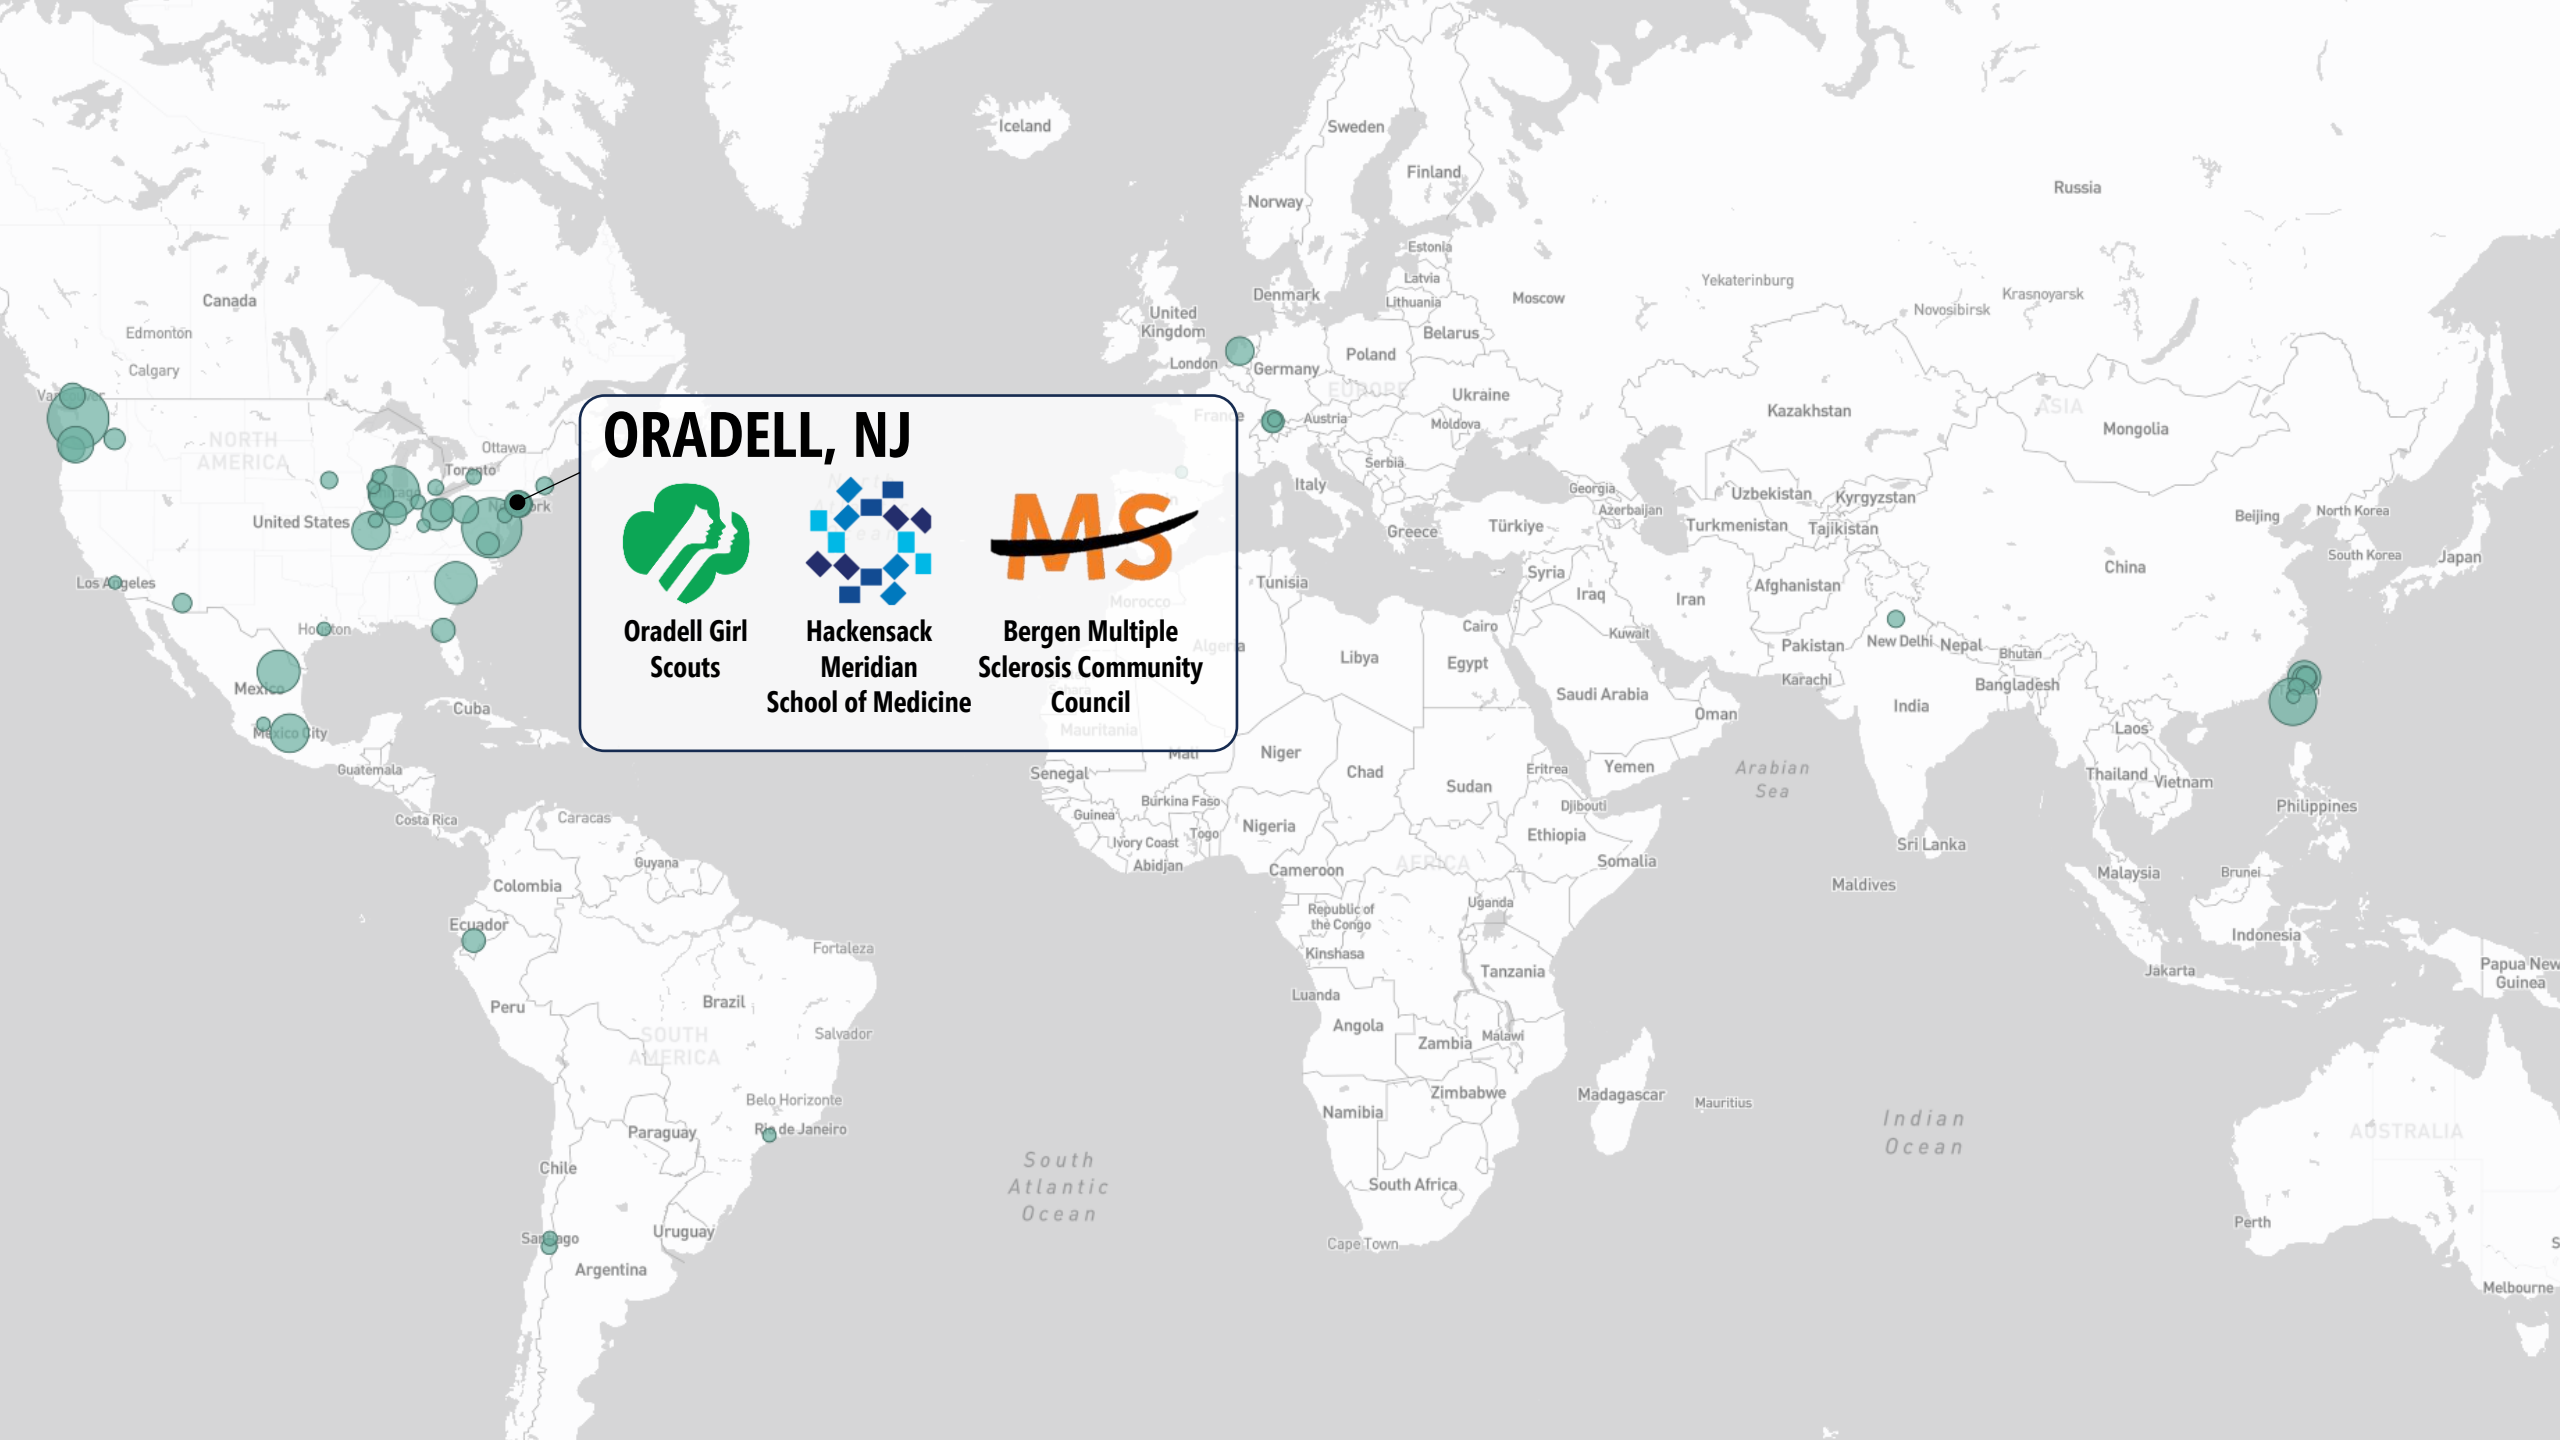
Informing Policy

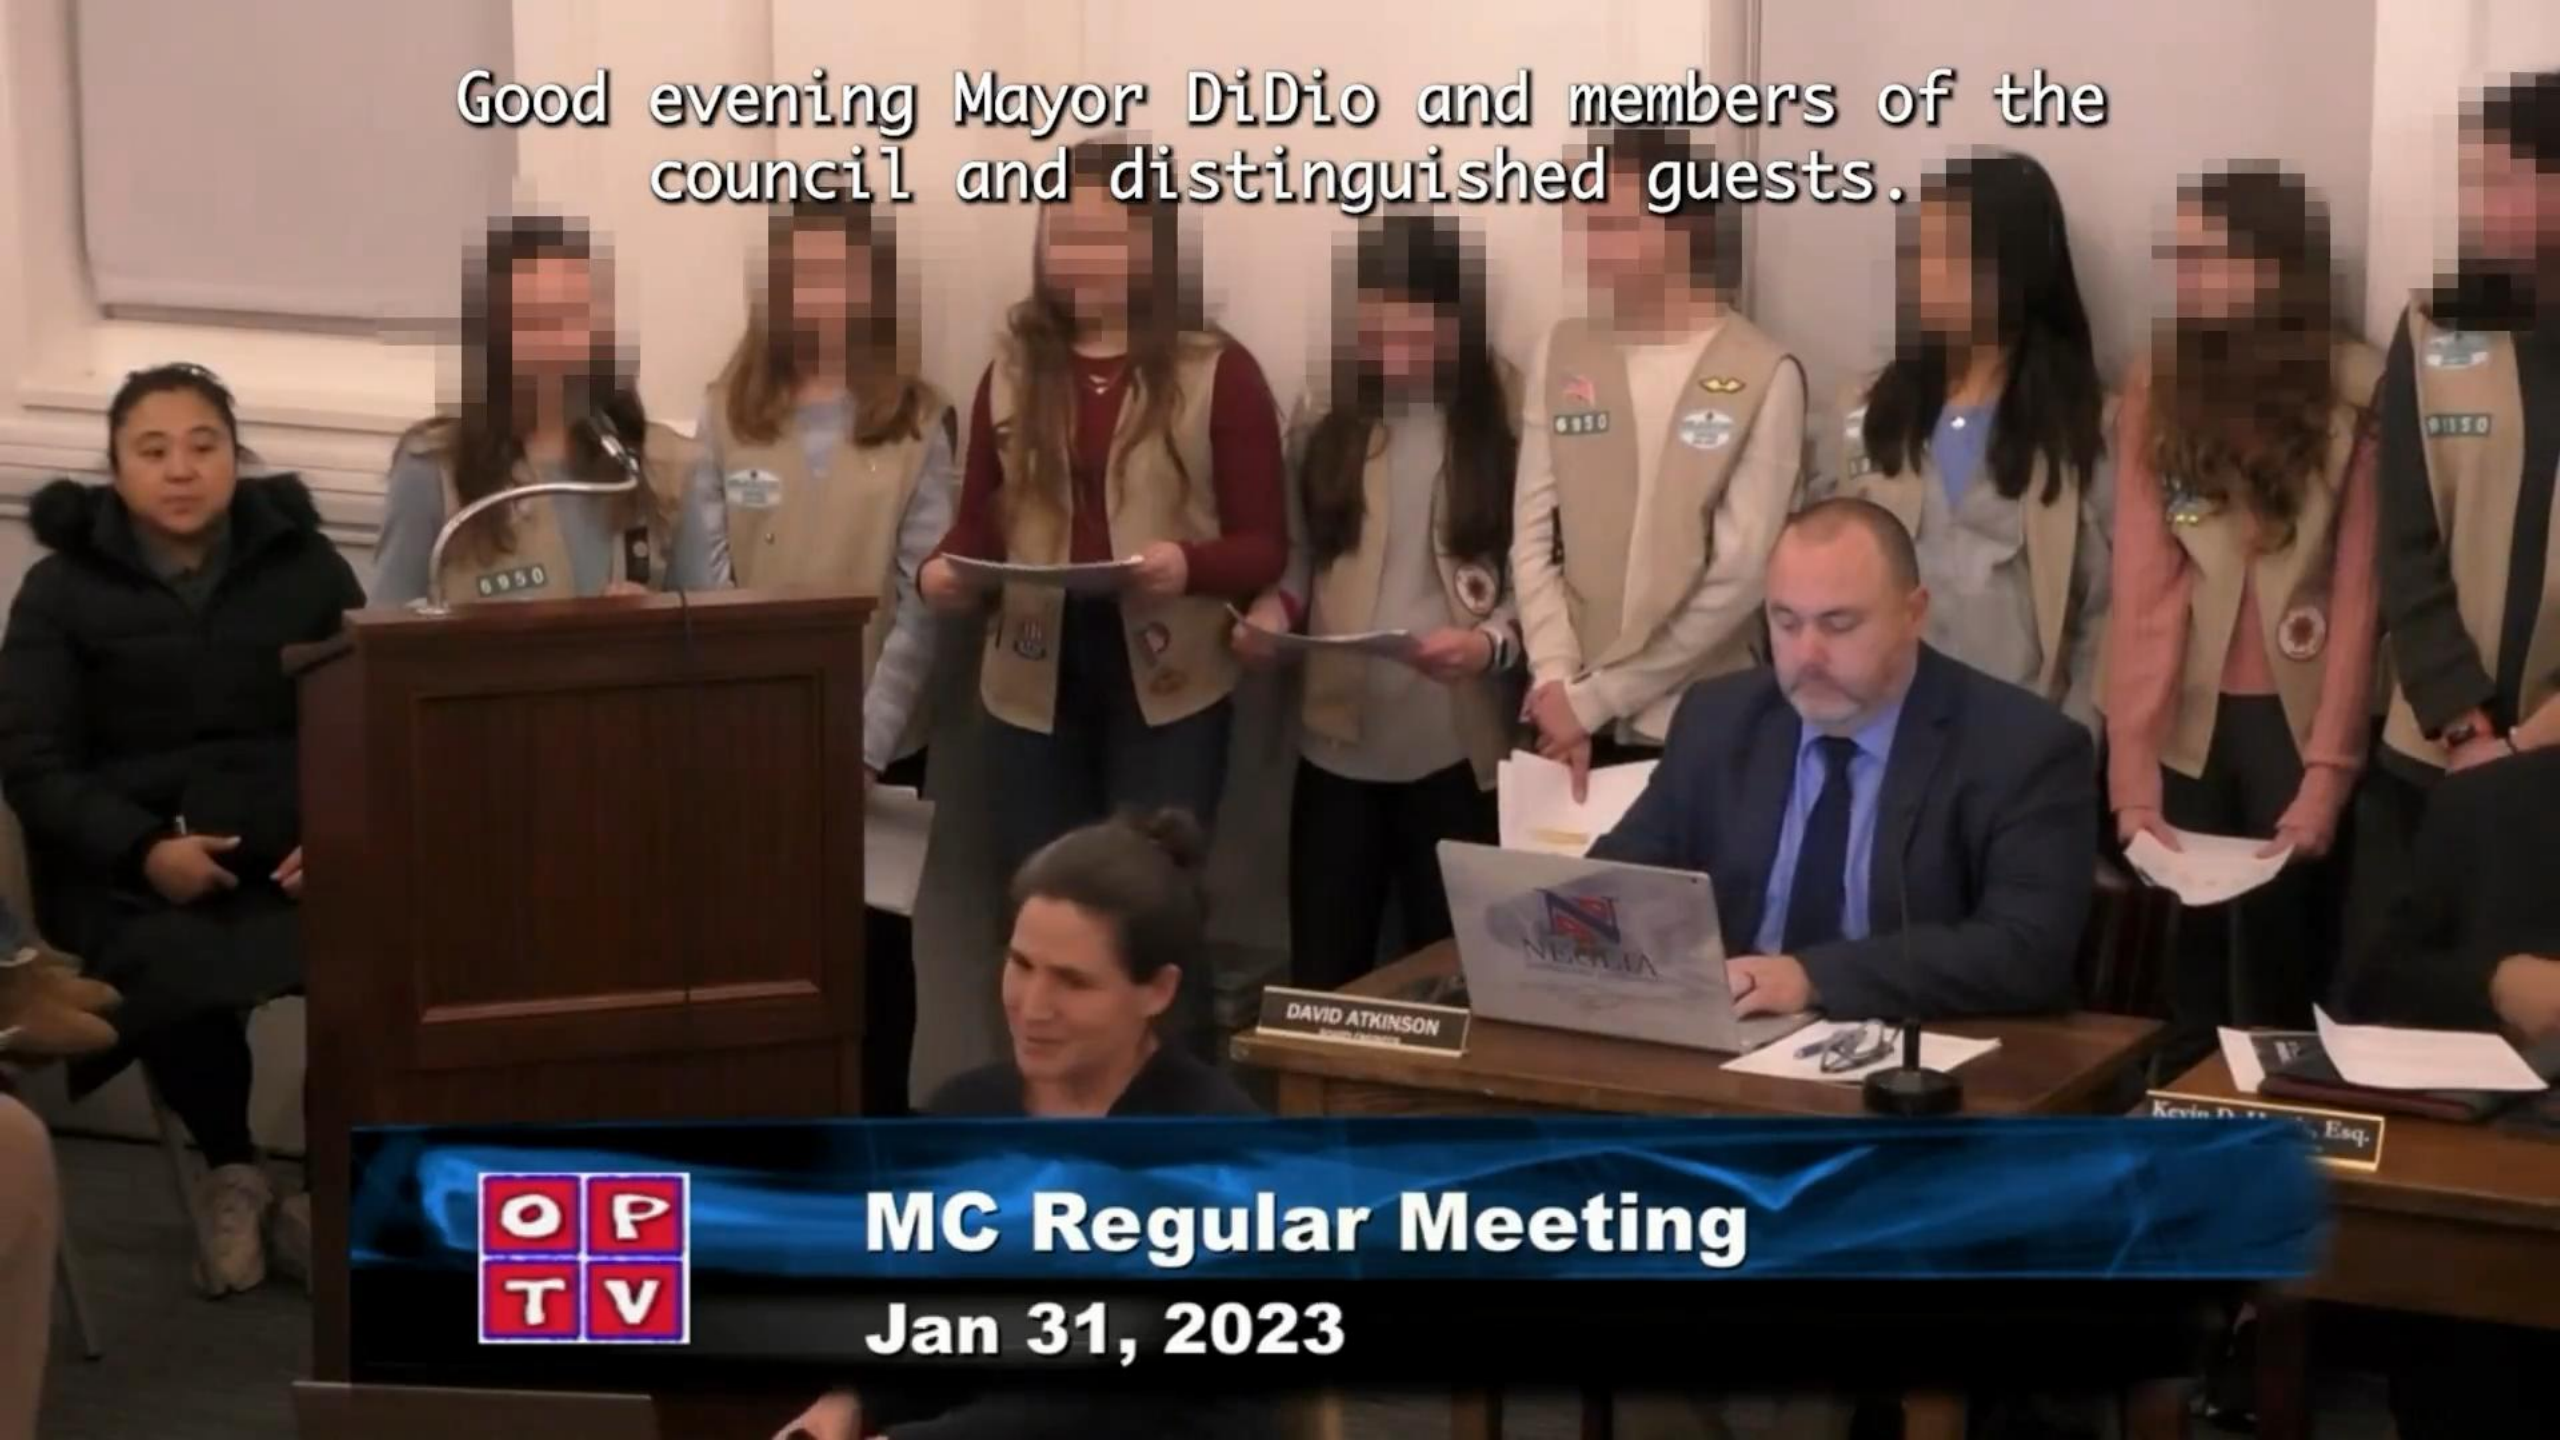
Urban Planning

Service Learning

Research Engine

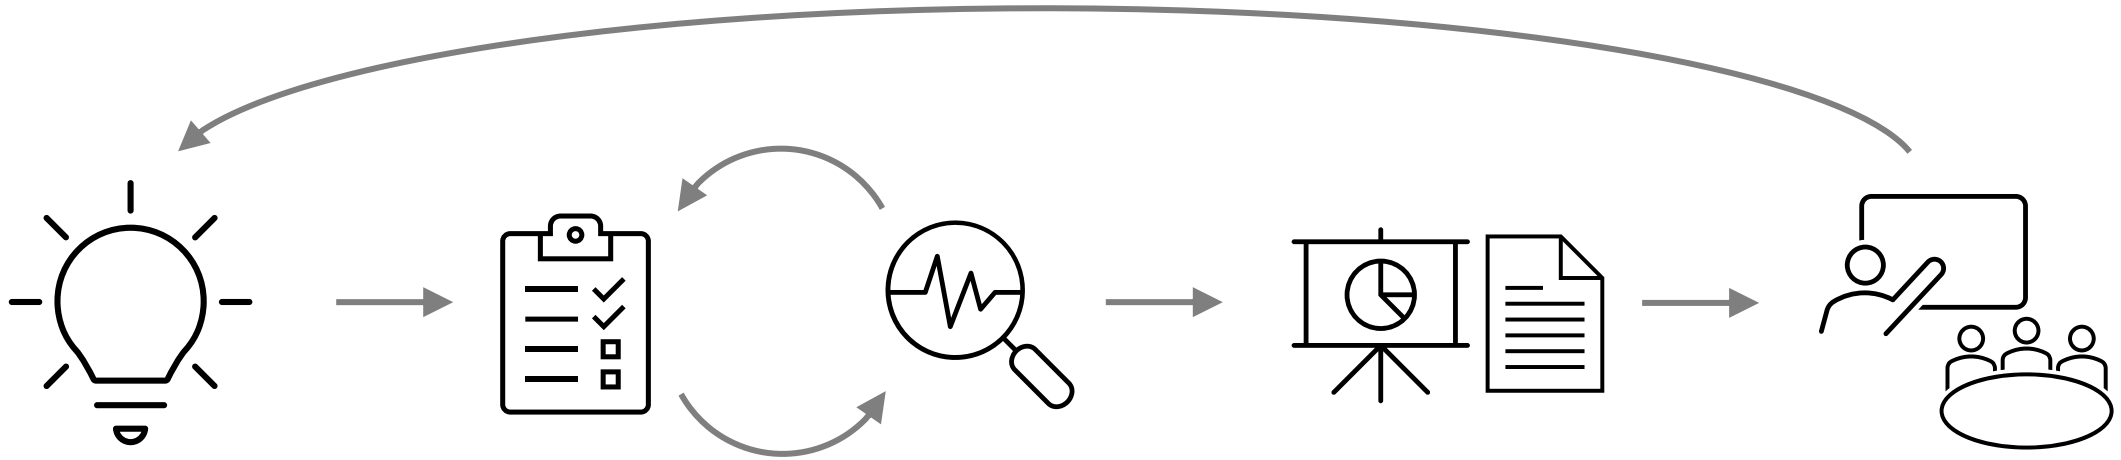

HOW CAN PROJECT SIDEWALK BE A PLATFORM FOR LEARNING WHILE ALSO PROVIDING VALUABLE DATA?


ORADELL, NJ
Oradell Girl Scouts
Hackensack Meridian School of Medicine
Bergen Multiple Sclerosis Community Council
Iceland
Sweden
Finland
Norway
Russia
Canada
Edmonton
Calgary
Ottawa
Toronto
United States
Los Angeles
Houston
Mexico
Mexico City
Cuba
Guatemala
Costa Rica
Caracas
Colombia
Guyana
Ecuador
Peru
Brazil
Fortaleza
Salvador
Belo Horizonte
Rio de Janeiro
Paraguay
Chile
Uruguay
Argentina
Santiago
South Atlantic Ocean
United Kingdom
London
Denmark
Estonia
Latvia
Lithuania
Moscow
Yekaterinburg
Novosibirsk
Krasnoyarsk
Belarus
Poland
Germany
France
Austria
Ukraine
Moldova
Serbia
Italy
Greece
Türkiye
Georgia
Azerbaijan
Kazakhstan
Uzbekistan
Turkmenistan
Kyrgyzstan
Tajikistan
Mongolia
Beijing
North Korea
South Korea
Japan
China
Syria
Iraq
Iran
Afghanistan
Pakistan
New Delhi
Nepal
Bhutan
Bangladesh
India
Karachi
Laos
Thailand
Vietnam
Philippines
Sri Lanka
Maldives
Malaysia
Brunei
Indonesia
Jakarta
Papua New Guinea
Arabian Sea
Kuwait
Oman
Saudi Arabia
Yemen
Eritrea
Djibouti
Ethiopia
Somalia
Tunisia
Libya
Egypt
Cairo
Morocco
Algeria
Mauritania
Mali
Niger
Chad
Sudan
Senegal
Guinea
Burkina Faso
Ivory Coast
Abidjan
Togo
Nigeria
Cameroon
Republic of the Congo
Kinshasa
Uganda
Tanzania
Luanda
Angola
Zambia
Malawi
Zimbabwe
Namibia
South Africa
Cape Town
Madagascar
Mauritius
Indian Ocean
AUSTRALIA
Perth
Melbourne


Good evening Mayor DiDio and members of the council and distinguished guests.
DAVID ATKINSON
OPTV
MC Regular Meeting
Jan 31, 2023

PROJECT SIDEWALK

SERVICE LEARNING



Li, C., *et al.* "I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk. CHI 2024

CHI'24

Li, Ma, Saugstad *et al.*
"I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk



"I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk

Chu Li*
University of Washington, USA
chuchuli@cs.washington.edu

Katrina Ma*
University of Washington, USA
ma.katrina91@gmail.com

Michael Saugstad
University of Washington, USA
saugstad@cs.washington.edu

Kie Fujii
Hackensack Meridian School of Medicine, USA
kie.fujii@hmhn.org

Molly Delaney
University of Illinois at Chicago, USA
mdelan9@uic.edu

Yochai Eisenberg
University of Illinois at Chicago, USA
yeisen2@uic.edu

Delphine Labbé
University of Illinois at Chicago, USA
dlabbe@uic.edu

Judy L. Shanley
Easterseals, USA
jshanley@easterseals.com

Devon Snyder
University of Illinois at Chicago, USA
dsnyde8@uic.edu

Florian P. Thomas
Hackensack Meridian School of Medicine, USA
florian.thomas@hmsom.edu

Jon E. Froehlich
University of Washington, USA
jonf@cs.washington.edu



Figure 1: A community-driven digital civics project using Project Sidewalk, an open-sourced accessibility tool, to conduct sidewalk accessibility assessments in the town of Oradell, NJ. Images show community members mapping, analyzing, and presenting the collected data to the City Council over the course of the project.

ABSTRACT

Despite decades of effort, pedestrian infrastructure in cities continues to be unsafe or inaccessible to people with disabilities. In this paper, we examine the potential of community-driven digital civics to assess sidewalk accessibility through a deployment study of an open-source crowdsourcing tool called Project Sidewalk. We explore Project Sidewalk's potential as a platform for civic learning and service. Specifically, we assess its effectiveness as a tool for community members to learn about human mobility, urban planning, and accessibility advocacy. Our findings demonstrate that community-driven digital civics can support accessibility advocacy and education, raise community awareness, and drive pro-social behavioral change. We also outline key considerations for deploying digital civic tools in future community-led accessibility initiatives.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05.
https://doi.org/10.1145/3613904.3642003

CCS CONCEPTS

• **Human-centered computing** → **Accessibility systems and tools**; **Interactive systems and tools**; • **Information systems** → **Crowdsourcing**.

KEYWORDS

accessibility, digital civics, community science, service learning, crowdsourcing



Figure 5: Timeline showing dates of partnership formations and key project events in the Oradell Sidewalk Project.

Hackensack Meridian School of Medicine), and a local *Girl Scouts* troop.[2]

3.2.2 In-person Field Audits. In September 2021, the advocacy group and the Scouts held several meetings to generate ideas on how to improve Oradell's urban accessibility for people with disabilities. They first discussed topics such as what constitutes an *"accessibility issue"* and general impressions of Oradell's accessibility levels across neighborhoods, and then decided on conducting in-person field audits to document accessibility barriers. From October to November 2021, advocacy group representatives and the Scouts took photographs of accessibility issues around their neighborhood, including barriers such as missing curb ramps, uneven surfaces, walkway obstructions (*e.g.*, utility poles, overgrown greenery), and obstructed entrances to municipal buildings (*e.g.*, library, post office, town hall). They identified ~30 issues with a location, picture, and description across ~1 miles of sidewalks (Figure 7).

In early November, an advocacy group member (P3 in Table 1) discovered Project Sidewalk online and initiated a meeting to discuss adopting this new approach for identifying accessibility issues. With the Project Sidewalk team's support, they decided to deploy and use this web tool in Oradell (https://oradell.projectsidewalk.org/).

3.2.3 Assessing Oradell with Project Sidewalk. After obtaining geographic neighborhood boundary data, Project Sidewalk was set up and deployed in Oradell in April 2022. Initial tests and contributions were from the advocacy group and Project Sidewalk team. To broaden community involvement and provide additional training, the organizers planned and held two mapathons and began advertising campaigns within their communities. Figure 9 shows a temporal plot of label and validation contributions over time.

First mapathon. The first mapathon was on April 22, 2022 at the administration building of the local government (Oradall Borough Hall). The event was organized by the local Scout troop leaders, representatives from the National Multiple Sclerosis Society, and students and faculty from Hackensack Medical. The mapathon's goal was to familiarize participants with Project Sidewalk, motivate contributions, and advance understanding of how the built environment can affect people with disabilities.



Figure 6: (A) Diagram showing the deeply intertwined interactions between stakeholders in the Oradell Sidewalk Project, including the local medical center (Hackensack Meridian), MS (Multiple Sclerosis) group, and Girl Scouts troop who started the initiative, then engaged with Project Sidewalk and their research team to assess all sidewalks in Oradell, NJ. (B) Image from the initial City Council meeting in March 2022. (C) Image from the first Scouts meeting with local advocates in September 2021.

[2]Girl Scouts is a youth organization for girls in the United States with a mission to "*build girls of courage, confidence, and character, who make the world a better place*" through activities such as camping, community service, and practical skill-building like first aid [87].



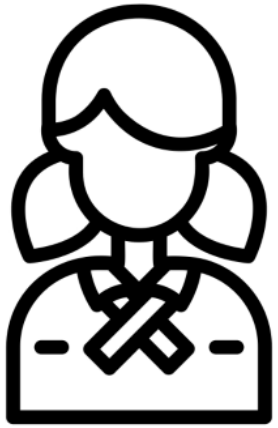
Figure 7: Examples of sidewalk accessibility issues photographed by advocacy group representatives and Scout members during in-person audits in their neighborhoods.

SERVICE LEARNING

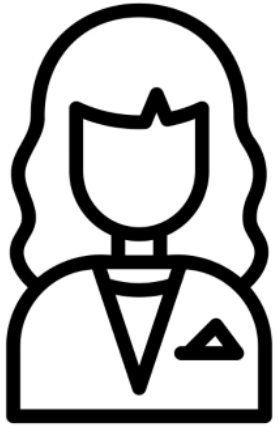
POST-HOC INTERVIEWS IN ORADELL (N=19)



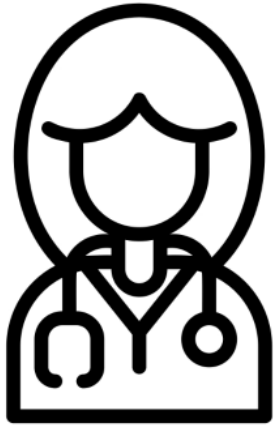
Disability Advocates



Girl Scouts



Parents



Medical Professionals

Li, C., *et al.* "I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk. CHI 2024

Before Project Sidewalk**, I never realized how important curb ramps** are for people with disabilities…

- GIRL SCOUT (GS4)

Post-study interview



Li, C., *et al.* "I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk. CHI 2024

Now when they walk down the street and see something [inaccessible], they quickly notice and say **"This is wrong!"**

- PARENT

Post-study interview

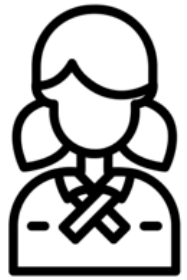


Li, C., *et al.* "I never realized sidewalks were a big deal": A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk. CHI 2024

“This experience gave me a new perspective that I can use **to help change the world**.”

- GIRL SCOUT (GS7)

Post-study interview



Li, C., *et al.* “I never realized sidewalks were a big deal”: A Case Study of a Community-Driven Sidewalk Audit Using Project Sidewalk. CHI 2024

McDevitt civics class earns state recognition for sidewalk accessibility project

BY LEA ZAHARONI • EDUCATION • JUNE 28, 2026

Four students from McDevitt Middle School have received recognition from the Massachusetts Department of Elementary and Secondary Education for their work on a student-led civics project addressing sidewalk accessibility in the city.



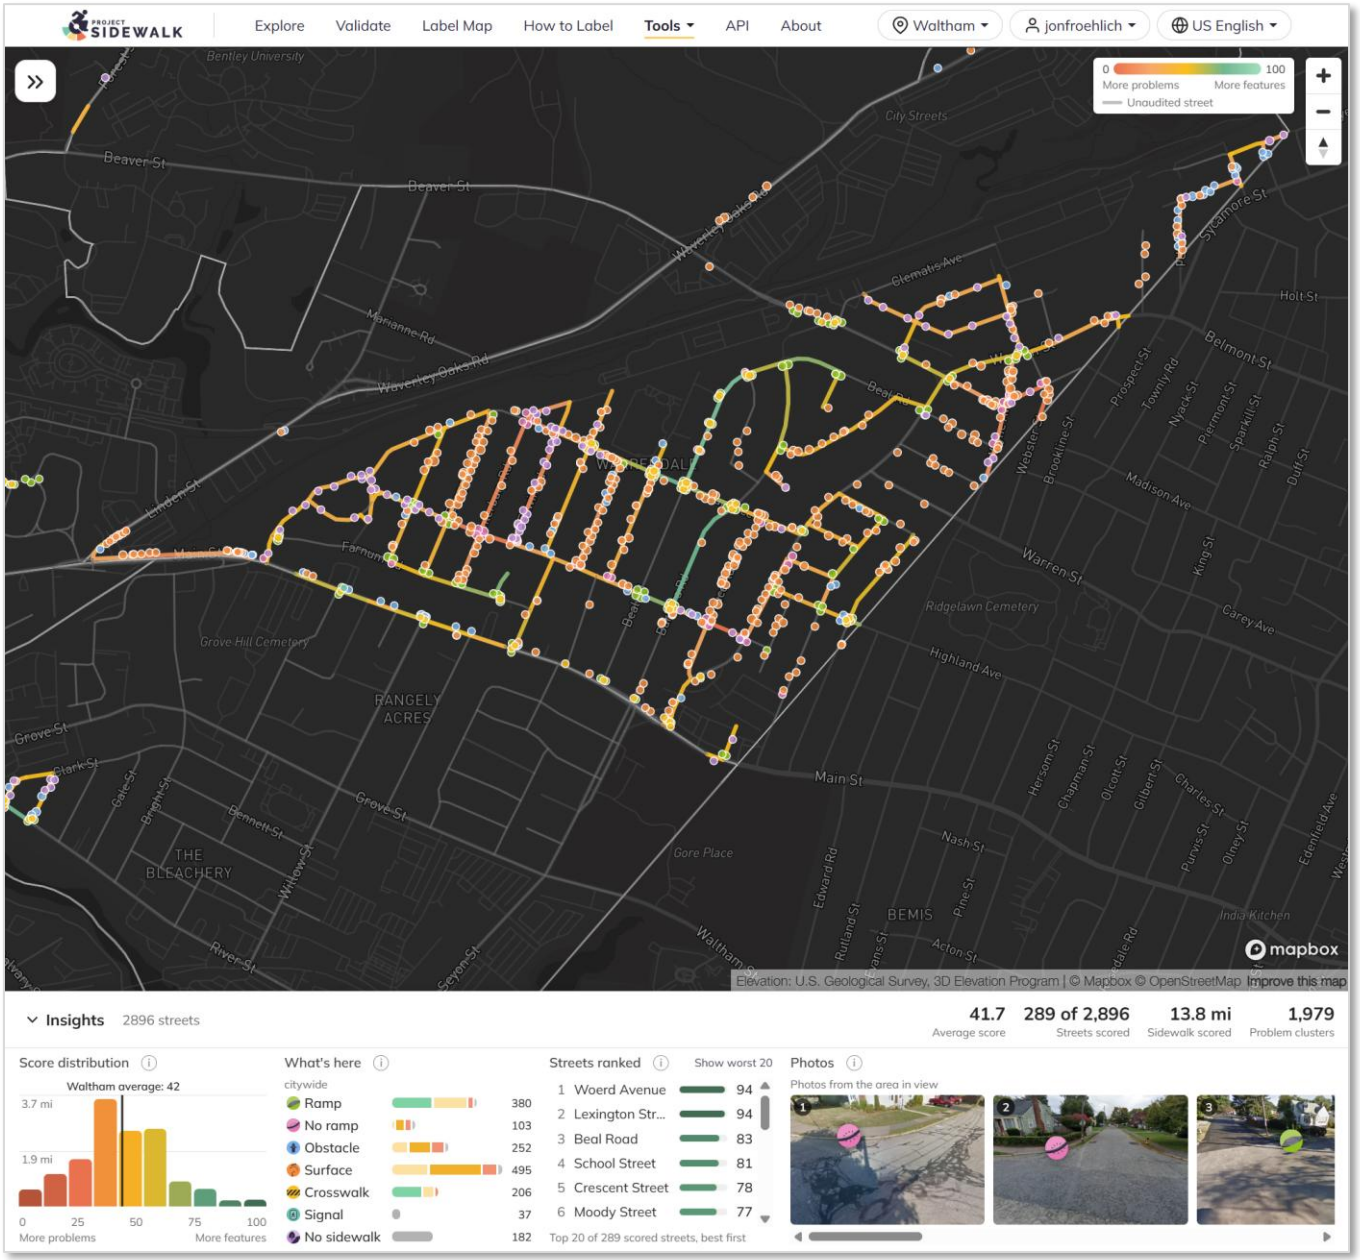
Students at the state civics showcase, from left to right: Isaac Hayward, Jacob Belding, Dylan Johanson, and Cameron Devine. Courtesy of Ryan Hanlon.

Jacob Belding, Isaac Hayward, Dylan Johanson and Cameron Devine, all members of Ryan Hanlon's eighth grade civics class, traveled to the Edward M. Kennedy Institute for the U.S. Senate in Boston to present their project at a statewide civics competition. They were chosen by their classmates to represent McDevitt Room 324.

Their project focuses on damaged and missing sidewalks across Waltham, which the class identified as a barrier to pedestrian safety.

The class began the project after spending a week examining problems in their school community and city. With this winter's heavy blizzards fresh in their minds, concerns about icy sidewalks and unsafe bicycling routes came up often.

Belding discovered an online tool called Project Sidewalk, developed by the Makeability Lab at the University of Washington's Paul G. Allen School of Computer Science. The platform uses Google Earth's street view to allow volunteers anywhere in the world to virtually walk city streets and label accessibility problems, including broken pavement, missing curb ramps and other obstacles.



https://walthamtimes.org/2026/06/28/mcdevitt-civics-class-earns-state-recognition-for-sidewalk-accessibility-project/

PROJECT SIDEWALK

SERVICE LEARNING



Professor Yochai Eisenberg

Disability and Human Development

Rehabilitation Sciences

UNIVERSITY OF ILLINOIS CHICAGO

Project Sidewalk Service-Learning Curriculum

Social Action and Media Curriculum

Inclusive Service-Learning Curriculum

Scan to learn more or email kmcke6@uic.edu

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

Service Learning

Research Engine

PROJECT SIDEWALK

TRANSLATIONAL IMPACT



PROJECTSIDEWALK.ORG

Informing Policy

Urban Planning

Service Learning

Research Engine

PROJECT SIDEWALK

TRANSLATIONAL IMPACT

Universidad Nacional Autónoma de México

PLAKSHA UNIVERSITY

UNIVERSITY OF SOUTH CAROLINA

University of Zurich[UZH]

SFU SIMON FRASER UNIVERSITY

USP Universidade de São Paulo

SMU

UIC UNIVERSITY OF ILLINOIS CHICAGO

WPI

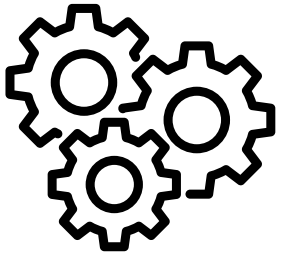
UNIVERSITEIT VAN AMSTERDAM

UF UNIVERSITY of FLORIDA College of Health and Human Performance

Informing Policy

Urban Planning

Service Learning



Research Engine

ZURICH, SWITZERLAND

Disability Community

Local Government

University of Zurich



UNIVERSITY OF ZURICH

URBAN SCIENCE RESEARCH WITH PROJECT SIDEWALK

Navigation Challenges in Urban Areas for Persons with Mobility Restrictions

Hoda Allahbakhshi[1]
Digital Society Initiative, University of Zurich, Switzerland
Department of Geography, University of Zurich, Switzerland
University Research Priority Program "Dynamics of Healthy Aging", University of Zurich, Switzerland

Annina Ardüser
Department of Geography, University of Zurich, Switzerland

Abstract

The Sustainable Development Goals (SDGs) promote making the world better for everyone, with a focus on creating cities that are inclusive and sustainable, as outlined in SDG 11. Spatial accessibility plays a pivotal role in fostering age-friendly and inclusive urban environments. However, there is still a lack of complete data on accessibility essential for providing mobility services to individuals with restricted mobility, mainly due to the high costs. While some participatory initiatives like OpenStreetMap (OSM) have made progress in this area, there is still a significant gap in data about sidewalk accessibility.

To address this gap, we used a citizen science approach to gather information and improve our understanding of sidewalk accessibility in District 1 of Zurich. Eighteen individuals from diverse population groups took part in our study. Using the Project Sidewalk web tool (PRSW), participants collected sidewalk features like curb ramps and surface problems by virtually inspecting street view images.

In this paper, we present preliminary results derived from participatory data collection. The findings show the variances in accessibility labels concerning their frequency, spatial distribution, and severity levels attributed by participants. Furthermore, we provide insights into the accuracy of the data, verified through validation by experts in geographical knowledge using PRSW.

Our approach allowed for broader participation and diverse perspectives in collecting sidewalk accessibility data. We believe that the provided dataset has the potential to address unanswered questions about spatial accessibility. For instance, the distribution of accessibility within specific population groups or across a city can be explored. This information can help policymakers develop interventions that tackle accessibility inequalities and ensure equitable access, especially for those with mobility impairments.

2012 ACM Subject Classification Social and professional topics

Keywords and phrases Navigation, Mobility-restrictions, Inclusive mobility, Spatial accessibility, Citizen science

Digital Object Identifier 10.4230/LIPIcs.COSIT.2024.22

Category Short Paper

Acknowledgements We thank the University of Zurich, the Digital Society Initiative, and Smart City Zurich for partially financing this research.

[1] corresponding author

© Hoda Allahbakhshi and Annina Ardüser;
licensed under Creative Commons License CC-BY 4.0
16th International Conference on Spatial Information Theory (COSIT 2024).
Editors: Benjamin Adams, Amy Griffin, Simon Scheider, and Grant McKenzie; Article No. 22; pp. 22:1–22:8
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Towards an Inclusive Urban Environment: A Participatory Approach for Collecting Spatial Accessibility Data in Zurich

Hoda Allahbakhshi
Digital Society Initiative, University of Zürich, Switzerland
Department of Geography, University of Zürich, Switzerland

Abstract

The unprecedented rate of urbanization, along with the increase in the aging and disabled populations, bring about an increasing demand for public services and an inclusive urban environment that allows easy access to those facilities. Spatial Accessibility is a measure to assess how inclusive a city is and how easily public facilities can be reached from a specific location through movement in physical space or built environment.

A detailed geodata source of accessibility features is needed for reliable spatial accessibility assessment, such as sidewalk width, surface type, and incline. However, such data are not readily available due to the huge implication costs. Remote crowdsourcing data collection using Street View Imagery, so-called 'virtual audits' have been introduced as a valid, cost-efficient tool for accessibility data enrichment at scales compared to conventional methods because it enables involving more participants, saving more time by avoiding field visits and covering a larger area.

Therefore, in our pilot project, ZuriACT: Zurich Accessible CiTy, with the help of digital tools that allow for virtual inspections and measurements of accessibility features, we want to contribute to collecting and enriching accessibility information in the city of Zurich embedded in a citizen science project that will have both scientific and social impacts.

With the help of additional accessibility data produced in this project, the issues of an inclusive urban environment can be demonstrated by mapping the potential spatial inequalities in access to public facilities for disabled or restricted people in terms of mobility. Thus, this project provides helpful insight into implementing policy interventions for overcoming accessibility biases to ensure equitable services, particularly for people with disabilities, and contributes to creating an inclusive and sustainable urban environment. It goes without saying that an inclusive city is beneficial and impacts the quality of life of not only the population groups mentioned above but also the society at large.

2012 ACM Subject Classification Social and professional topics

Keywords and phrases Spatial accessibility, virtual audits, digital tools, mobility disability, citizen science, inclusive city, Zurich

Digital Object Identifier 10.4230/LIPIcs.GIScience.2023.13

Category Short Paper

Funding I thank the University of Zurich, the Digital Society Initiative, and Smart City Zurich for partially financing this research.

1 Introduction

It is projected that by 2050, about 70 percent of the world's population will live in urban environments, 15 percent of them will live with disabilities [10]. Moreover, the prediction shows that by 2050, the number of older people will reach 2 billion worldwide [12]. The unprecedented rate of urbanization, along with the increase in the aging and disabled populations, bring about an increasing demand for public services and access to those facilities. Depending on the infrastructure and design, the urban environment and physical

© Hoda Allahbakhshi;
licensed under Creative Commons License CC-BY 4.0
12th International Conference on Geographic Information Science (GIScience 2023).
Editors: Roger Beecham, Jed A. Long, Dianna Smith, Qunshan Zhao, and Sarah Wise; Article No. 13; pp. 13:1–13:6
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

AGILE: GIScience Series, 6, 23, 2025. https://doi.org/10.5194/agile-giss-6-23-2025
Proceedings of the 28th AGILE Conference on Geographic Information Science, 10–13 June 2025.
Eds.: Auriol Degbelo, Serena Coetzee, Carsten Keßler, Monika Sester, Sabine Timpf, Lars Bernard
This contribution underwent peer review based on a full paper submission.
© Author(s) 2025. This work is distributed under the Creative Commons Attribution 4.0 License.

Developing enriched pedestrian networks using accessibility features

Alexandra-Ioana Georgescu[1,2,3], Hoda Allahbakhshi[1,2,3], Robert Weibel[2,3]

[1] Digital Society Initiative, University of Zurich, Zurich, Switzerland
[2] Department of Geography, University of Zurich, Zurich, Switzerland
[3] Healthy Longevity Centre, University of Zurich, Zurich, Switzerland

Correspondence: Alexandra-Ioana Georgescu (alexandra-ioana.georgescu@geo.uzh.ch)

Abstract. Accessibility studies often focus on the general population, overlooking individual differences in mobility capacities and the role of external factors. Microscale street elements such as stairs, high kerbs, and sidewalk cracks can significantly impact urban mobility for individuals with restricted movement capacities. This study introduces a workflow to integrate different detailed accessibility information, such as barriers and facilitators for pedestrian mobility, with sidewalk data to create an enriched pedestrian network. Using this network, we evaluate each segment by computing an impedance score to quantify accessibility. Furthermore, we demonstrate how the network can be tailored to individual mobility needs and highlight its potential as a decision-making tool for urban planners and civil engineers to identify and prioritise targeted interventions for vulnerable populations.

Submission Type. Analysis

BoK Concepts. [AM11] Network analysis, [AM11-7] Accessibility modelling, [AM5-3] Spatial cluster analysis

Keywords. Spatial accessibility, Microscale elements, Mobility restriction, Pedestrian networks, Urban environment

1 Introduction

Ensuring equitable and adequate access for all to public facilities in urban areas is an essential target included in various Sustainable Development Goals such as 10 and 11 (United Nations, 2022). This is especially important since by 2050 it is projected that approximately two-thirds of the world's population is expected to reside in urban areas (United Nations, 2019), and currently 16% of the global population worldwide is believed to live with some sort of disability (World Health Organisation, 2022).

Spatial accessibility refers to how easily destinations, such as public facilities, can be reached through movement in physical space (Allahbakhshi, 2023). It is highly context-dependent and influenced by various factors (Lid & Solvang, 2016), including microscale street elements such as poles, stairs, cobblestones, kerb ramps, or pedestrian signals. These elements can either restrict or facilitate pedestrian movement, significantly impacting access to certain areas and services (Hammel et al., 2015) and contributing to social exclusion (Svensson, 2010).

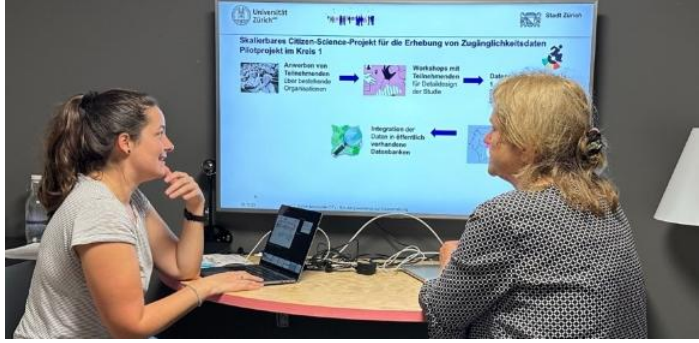
Physical barriers are especially important for individuals with mobility restrictions, as they can increase perceived distance (Vale et al., 2016), reduce the number of available opportunities (Achuthan et al., 2010) and force individuals to stay at home or in familiar places (Mao & Chen, 2022). Additionally, as individuals have different mobility needs, the impact of these elements varies (Georgescu et al., 2024). For example, a high kerb may be a barrier for a wheelchair user but can serve as a facilitator, marking the edge of the sidewalk, for a visually impaired individual. Therefore, such elements become vital in assessing spatial accessibility.

However, despite their significance, research on microscale elements and their impact on individuals with varying mobility capacities remains limited. Most of these studies rely on field visits to collect data on the various accessibility features that affect pedestrian movement. These features are then recorded in spatial databases, where each element is linked to its corresponding

UNIVERSITY OF ZURICH

SMART CITY HUB AWARD





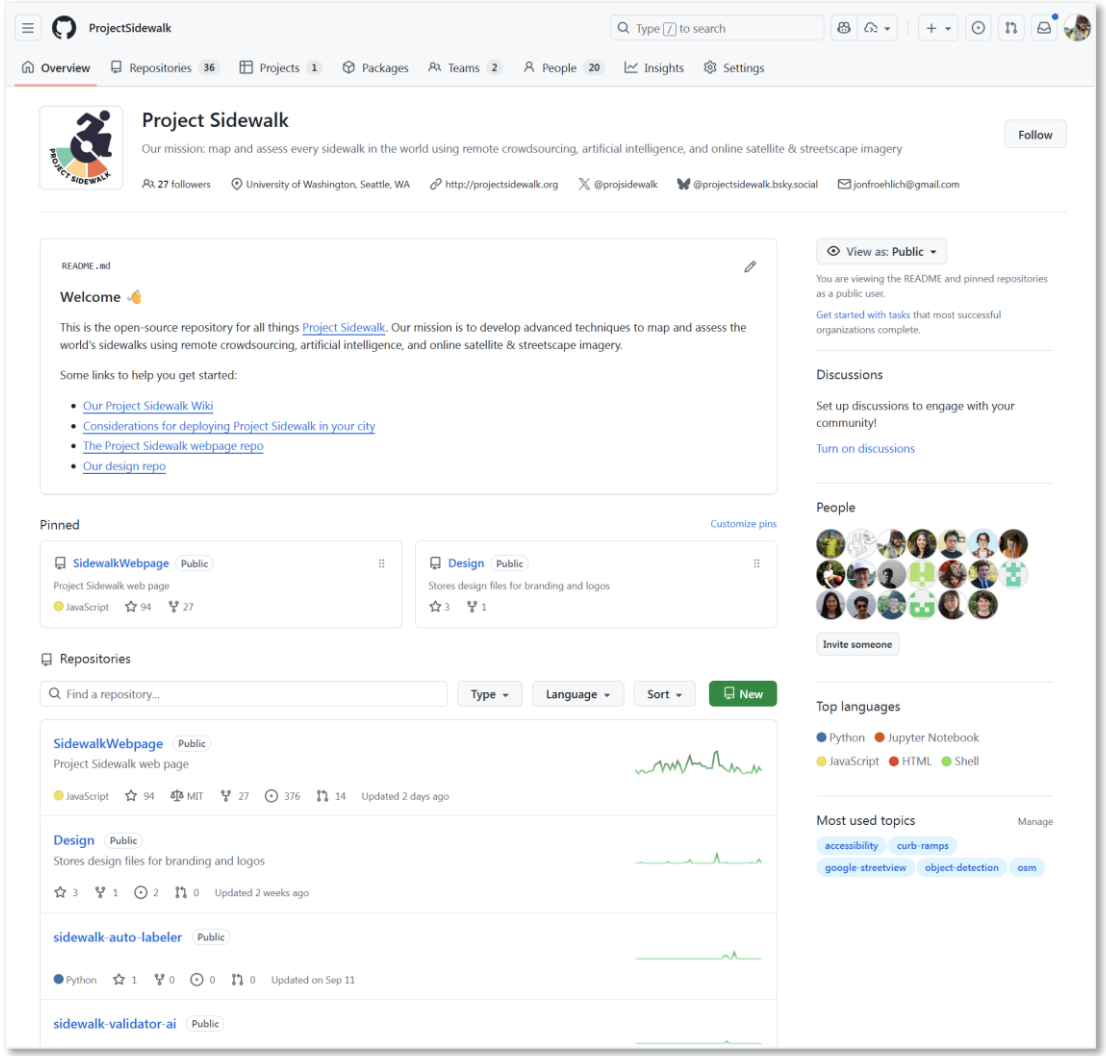




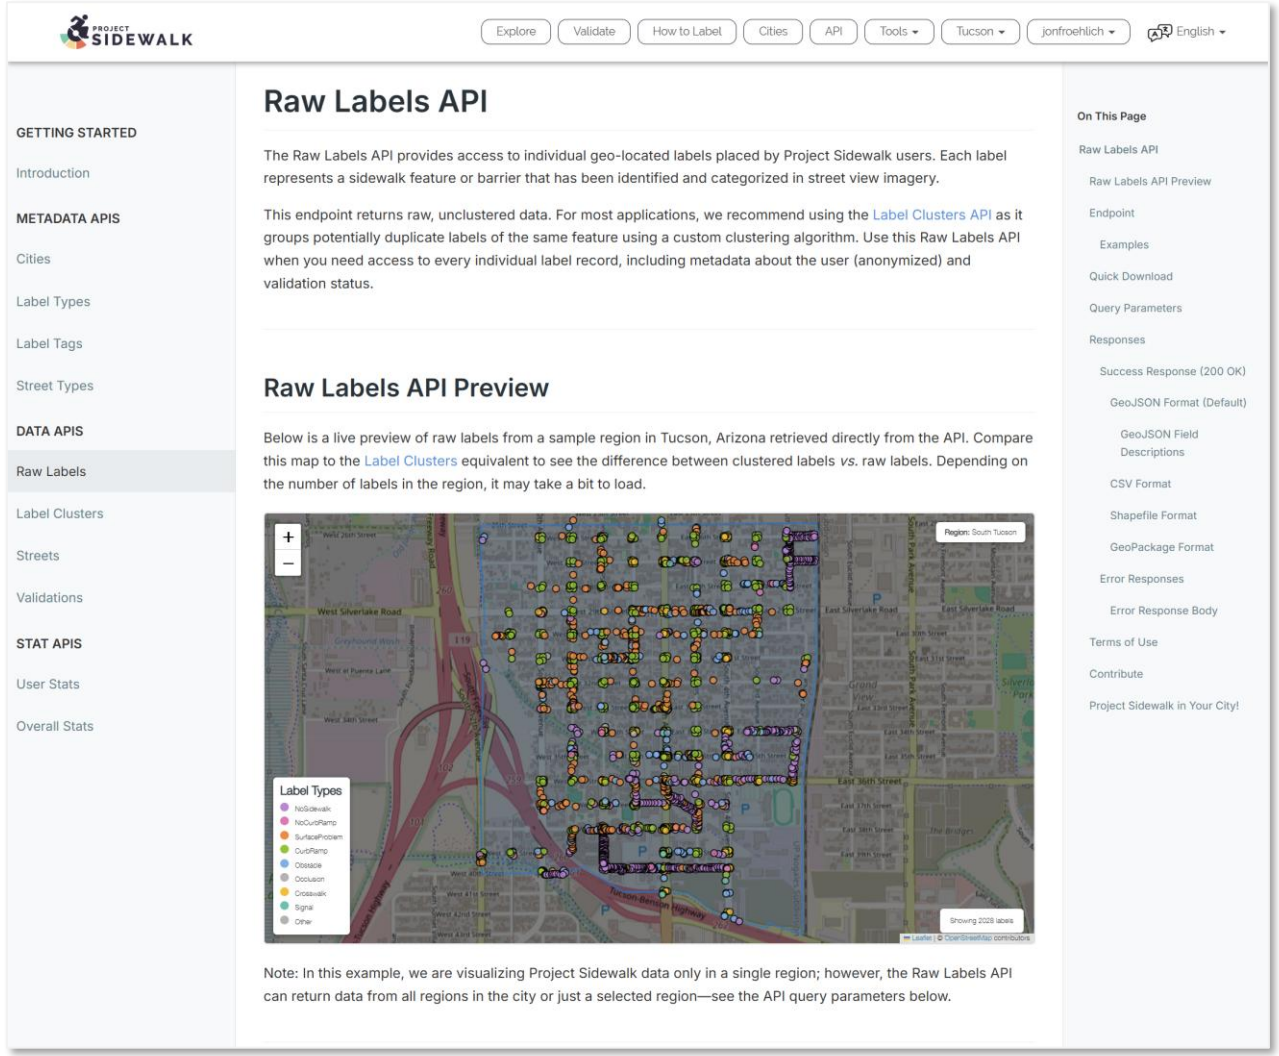


PROJECT SIDEWALK

OPEN SOURCE / DATA



https://github.com/ProjectSidewalk/

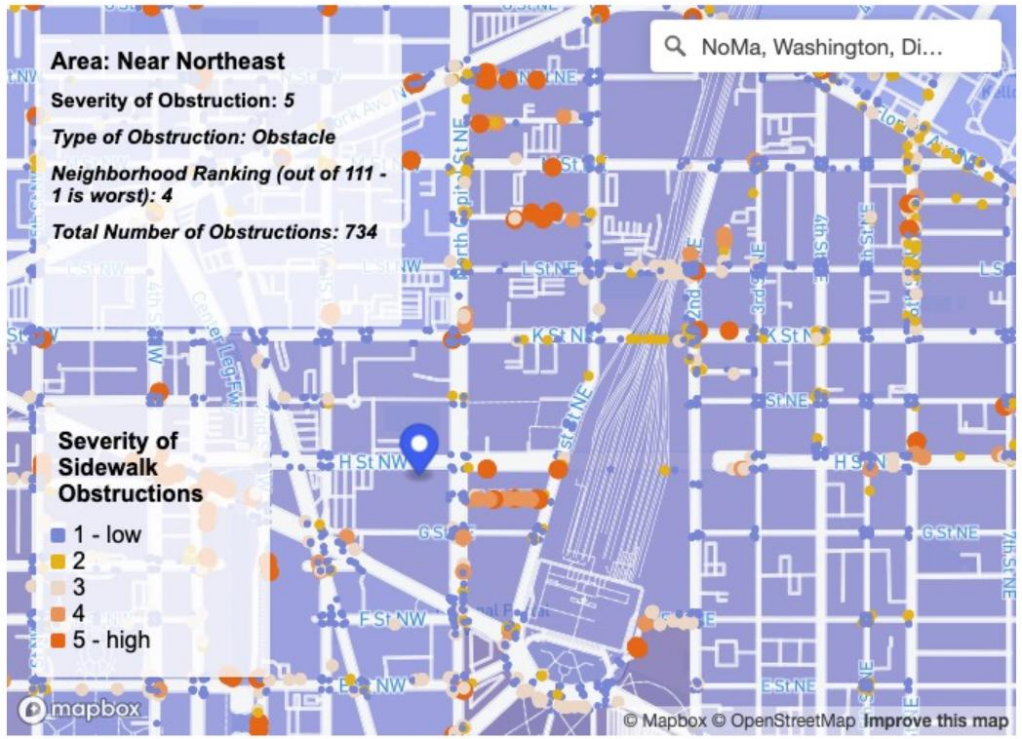


https://projectsidewalk.org/api

A city is only as accessible as its sidewalks. This map shows DC's are often blocked.

WALKING By Barbara Moreno (Guest Contributor) September 10, 2019 32

SHARE



A snapshot of sidewalk obstructions on NoMa's streets. by the author.

When Washingtonians like myself look for new apartments, we pay close attention to the walk score of a neighborhood. Any score upwards of 90 on a hundred point scale marks an area as a "walker's paradise," meaning major needs such as grocery stores and transit are within walking distance. However, what is *not* factored into the walkability score is the actual condition of the sidewalks.

ALSO OF INTEREST

WALKING
A pedestrian-only block in Alexandria may become a reality this spring 8

TRANSIT ANALYSIS
The good, the bad and the unexplained: what you need to know about the WMATA budget 27

DEVELOPMENT
Innovation Center Metro won't get a corporate name (for now), but a lot is already happening there. 17

TRANSIT
Baltimore's transit system is not meeting residents' needs. Can this plan change that? 17

Get daily updates via email

SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, in your vision
6. The **people matter** the most
7. Engage and serve **community**





Jon E. Froehlich
Lab Director
UW CS



Mikey Saugstad
Research Scientist
UW CS



Yochai Eisenberg
Assistant Professor
UIC Disab



Judy L. Shanley
Director
E Natio



Delphine Labbé
Assistant Professor
UIC Disab



Chu Li
PhD Student
UW CS



Manaswi Saha
PhD Student
UW CSE



David Jacobs
Professor
UMD CS



Kotaro Hara
PhD Student
UMD CS



Florian P. Thomas
Professor
HMSM Neuro



Katrina Ma
MS Student
UW HCDE



Joy Hammel
Professor
UIC Disab



Anthony Li
Undergrad
UMD CS



Robert Moore
Undergrad
UMD CS



Zachary Lawrence
Undergrad
UMD CS



Sho Kiami
Undergrad
UW CS



Ron Pechuk
Undergrad
UW CS



Devon Snyder
MS Student
UIC Depar



Aileen Zeng
Undergrad
UW CS



Sean Pannella
Undergrad
UMD CS



Victoria Le
Undergrad
UMD CS



Kie Fujii
PhD Student
HMSM Medic



Jin Sun
PhD Student
UMD CS



Minchu Kulkarni
Research Scientist
UW CS



Ladan Najafizadeh
MS Student
UMD CS



Aroosh Kumar
Undergrad
UW CS



Kevin Wu
Undergrad
UW CS



Nick Chung
Undergrad
UW CSE



David Phan
Undergrad
UW EE



Zoe Kaputa
Undergrad
UW CS



Galen Weld
PhD Student
UW CS



Jaye Ahn
MS Student
UW MHCI+D



Dylan Bunarto
Undergrad
UW CS



Alexander Zhang
Undergrad
UMD CS



Shiyu (Sherry) Wang
MS Student
UW HCDE



Udit Patwal
Undergrad
UW



Michael Duan
Undergrad
UW CS



Matthew Johnson
Undergrad
UW CS



Marianne Aubin Le Quéré
MS Student
UW HCDE



Jey Yoo
MS Student
UW MHCI+D

Tanaya Sharma
High School Student
NCHS

Alex Chen
MS Student
UW MHCI+D

Teja Maddali
PhD Student
UMD CS

Richard McGovern
MS Student
UW iSchool

Ilia Savin
MS Student
UW HCDE

Kevin Chen
High School Student
CHS

Allison Gu
Undergrad
UW CS

Jeremy Freiburger
Undergrad
UW CSE

Esther Jang
PhD Student
UW CSE

Jules Ropars
Undergrad
UW CSE

Jiamae Wang
Undergrad
UW CS

Leon Li
Undergrad
UW CS

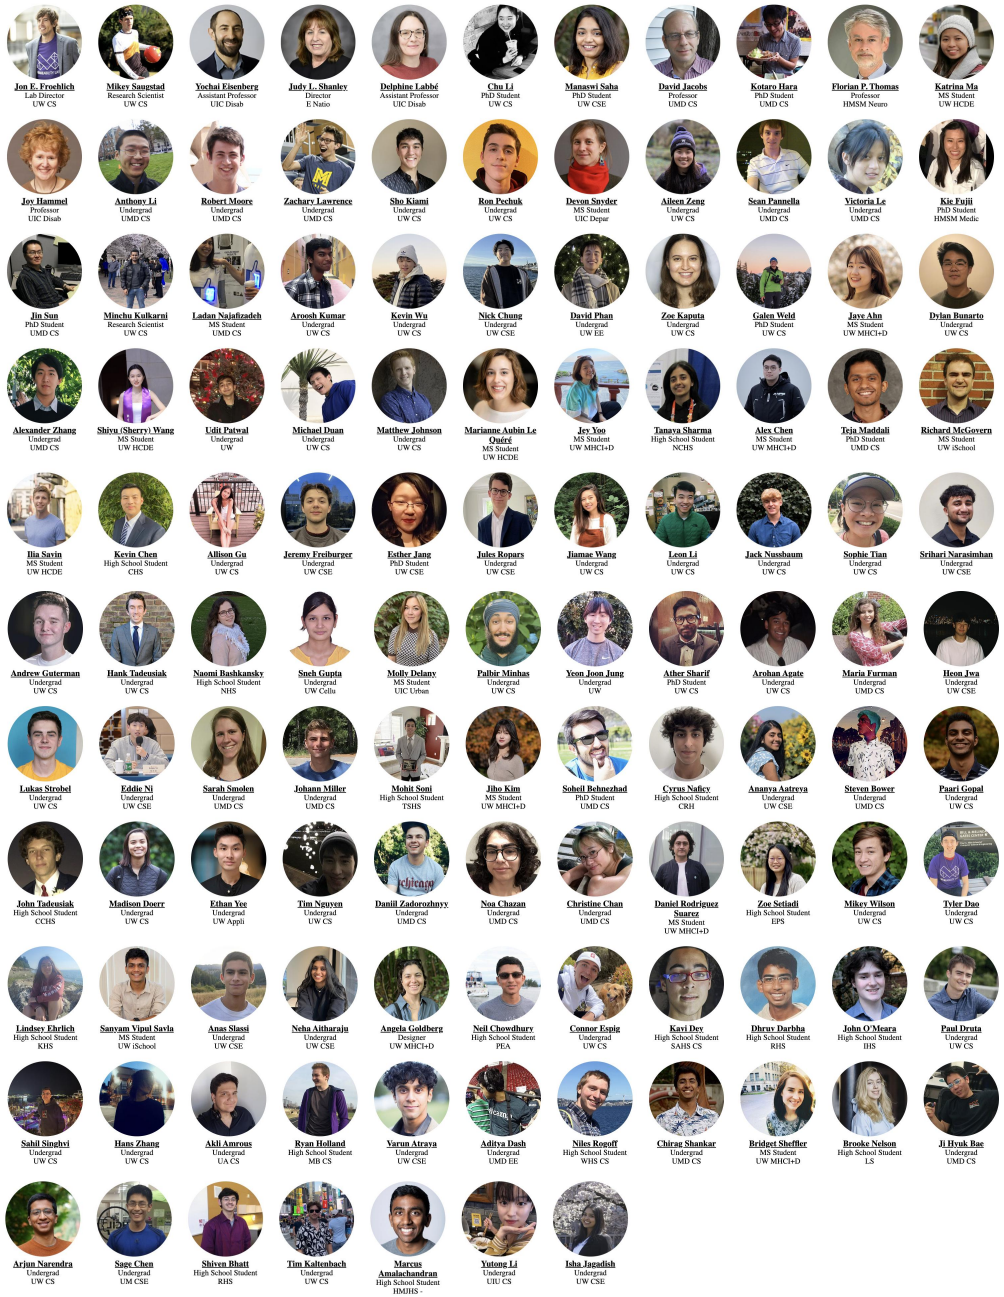
Jack Nussbaum
Undergrad
UW CS

Sophie Tian
Undergrad
UW CS

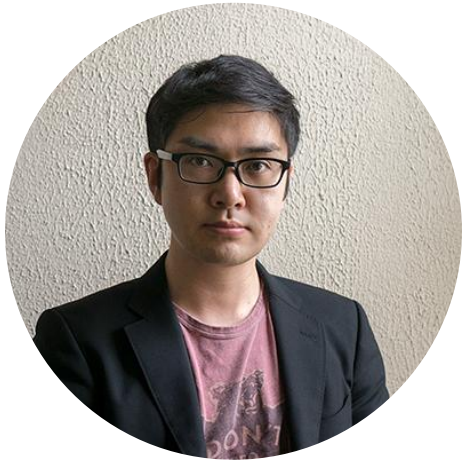
Srihari Narasimhan
Undergrad
UW CSE

PROJECT SIDEWALK

BUILT BY STUDENTS, FACULTY, & OUR COMMUNITY PARTNERS



PROJECT SIDEWALK



Kotaro Hara
UMD CS PhD'16
2012-2016
Now: Prof at SMU



Manaswi Saha
UW CSE PhD'22
2016-2022
Now: Accenture Labs



Chu Li
UW CSE PhD
2022-present
On Job Market!



Yochai Eisenberg
UIC Professor
Disability & Public Health
2020-present



Mikey Saugstad
Research Engineer
2017-present

MAKEABILITY LAB

SOCIAL IMPACT AREAS



ACCESSIBILITY



HEALTH
& WELLNESS



ENVIRONMENTAL
SUSTAINABILITY



STEM
EDUCATION

MAKEABILITY LAB

MY PHD STUDENTS



Arnavi Chheda-Kothary



Kotaro Hara



Liang He



Jared Hwang



Dhruv Jain



Seokbin Kang



Jaewook Lee



Chu Li



Matthew L Mauriello



Manaswi Saha



Mikey Saugstad



Lee Stearns



Xia Su

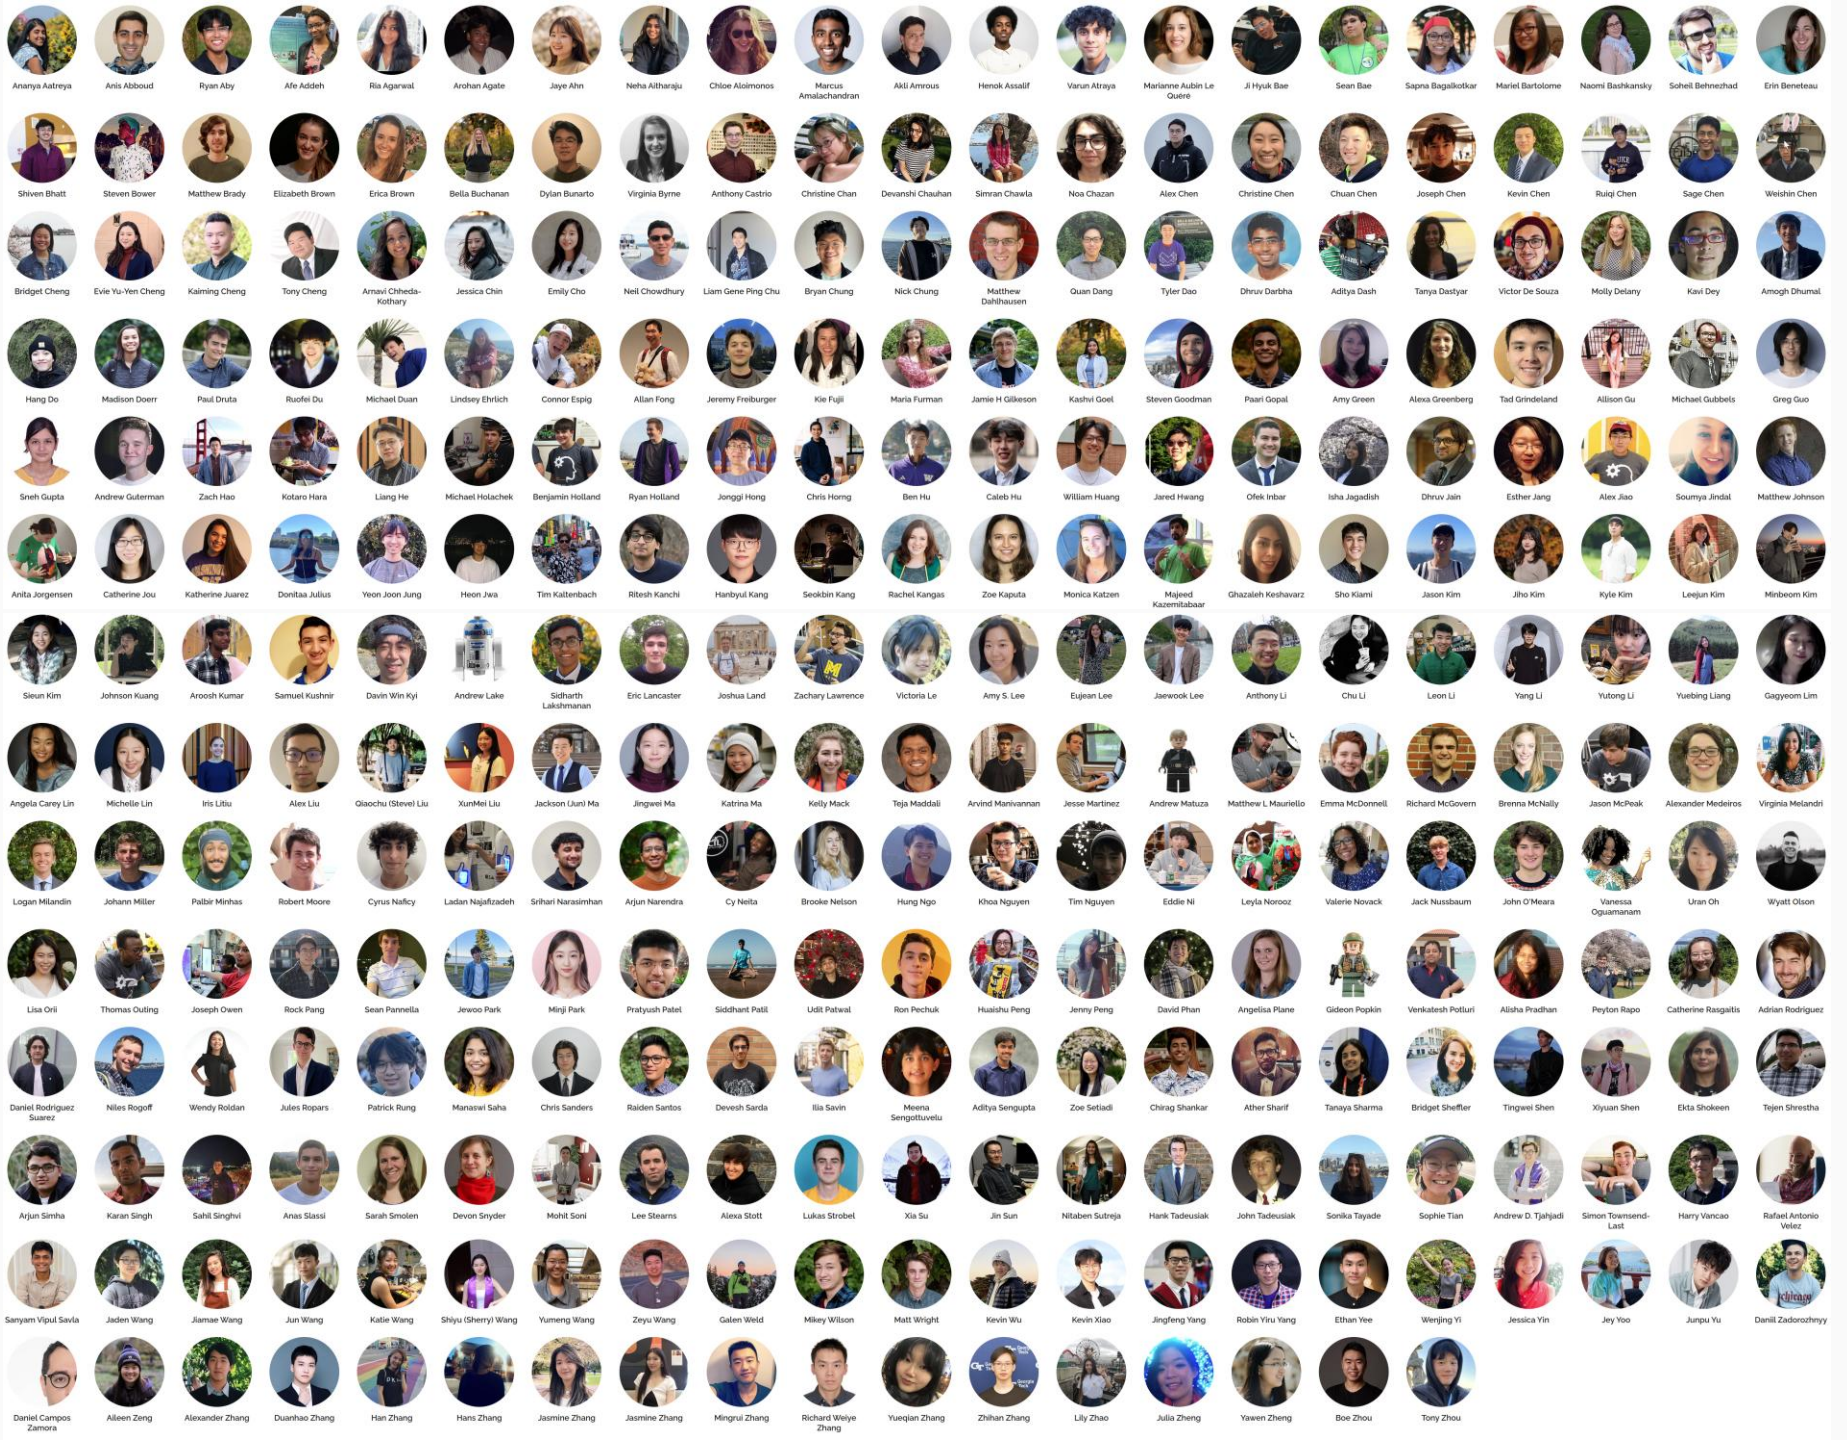


Daniel Campos Zamora

MAKEABILITY LAB

290+ STUDENTS

MAKEABILITY LAB



SIGCHI SOCIETAL IMPACT

SEVEN LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your vision
6. The **people matter** the most
7. Engage and serve **community**

SIGCHI SOCIETAL IMPACT

~~SEVEN~~ EIGHT LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your vision
6. The **people matter** the most
7. Engage and serve **community**
8. **Better to try and fail** than to not try at all

The biggest **regrets**

An official website of the United States government Here's how you know



NSF Graduate Research Fellowship Program (NSF GRFP)

Out now: View the NSF GRFP Notice of Funding Opportunity.

Policies and procedures for recipients: View the GRFP Administrative Guide.

NSF Graduate Research Fellowship Program (NSF GRFP)

Applicants +

Reference Writers

Reviewers

The purpose of the U.S. National Science Foundation Graduate Research Fellowship Program (NSF GRFP) is to help ensure the quality, vitality, and strength of the scientific and engineering workforce of the United States.

2027 NSF GRFP Notice of Funding Opportunity

Applicant deadlines

Feedback

An official website of the United States government Here's how you know



NSF Graduate Research Fellowship Program (NSF GRFP)

BETTER TO TRY AND FAIL THAN TO NOT TRY AT ALL

Out now: View the NSF GRFP Notice of Funding Opportunity.

Policies and procedures for recipients: View the GRFP Administrative Guide.

NSF Graduate Research Fellowship Program (NSF GRFP)

Applicants +

Reference Writers

Reviewers

The purpose of the U.S. National Science Foundation Graduate Research Fellowship Program (NSF GRFP) is to help ensure the quality, vitality, and strength of the scientific and engineering workforce of the United States.

2027 NSF GRFP Notice of Funding Opportunity

Applicant deadlines

Feedback

SIGCHI SOCIETAL IMPACT

~~SEVEN~~ EIGHT LESSONS LEARNED

1. Always **another mountain**
2. **Choose** your mountains carefully
3. **Persistence** is key
4. Know & use your **leverage points**
5. **Believe in yourself**, believe in your vision
6. The **people matter** the most
7. Engage and serve **community**
8. **Better to try and fail** than to not try at all



2024 ACM SIGCHI Societal Impact Award External

[redacted]

Tue, Feb 20, 2024, 3:24 PM

to me

Dear Jon,

that you that your nomination was not selected for the 2024 ACM SIGCHI

nominations for this award was extremely strong.

ider submitting

2025 ACM SIGCHI Societal Impact Award External

[redacted]

to me

Feb 25, 2025, 1:43 PM

Dear Jon,

I regret to inform you that you that your nomination was not selected for the 2025 ACM SIGCHI Societal Impact Award. The pool and caliber of the nominations for this award was extremely strong. We appreciate your interest in the Societal Impact Award and hope that you will consider submitting nominations in the future.

Sincerely,
The SIGCHI Societal Impact Award Committee

2024 ACM SIGCHI Societal Impact Award External

Tue, Feb 20, 2024, 3:24 PM

2025 ACM SIGCHI Societal Impact Award External

2026 ACM SIGCHI Societal Impact Award External

Elizabeth M Gerber <egerber@northwestern.edu>

to jonf@cs.uw.edu, Kentaro

Sun, Jan 25, 3:12 PM

Dear Jon,

I am pleased to let you know that you will receive the ACM SIGCHI Societal Impact Award at CHI 2026. You have earned this award for the novelty and depth of your research, efforts in scaling the work for direct impact on society, and leadership in our field.

We hope you will celebrate with friends and close colleagues, but ask that you not share widely until ACM SIGCHI announces all the award winners for 2026.

Kentaro Toyama, ACM SIGCHI Awards Chair, will follow up with more details regarding the honorarium, talk, and travel details.

Congratulations again for your incredible work! Please let me know if you have any questions.

Sincerely,



MAKEABILITY LAB

DESIGN USE BUILD

CREATE
Center for Research and Education on Accessible Technology and Experiences

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

UNIVERSITY of WASHINGTON