Characterizing Physical World Accessibility at Scale

Combining Google Street View, Crowdsourcing, and Automated Methods to Collect Accessibility Data about the Built Environment

Kotaro Hara

Carnegie Mellon University

UNIVERSITY OF MARYLAND

COMPUTER SCIENCE UNIVERSITY OF MARYLAND

HCIL

makeability lab

The Americans with Disabilities Act

mandates that new constructions and
alterations of the built environment
are to be **accessible for everyone.**

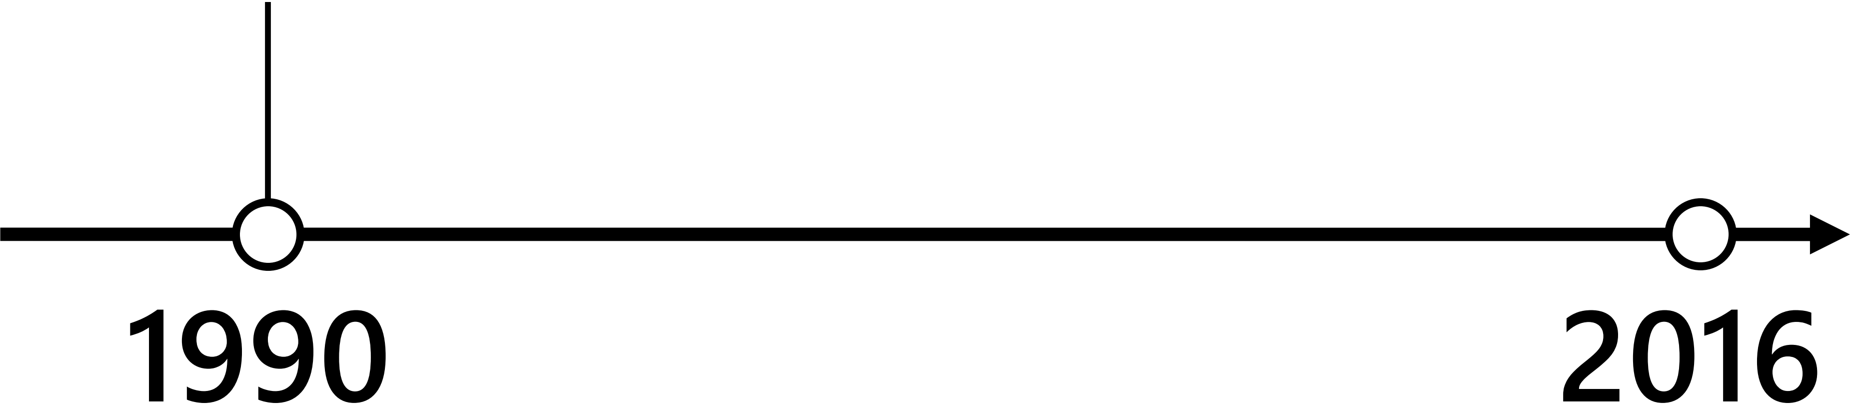

FRONT PAGE
POLITICS
ENTERTAINMENT
WHAT'S WORKING
HEALTHY LIVING
WORLDPOST
HIGHLINE
HUFFPOST LIVE
ALL SECTIONS
25 Years After The A... Navigating New York City Is
Still Daunting For T...
07/27/2015 03:55 pm ET
Savannah OLeary
Multimedia Producer
02:15
Lisa's
Sections
The Washington Post
+ More
Sign In
Subscribe
Wonkblog
The inequality of sidewalks
9
Save for Later
Reading List
By Emily Badger January 15
Follow @emilymbadger
Most Read
1 Legal marijuana is finally doing what the drug war couldn't
2 The magical thing eating chocolate does to your brain
3 The most racist places in America, according to Google
4 Georgians form human chain to protest talks with Gazprom
5 The key difference between how Trump and Romney made their money

2016

30.6

million U.S. adults with mobility impairment


15.2
million use an assistive
mobility aid

Missing Curb Ramp

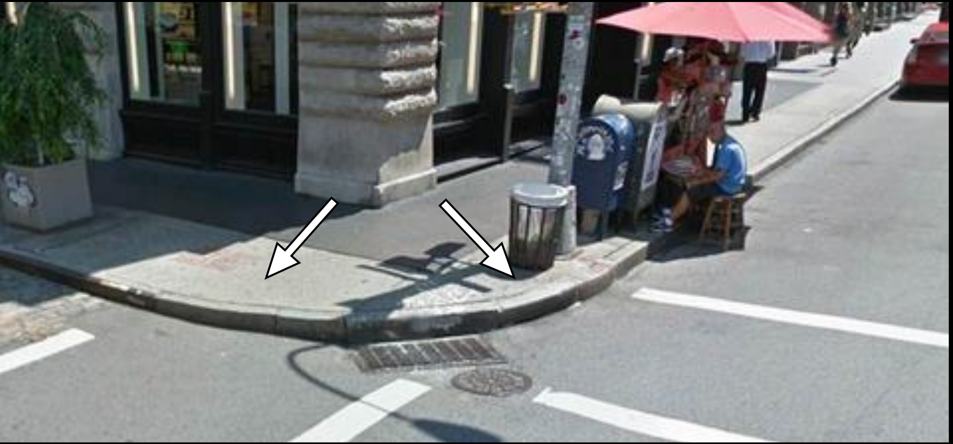

Obstacle

Surface Problem

No Sidewalk

Meyers *et al* [Soc. Sci. & Med. 2002], Rimmer *et al.* [AJPM 2004],

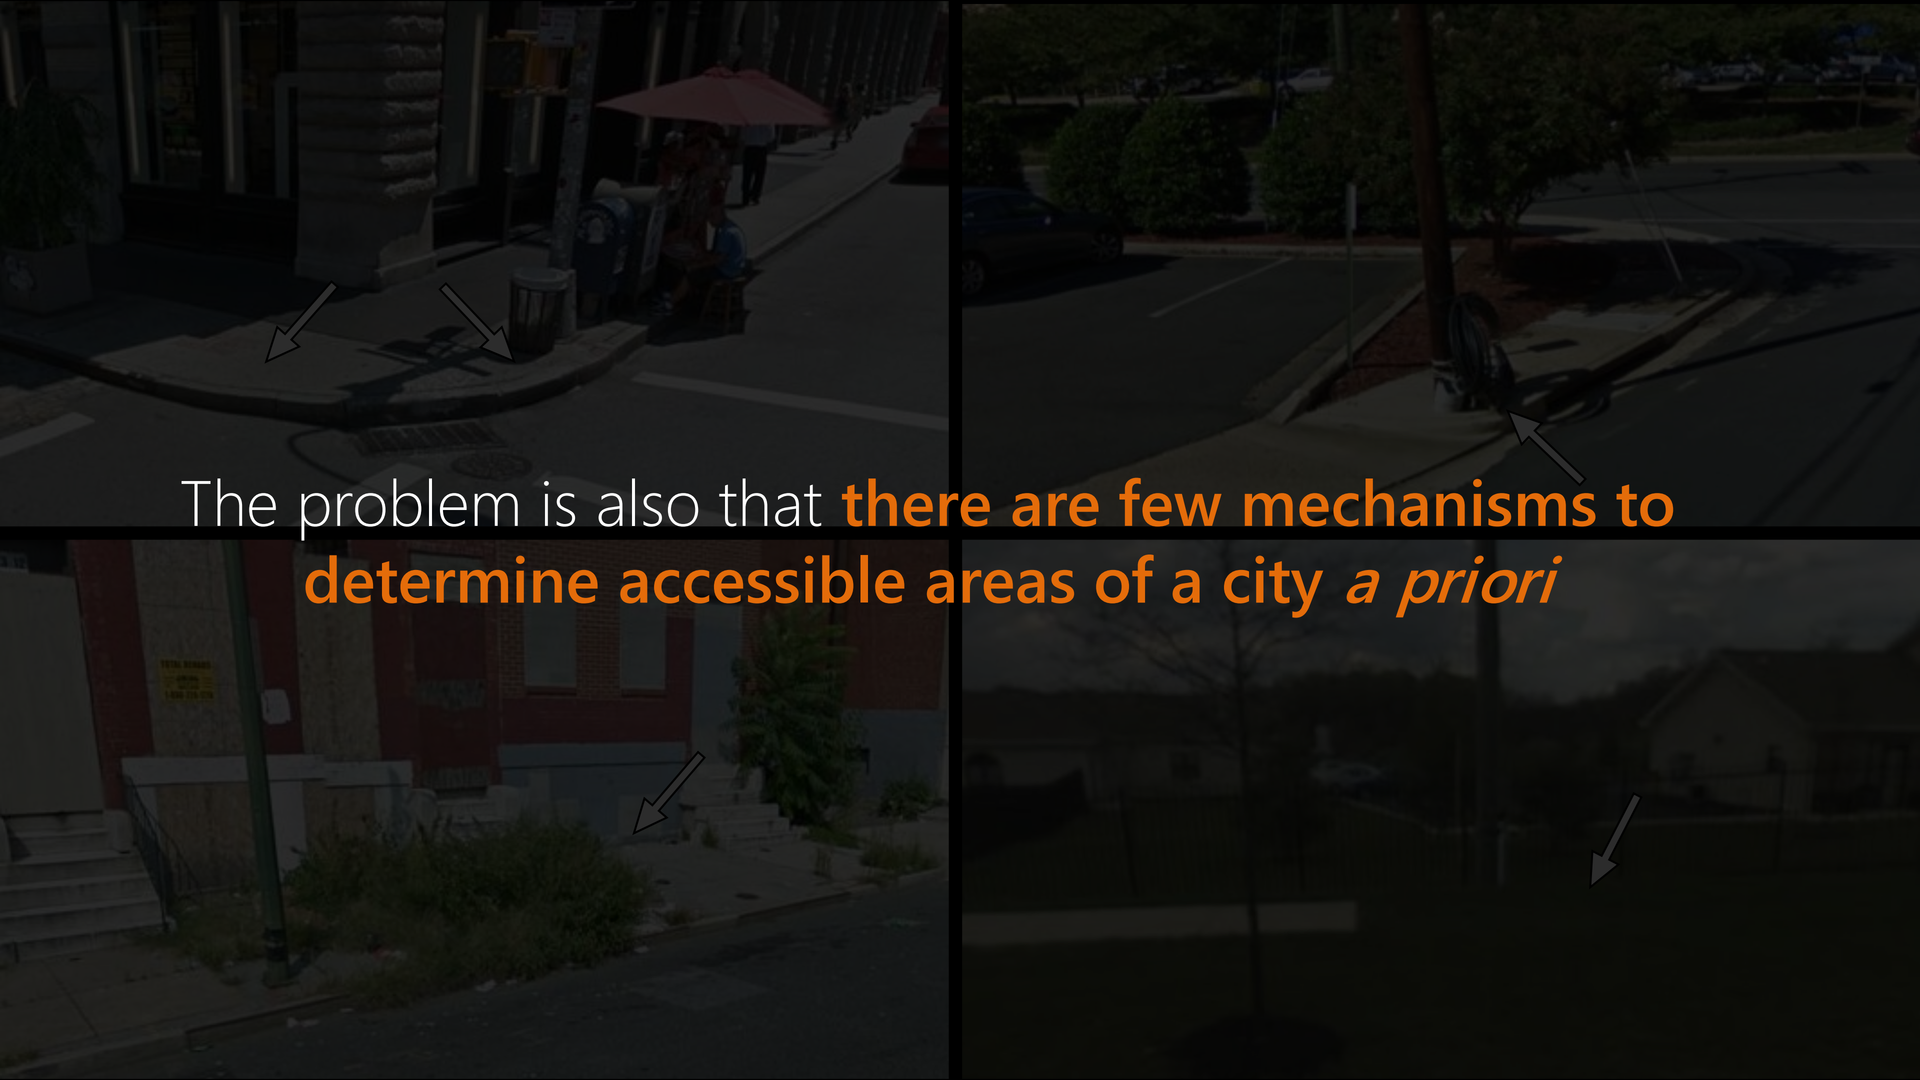
The problem is also that **there are few mechanisms to determine accessible areas of a city *a priori***

The National Council on Disability noted that there is **no comprehensive information** on “the degree to which sidewalks are accessible” in cities.

National Council on Disability, 2007
The impact of the Americans with Disabilities Act: Assessing the progress toward achieving the goals of the ADA

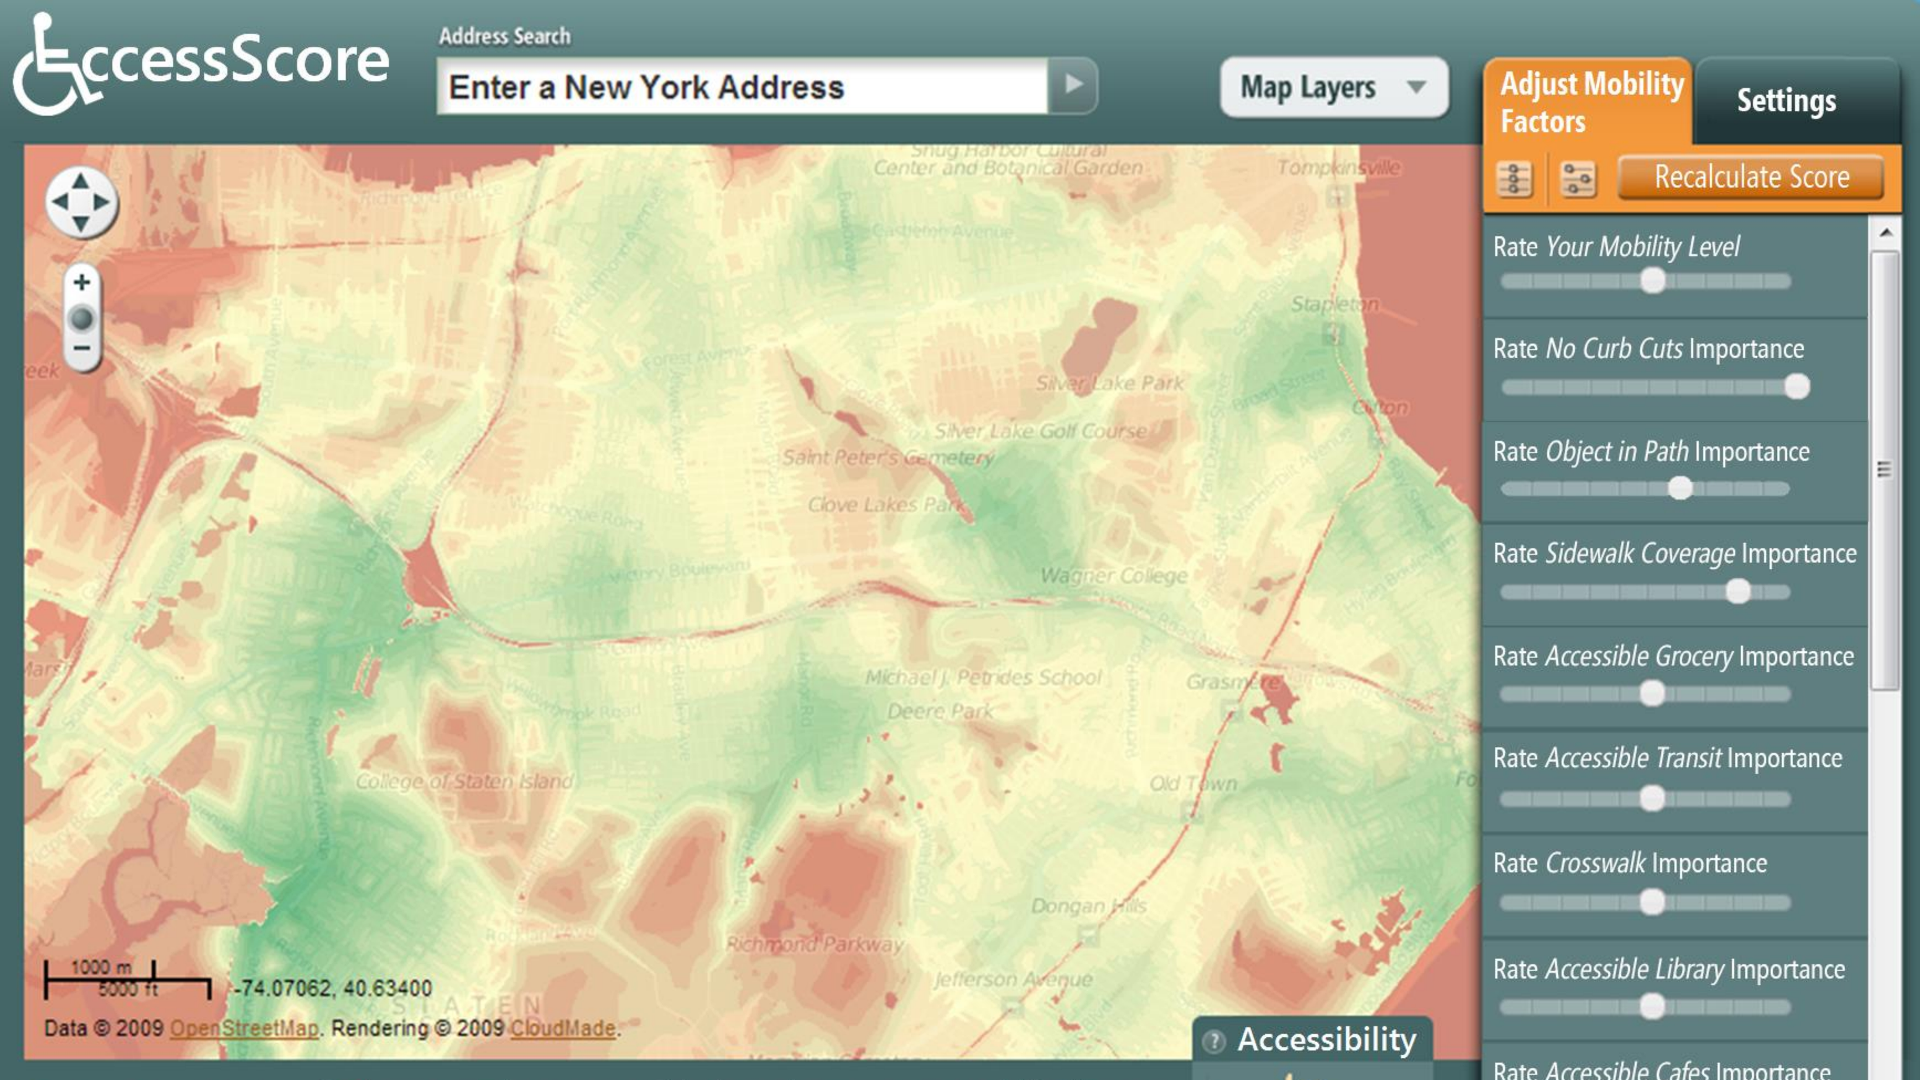
AccessScore
Address Search
Enter a New York Address
Map Layers
Adjust Mobility Factors
Settings
Recalculate Score
Rate Your Mobility Level
Rate No Curb Cuts Importance
Rate Object in Path Importance
Rate Sidewalk Coverage Importance
Rate Accessible Grocery Importance
Rate Accessible Transit Importance
Rate Crosswalk Importance
Rate Accessible Library Importance
Rate Accessible Cafes Importance
Snug Harbor Cultural Center and Botanical Garden
Tompkinsville
Stapleton
Silver Lake Park
Silver Lake Golf Course
Saint Peter's Cemetery
Clove Lakes Park
Wagner College
Michael J. Petrides School
Deere Park
Grasmere
College of Staten Island
Old Town
Dongan Hills
Richmond Parkway
Jefferson Avenue
1000 m
5000 ft
-74.07062, 40.63400
Data © 2009 OpenStreetMap. Rendering © 2009 CloudMade.
Accessibility

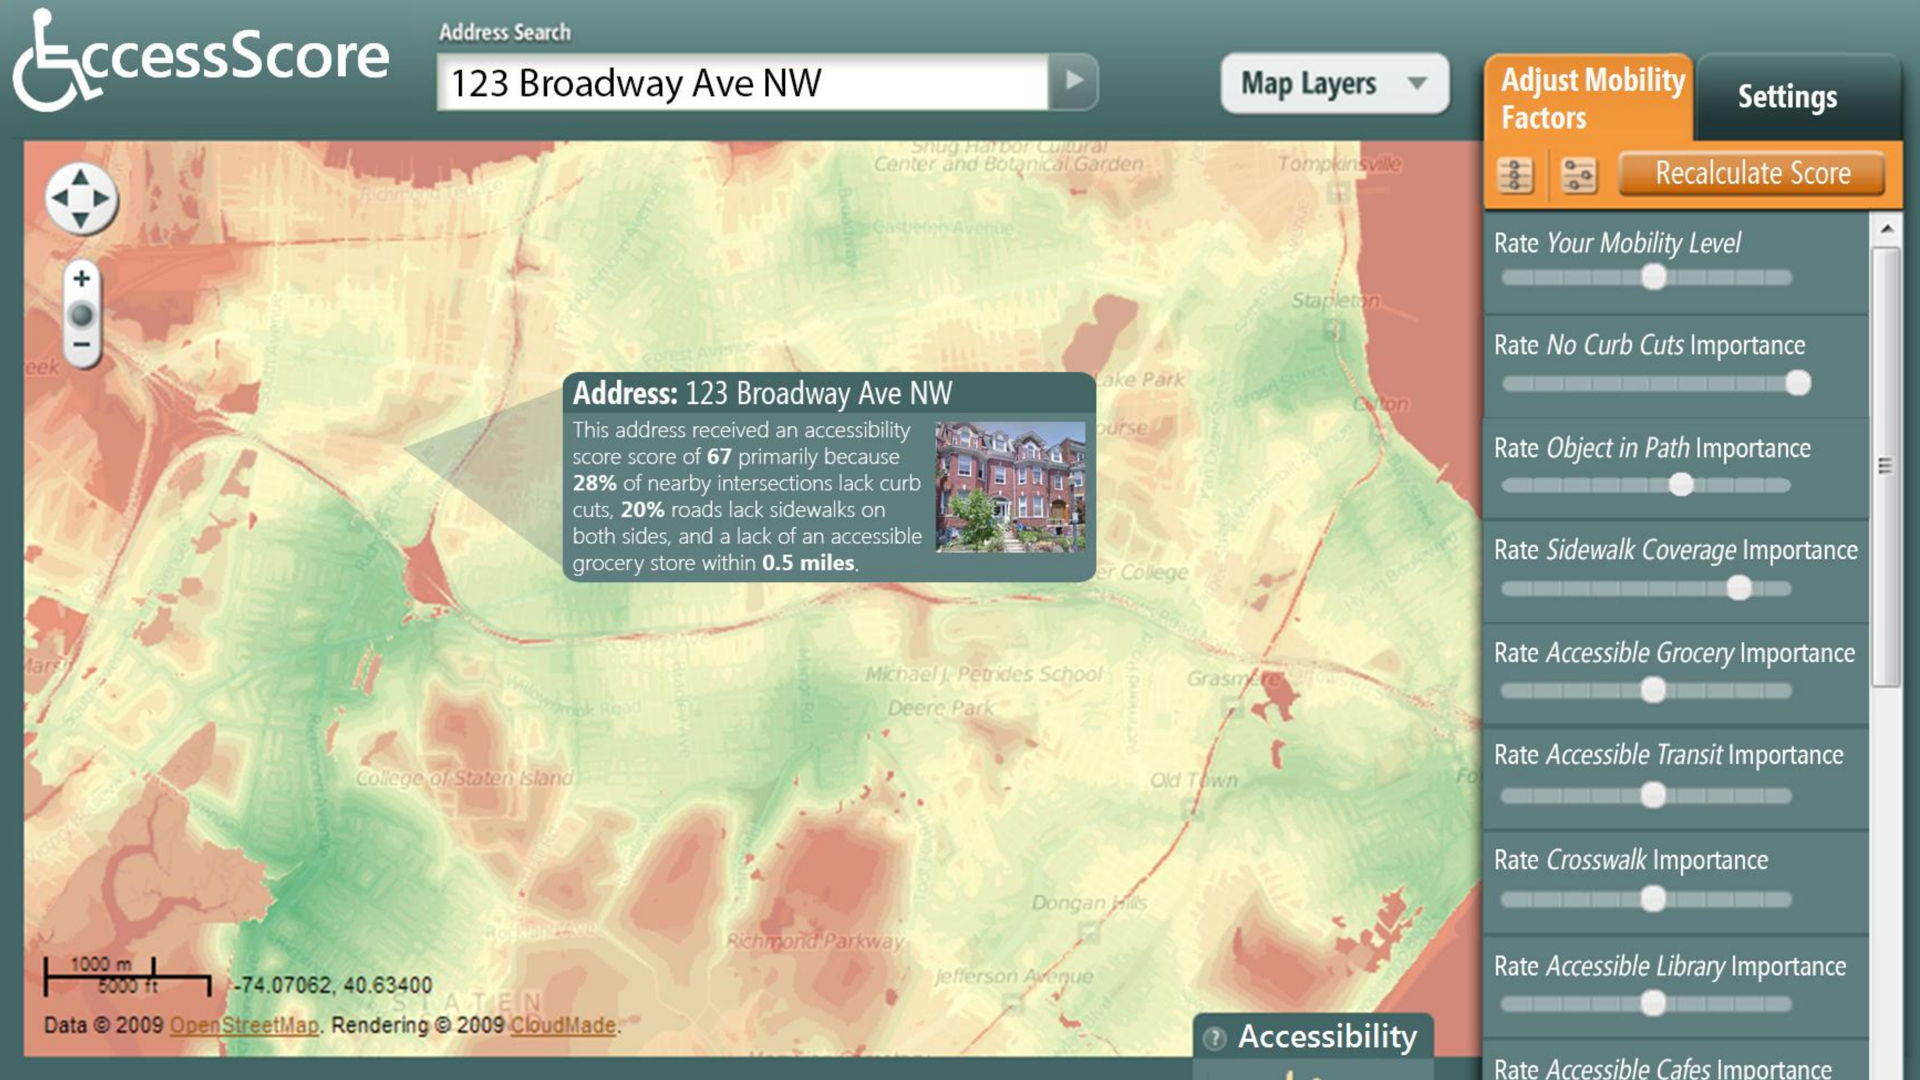

AccessScore
Address Search
123 Broadway Ave NW
Map Layers
Adjust Mobility Factors
Settings
Recalculate Score
Rate Your Mobility Level
Rate No Curb Cuts Importance
Rate Object in Path Importance
Rate Sidewalk Coverage Importance
Rate Accessible Grocery Importance
Rate Accessible Transit Importance
Rate Crosswalk Importance
Rate Accessible Library Importance
Rate Accessible Cafes Importance
Address: 123 Broadway Ave NW
This address received an accessibility score score of 67 primarily because 28% of nearby intersections lack curb cuts, 20% roads lack sidewalks on both sides, and a lack of an accessible grocery store within 0.5 miles.
Snug Harbor Cultural Center and Botanical Garden
Tompkinsville
Stapleton
Lake Park
Michael J. Petrides School
Deere Park
Grasmere
College of Staten Island
Old Town
Dongan Hills
Richmond Parkway
Jefferson Avenue
1000 m
5000 ft
-74.07062, 40.63400
Data © 2009 OpenStreetMap. Rendering © 2009 CloudMade.
Accessibility

Accessibility-aware Navigation

Our vision is to design methods to **scalably collect street-level accessibility data** that enables technologies that support mobility impaired people to navigate the physical environment

Traditional Walkability Audits

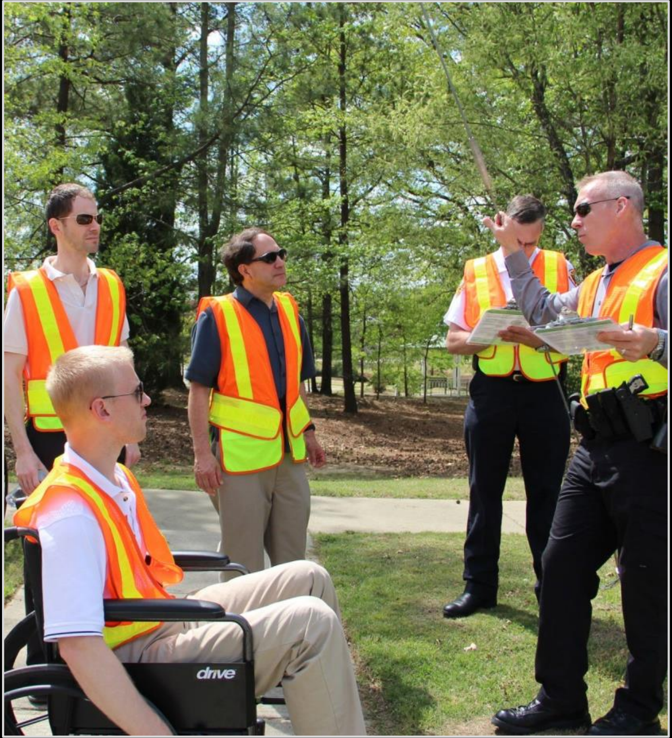

Walkability Audit
Wake County, North Carolina

Walkability Audit
Wake County, North Carolina

Safe Routes to School Walkability Audit
Rock Hill, South Carolina

Background

T... ty Audits

Take a walk and use this checklist to rate your neighborhood's walkability.

How walkable is your community?

Location of walk ____________________

____

Rating Scale: 1 awful · 2 many problems · 3 some problems · 4 good · 5 very good · 6 excellent

1. Did you have room to walk?

- ☐ Yes
- ☐ Some problems:
  - ☐ Sidewalks or paths started and stopped
  - ☐ Sidewalks were broken or cracked
  - ☐ Sidewalks were blocked with poles, signs, shrubbery, dumpsters, etc.
  - ☐ No sidewalks, paths, or shoulders
  - ☐ Too much traffic
  - ☐ Something else ____________
  - Locations of problems: ____________

Rating: (circle one) 1 2 3 4 5 6

2. Was it easy to cross streets?

- ☐ Yes
- ☐ Some problems:
  - ☐ Road was too wide
  - ☐ Traffic signals made us wait too long or did not give us enough time to cross
  - ☐ Needed striped crosswalks or traffic signals
  - ☐ Parked cars blocked our view of traffic
  - ☐ Trees or plants blocked our view of traffic
  - ☐ Needed curb ramps or ramps needed repair
  - ☐ Something else ____________
  - Locations of problems: ____________

Rating: (circle one) 1 2 3 4 5 6

3. Did drivers behave well?

- ☐ Yes
- ☐ Some problems: Drivers...
  - ☐ Backed out of driveways without looking
  - ☐ Did not yield to people crossing the street
  - ☐ Turned into people crossing the street
  - ☐ Drove too fast
  - ☐ Sped up to make it through traffic lights or drove through traffic lights?
  - ☐ Something else ____________
  - Locations of problems: ____________

Rating: (circle one) 1 2 3 4 5 6

4. Was it easy to follow safety rules? Could you and your child...

- ☐ Yes ☐ No Cross at crosswalks or where you could see and be seen by drivers?
- ☐ Yes ☐ No Stop and look left, right and then left again before crossing streets?
- ☐ Yes ☐ No Walk on sidewalks or shoulders facing traffic where there were no sidewalks?
- ☐ Yes ☐ No Cross with the light?
- Locations of problems: ____________

Rating: (circle one) 1 2 3 4 5 6

5. Was your walk pleasant?

- ☐ Yes
- ☐ Some unpleasant things:
  - ☐ Needed more grass, flowers, or trees
  - ☐ Scary dogs
  - ☐ Scary people
  - ☐ Not well lighted
  - ☐ Dirty, lots of litter or trash
  - ☐ Something else ____________
  - Locations of problems: ____________

Rating: (circle one) 1 2 3 4 5 6

How does your neighborhood stack up? Add up your ratings and decide.

1. ____
2. ____
3. ____
4. ____
5. ____

Total ____

- 26–30 Celebrate! You have a great neighborhood for walking.
- 21–25 Celebrate a little. Your neighborhood is pretty good.
- 16–20 Okay, but it needs work.
- 11–15 It needs lots of work. You deserve better than that.
- 5–10 Call out the National Guard before you walk. It's a disaster area.

Now that you've identified the problems, go to the next page to find out how to fix them.

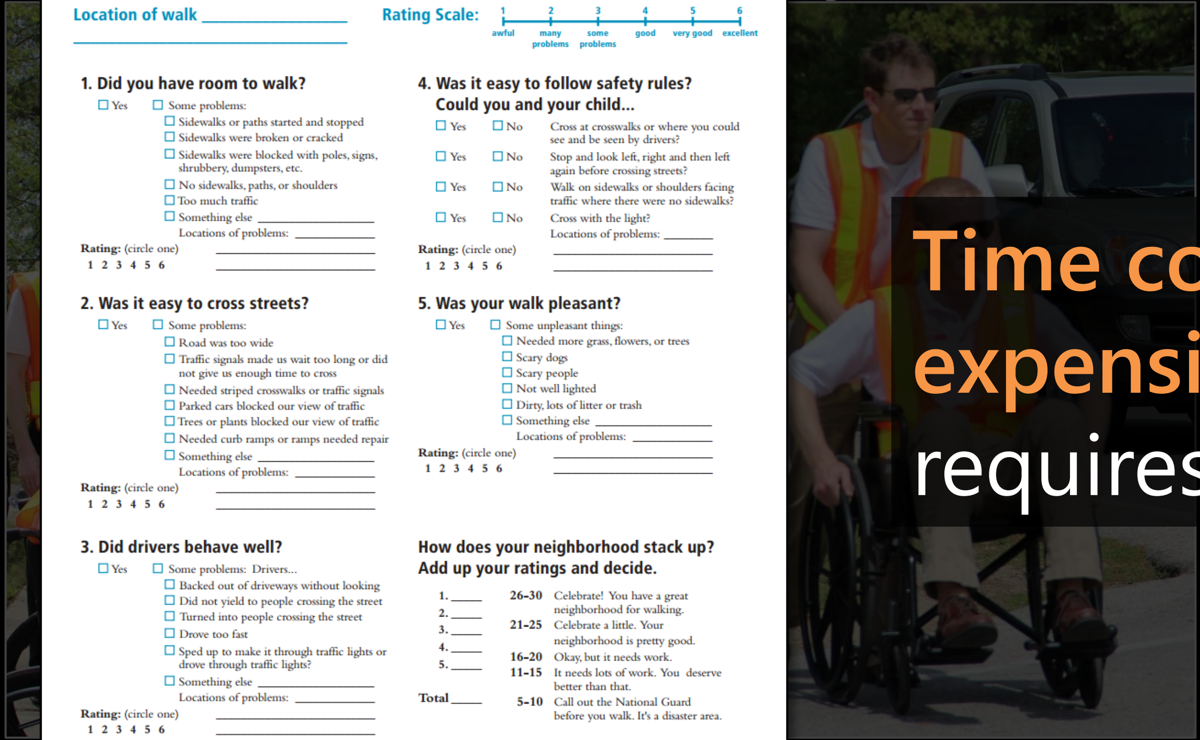

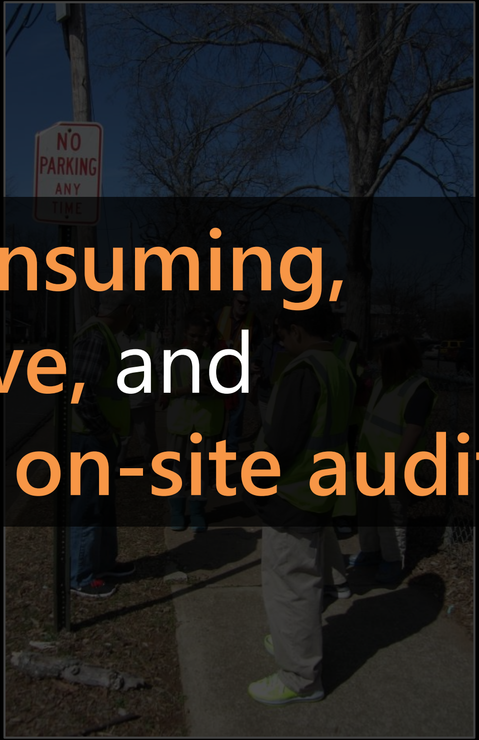

Time consuming, expensive, and requires on-site audit

Wa...
Wa...

...ility Audit
...County, North Carolina

Safe Routes to School Walkability Audit
Rock Hill, South Carolina

Mobile Reporting Solutions

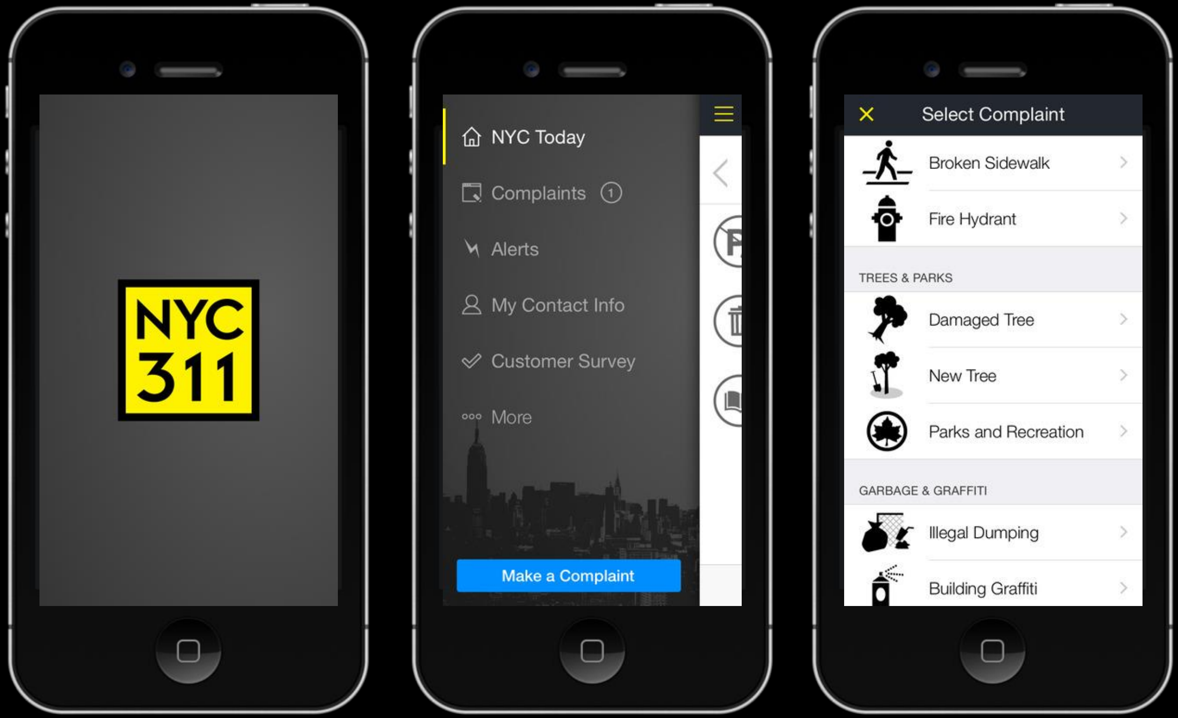

http://www1.nyc.gov/311/index.page

Mobile Reporting Solutions

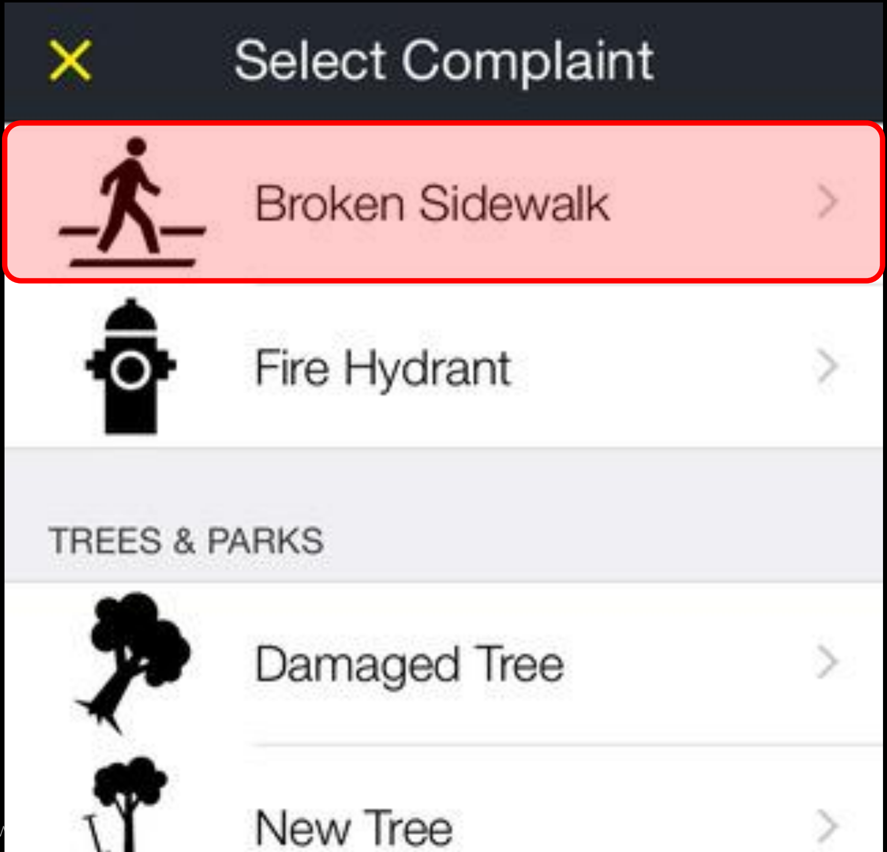

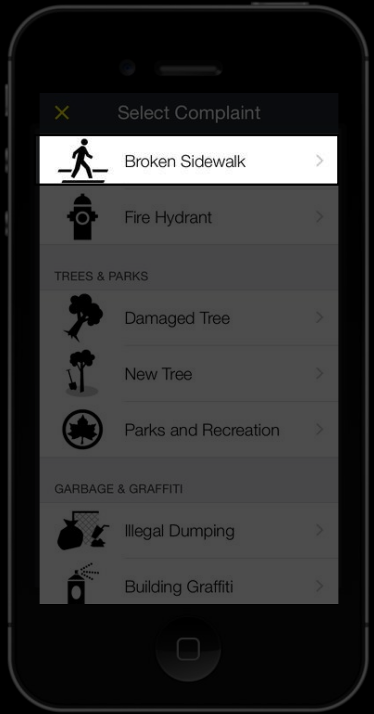

http://v

Mobile Reporting Solutions

Similar to physical audits, these tools are built for ***in situ* reporting** and do not support remote, virtual inquiry—which **limits scalability**

Not designed for **accessibility data collection**

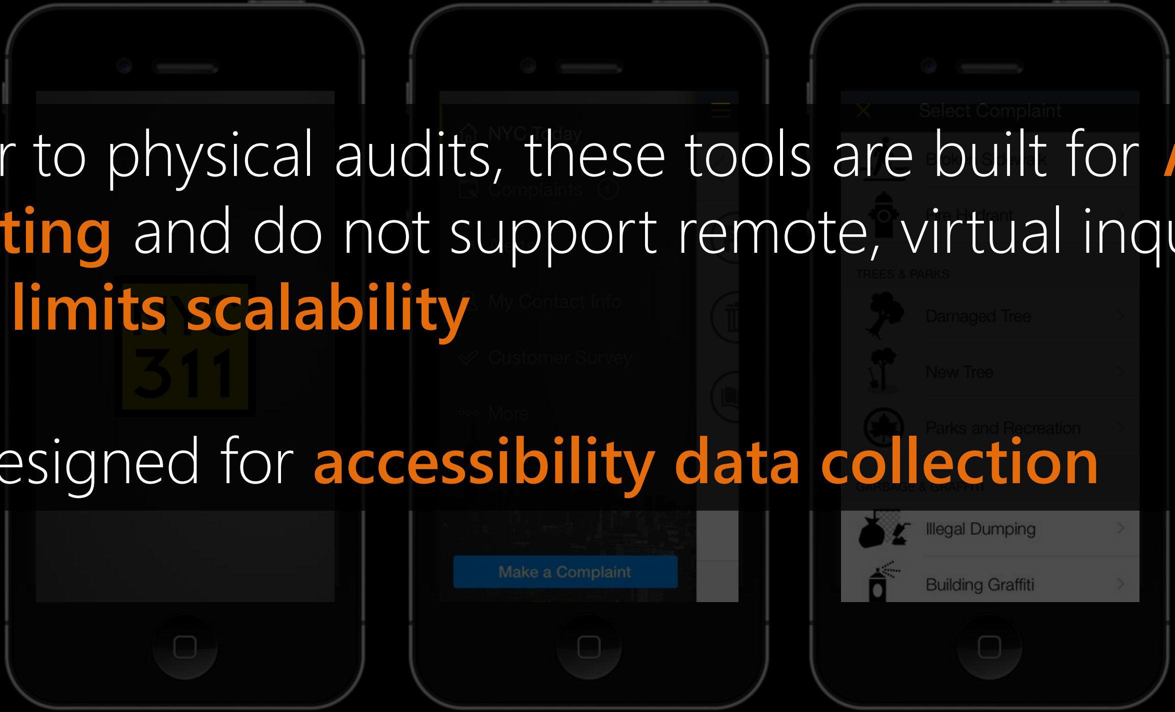

http://www1.nyc.gov/311/index.page

Mark & Find Accessible Businesses

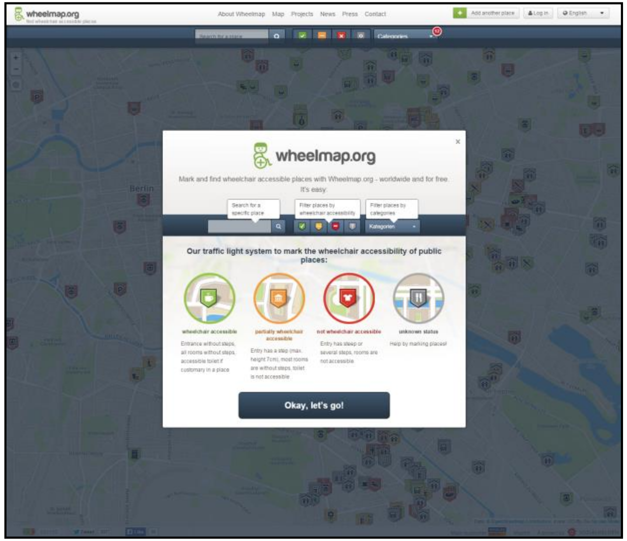

wheelmap.org

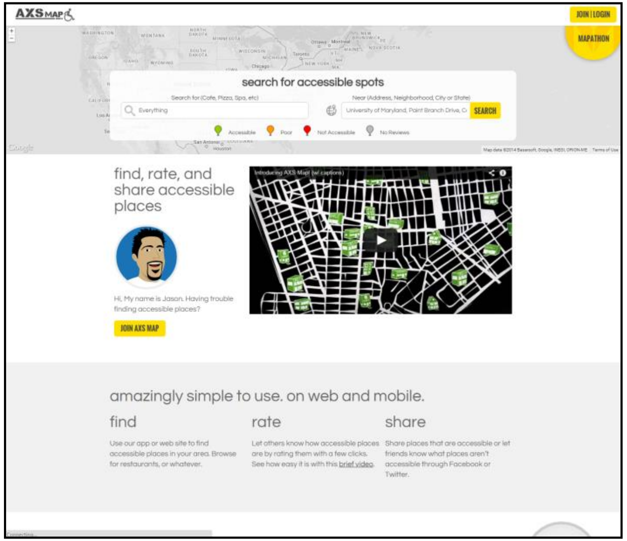

axsmap.com

Focuses on businesses rather than streets/sidewalks

Model is still to report on places you've visited

Our Approach: Remotely collect street-level accessibility information from Google Street View (GSV) using crowdsourcing and computation

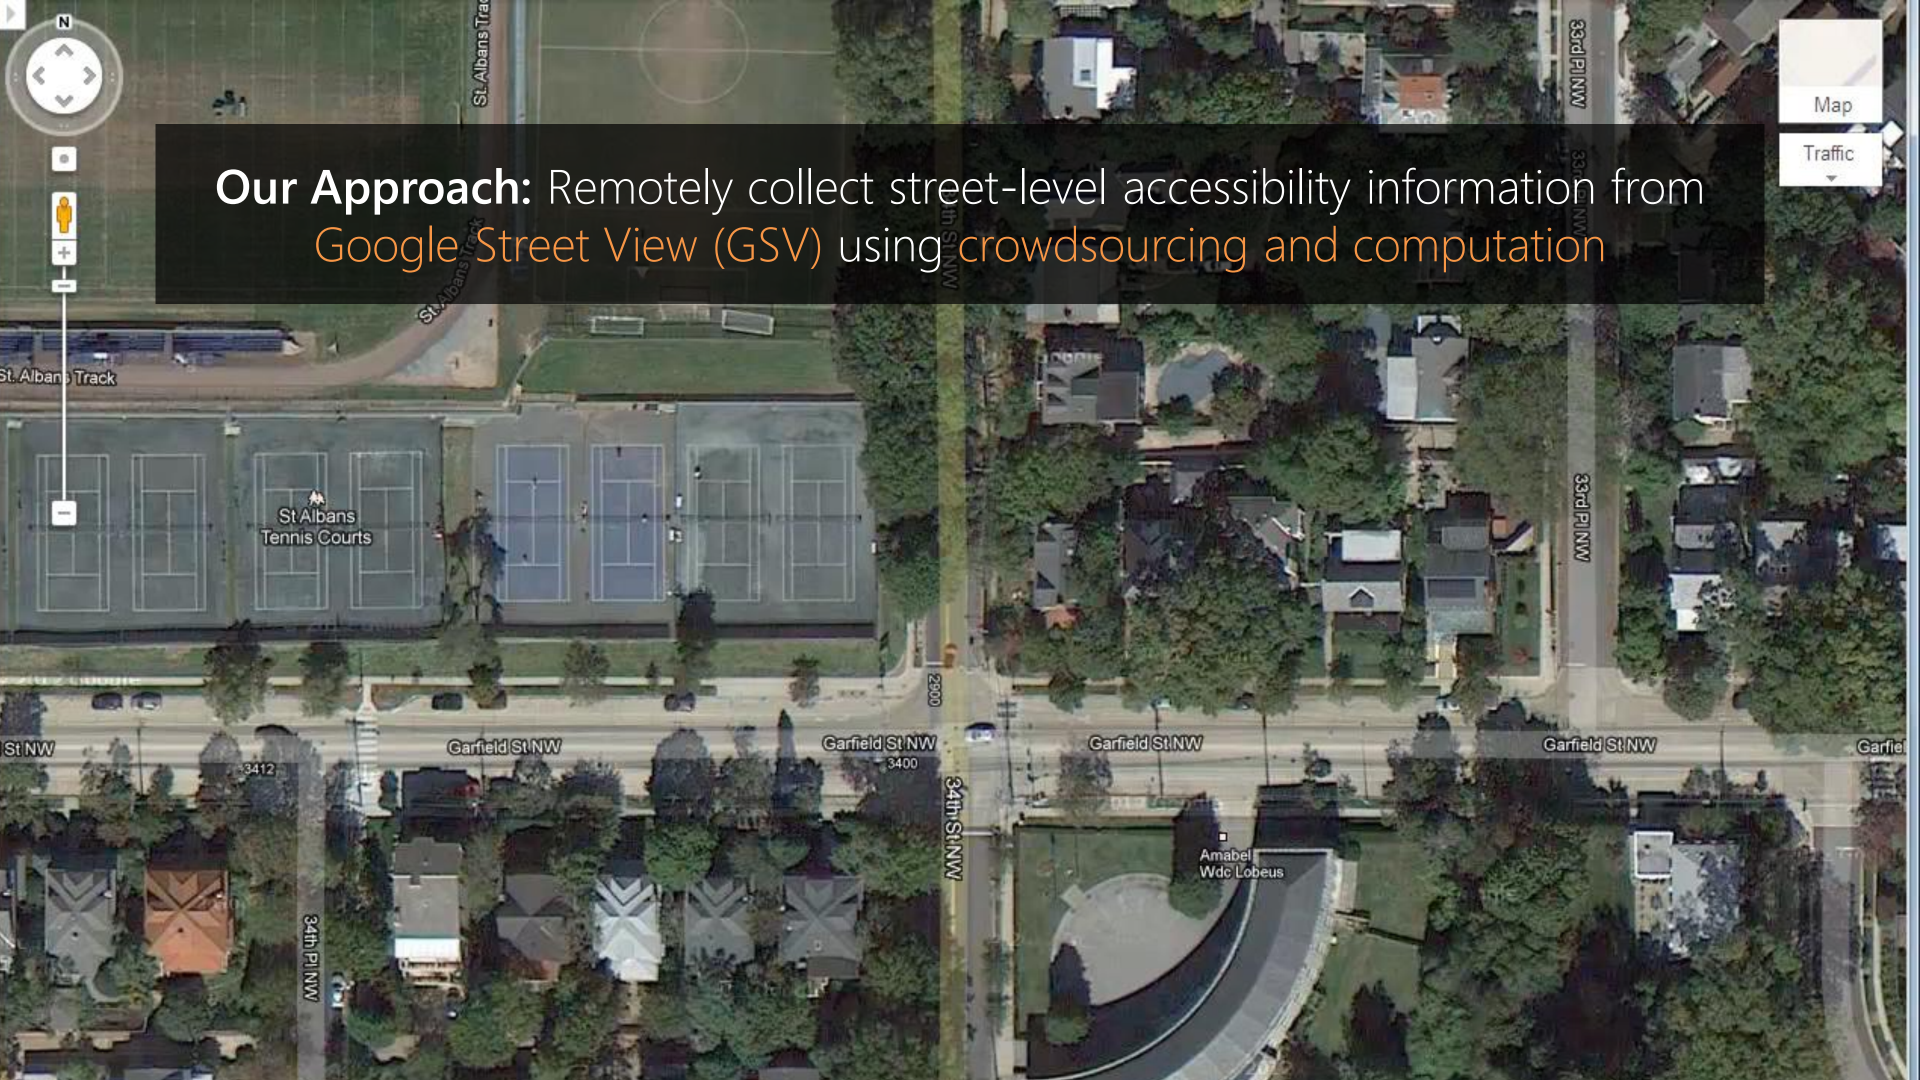

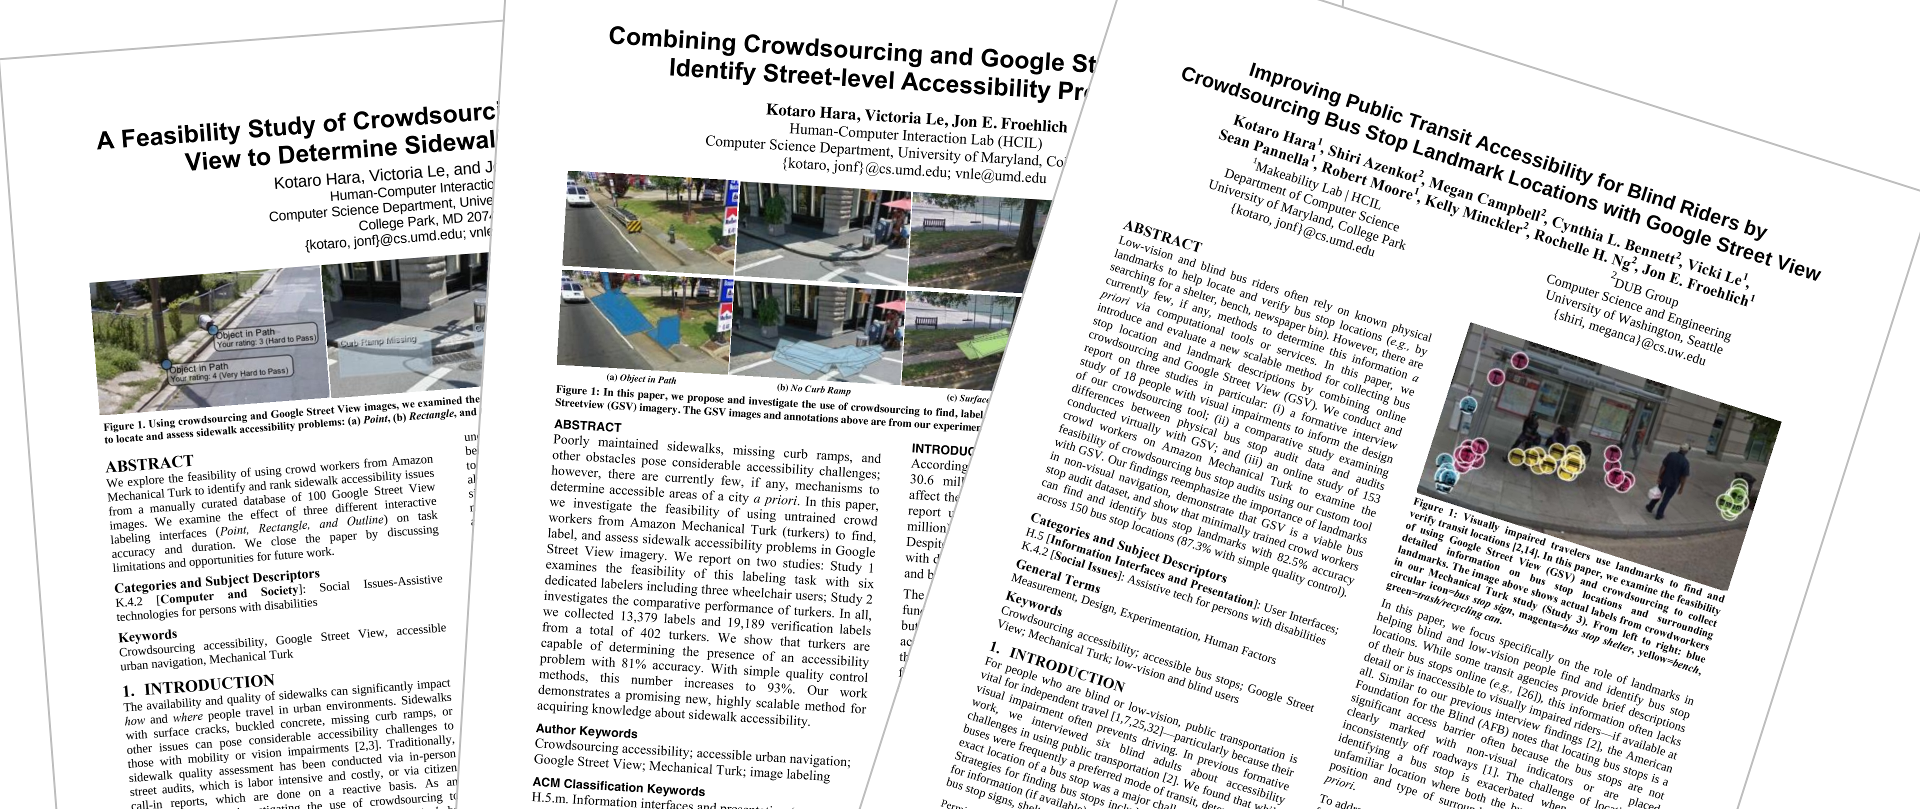

Crowdsourcing Accessibility Data from Google Street View

Hara K., Le V., & Froehlich. J.E. 2012, 2013; Hara K., *et al*. 2013

What **accessibility problems** exist in this image?

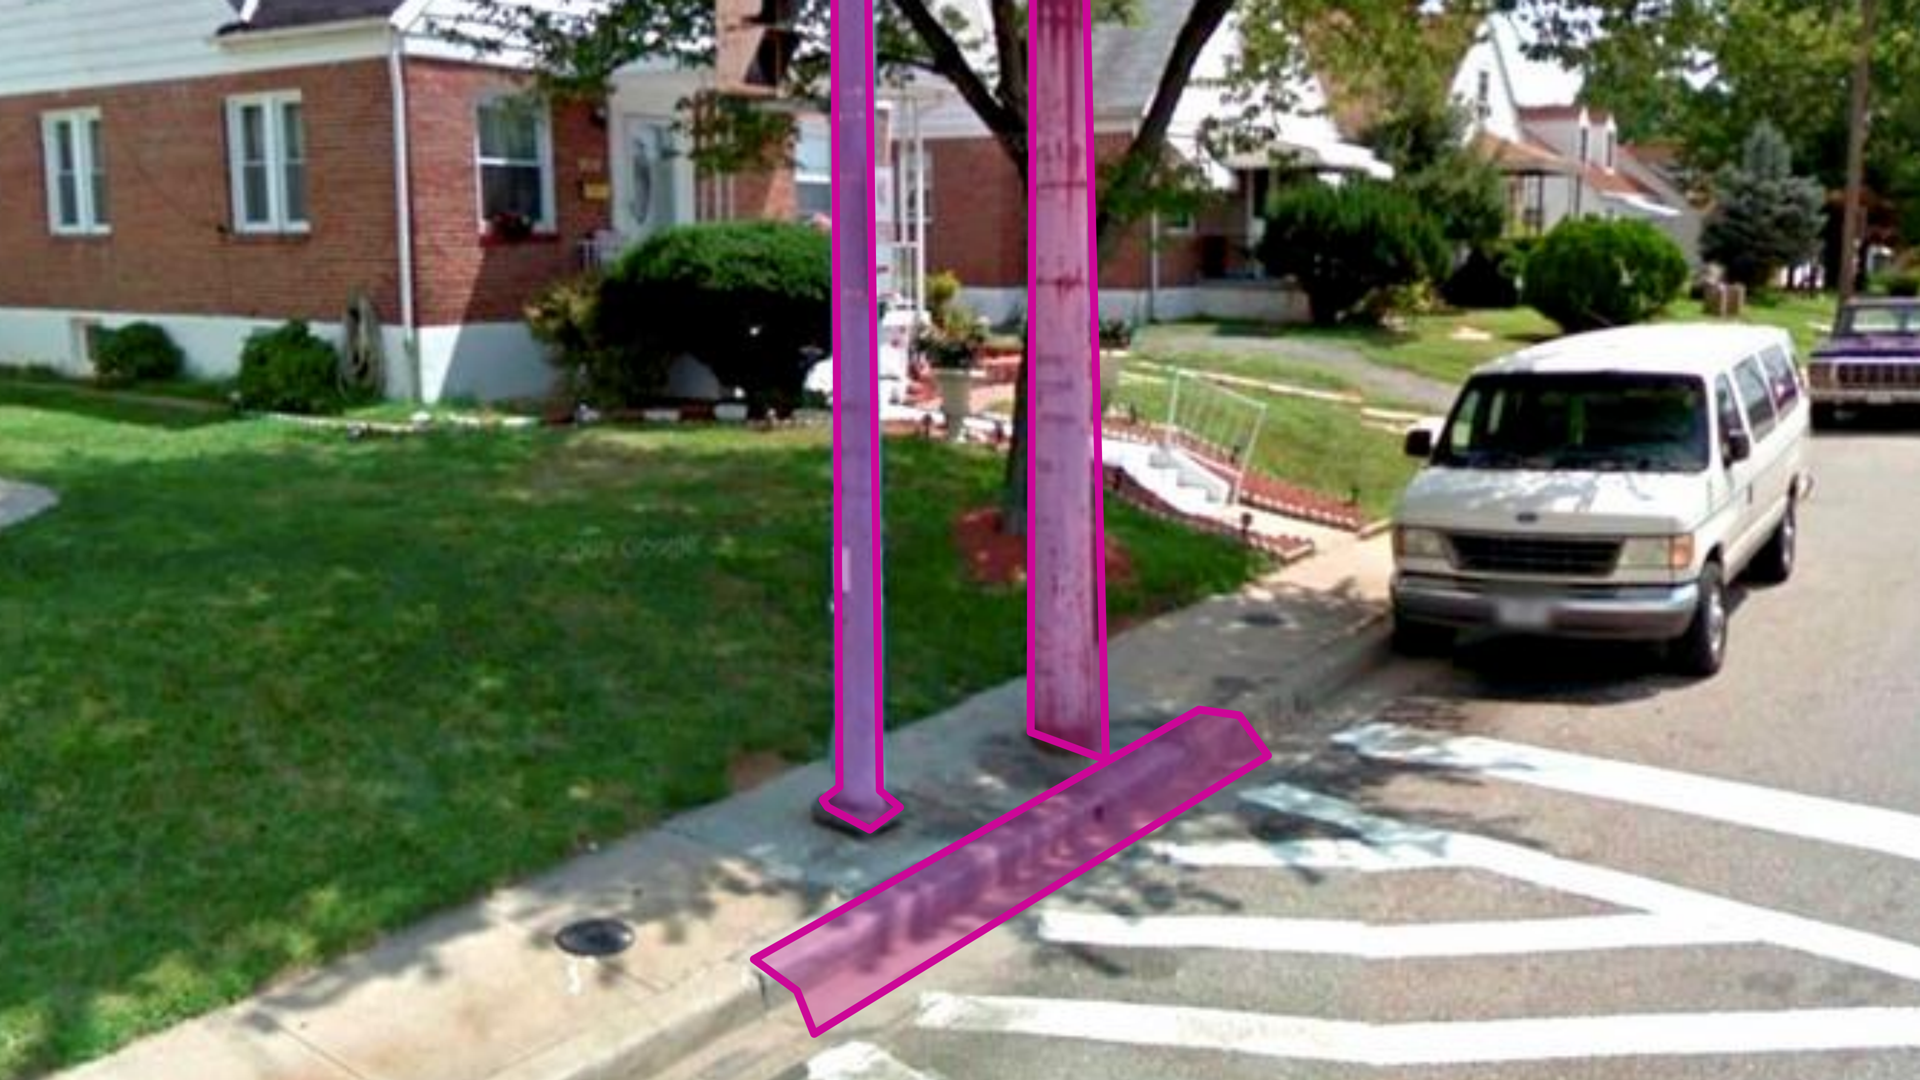

Labeling Interface

Labeling Interface

Crowdsourced Data Accuracy

We could collect street-level accessibility data from static Google Street View using crowdsourcing with **81% accuracy** and this figure went up to **93% with majority voting**

Hara K., Le V., Froehlich J.E. [ASSETS 2012]; Hara K., Le V., Froehlich J.E. [CHI 2013]

Washington, D.C.

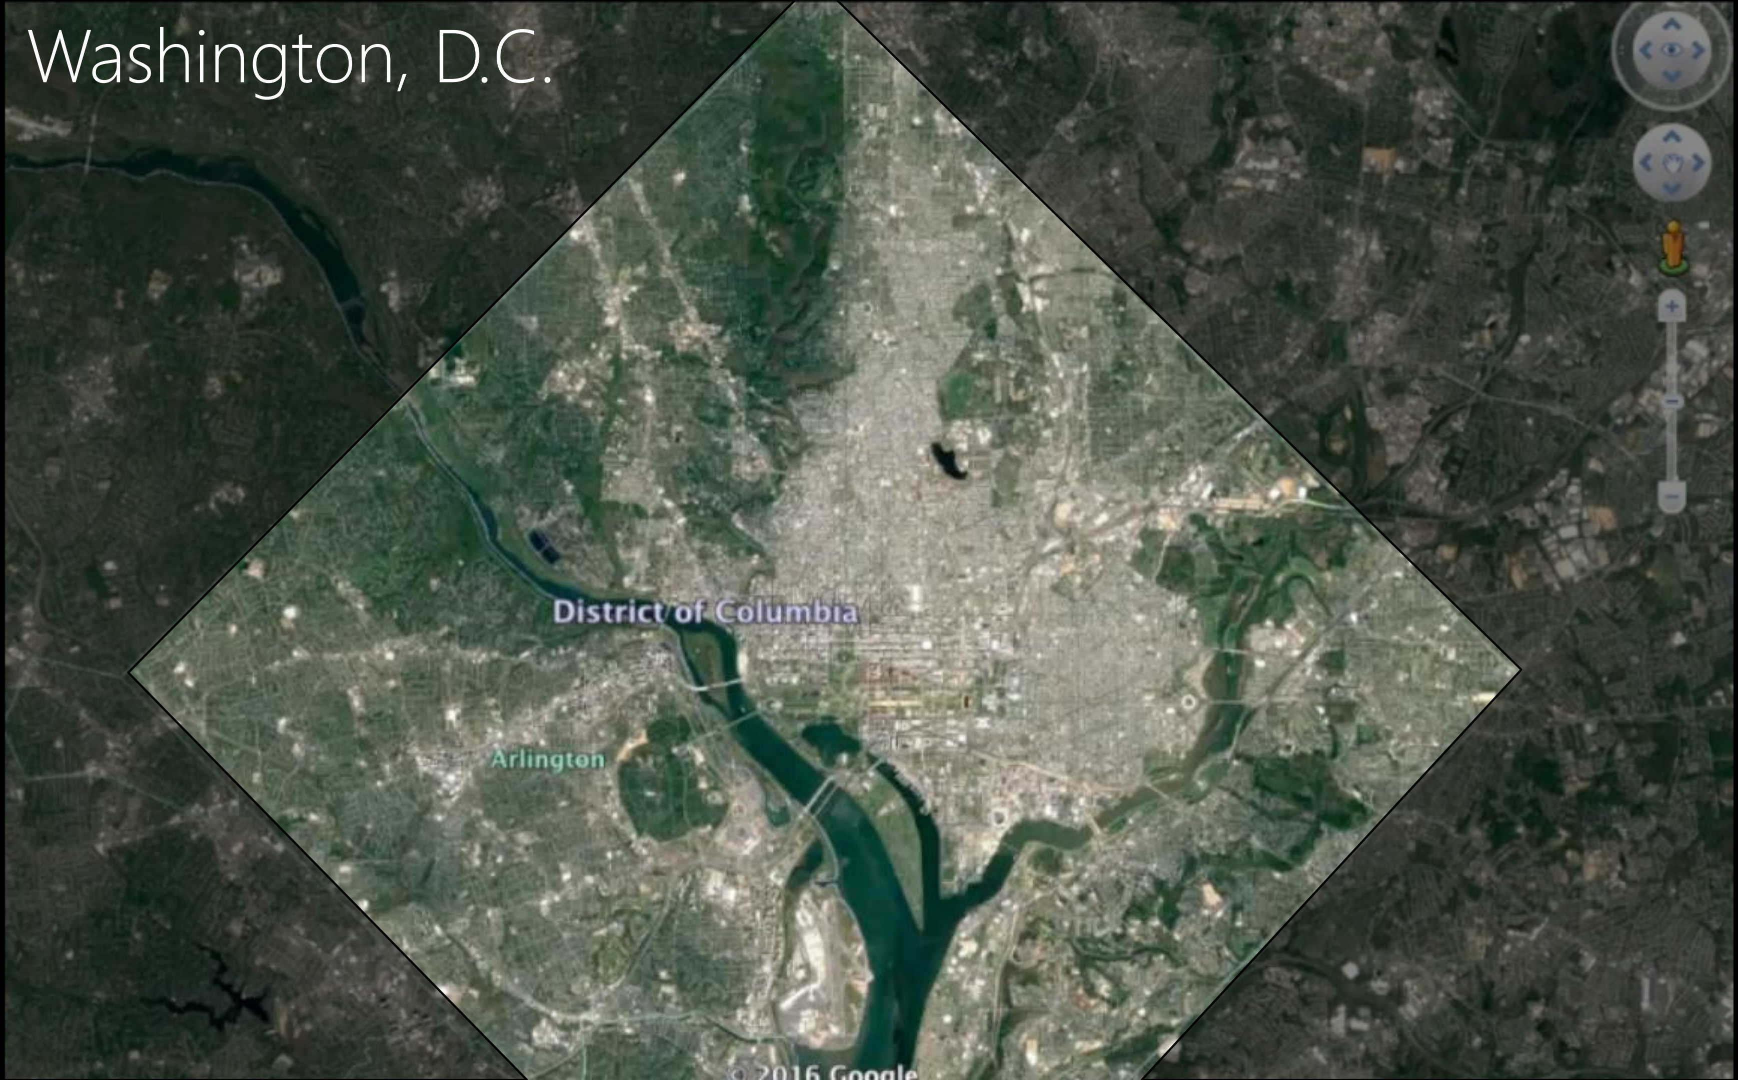

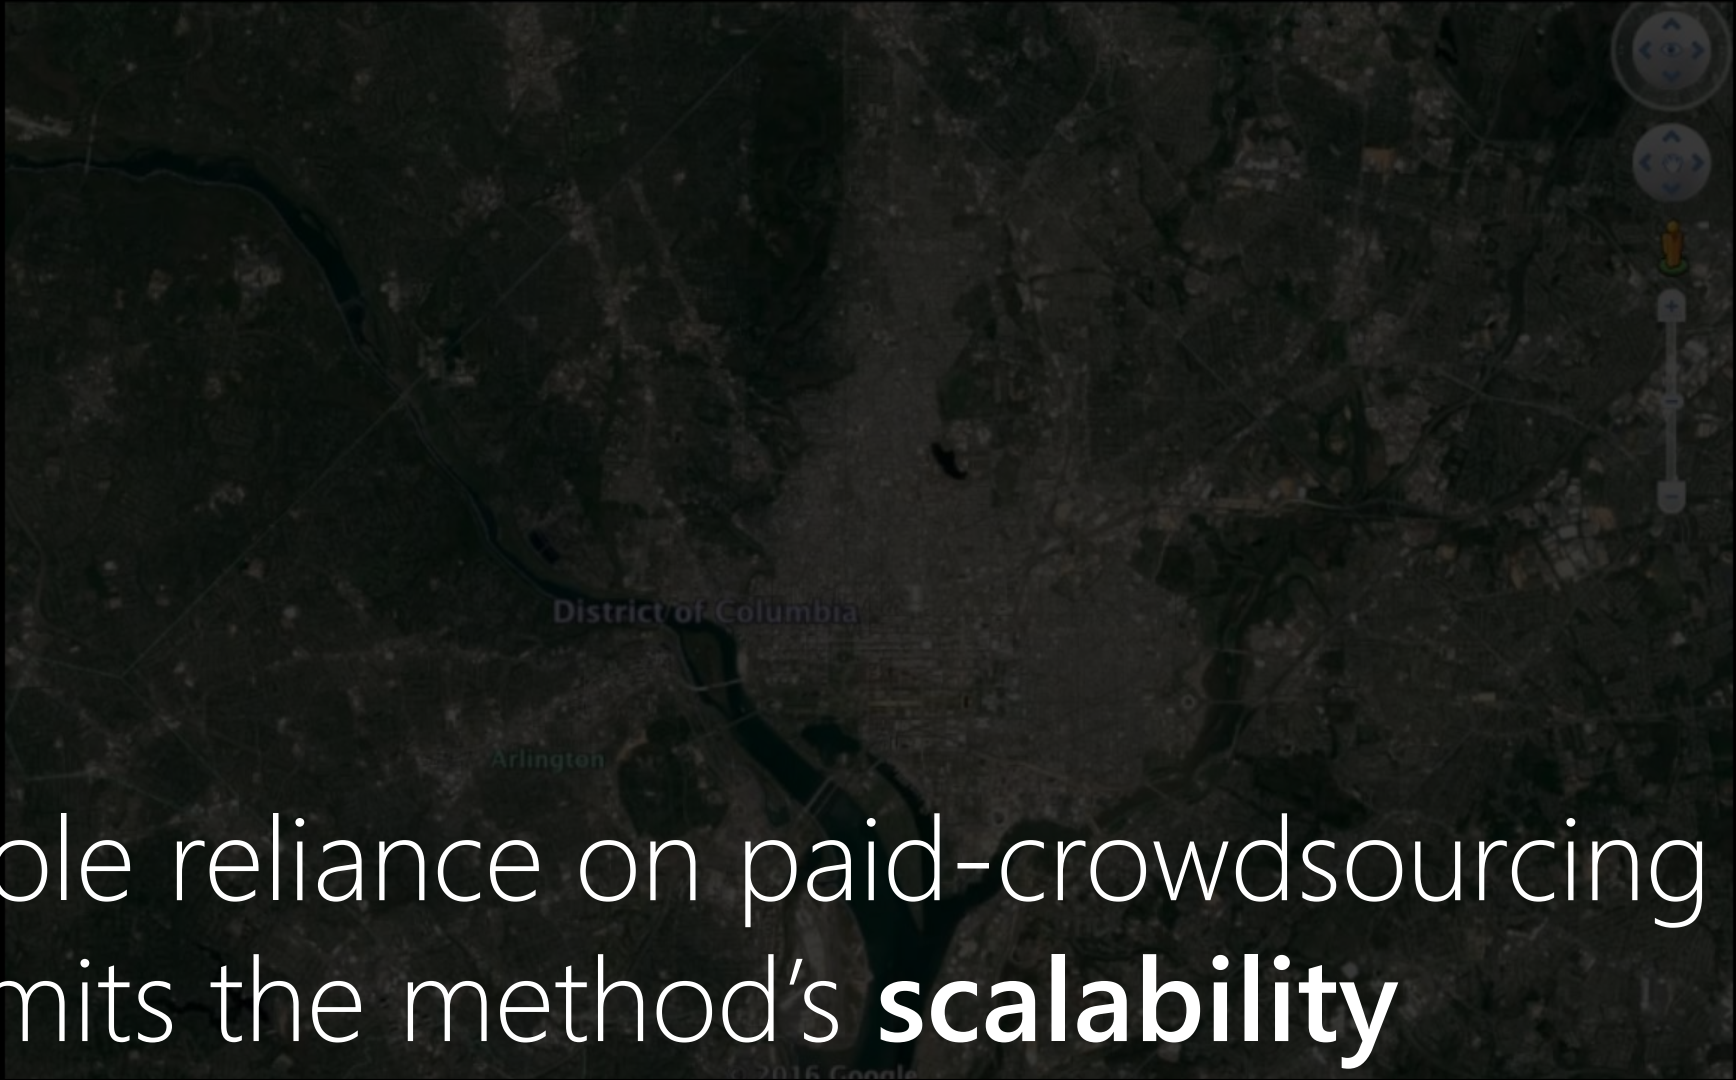

Sole reliance on paid-crowdsourcing limits the method's **scalability**

Scalable Data Collection Methods

Volunteered
Data Collection

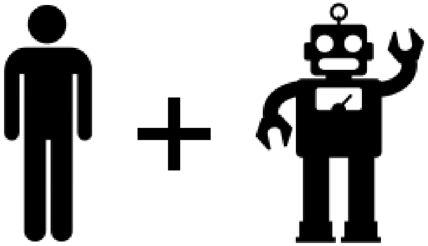

Semi-automated
Data Collection

Related Work: Quantifying Neighborhood Environment Using Computer Vision

[Arietta, S. *et al*, 2014, Naik, N. *et al*, 2014, Quercia, D. *et al*, 2014]

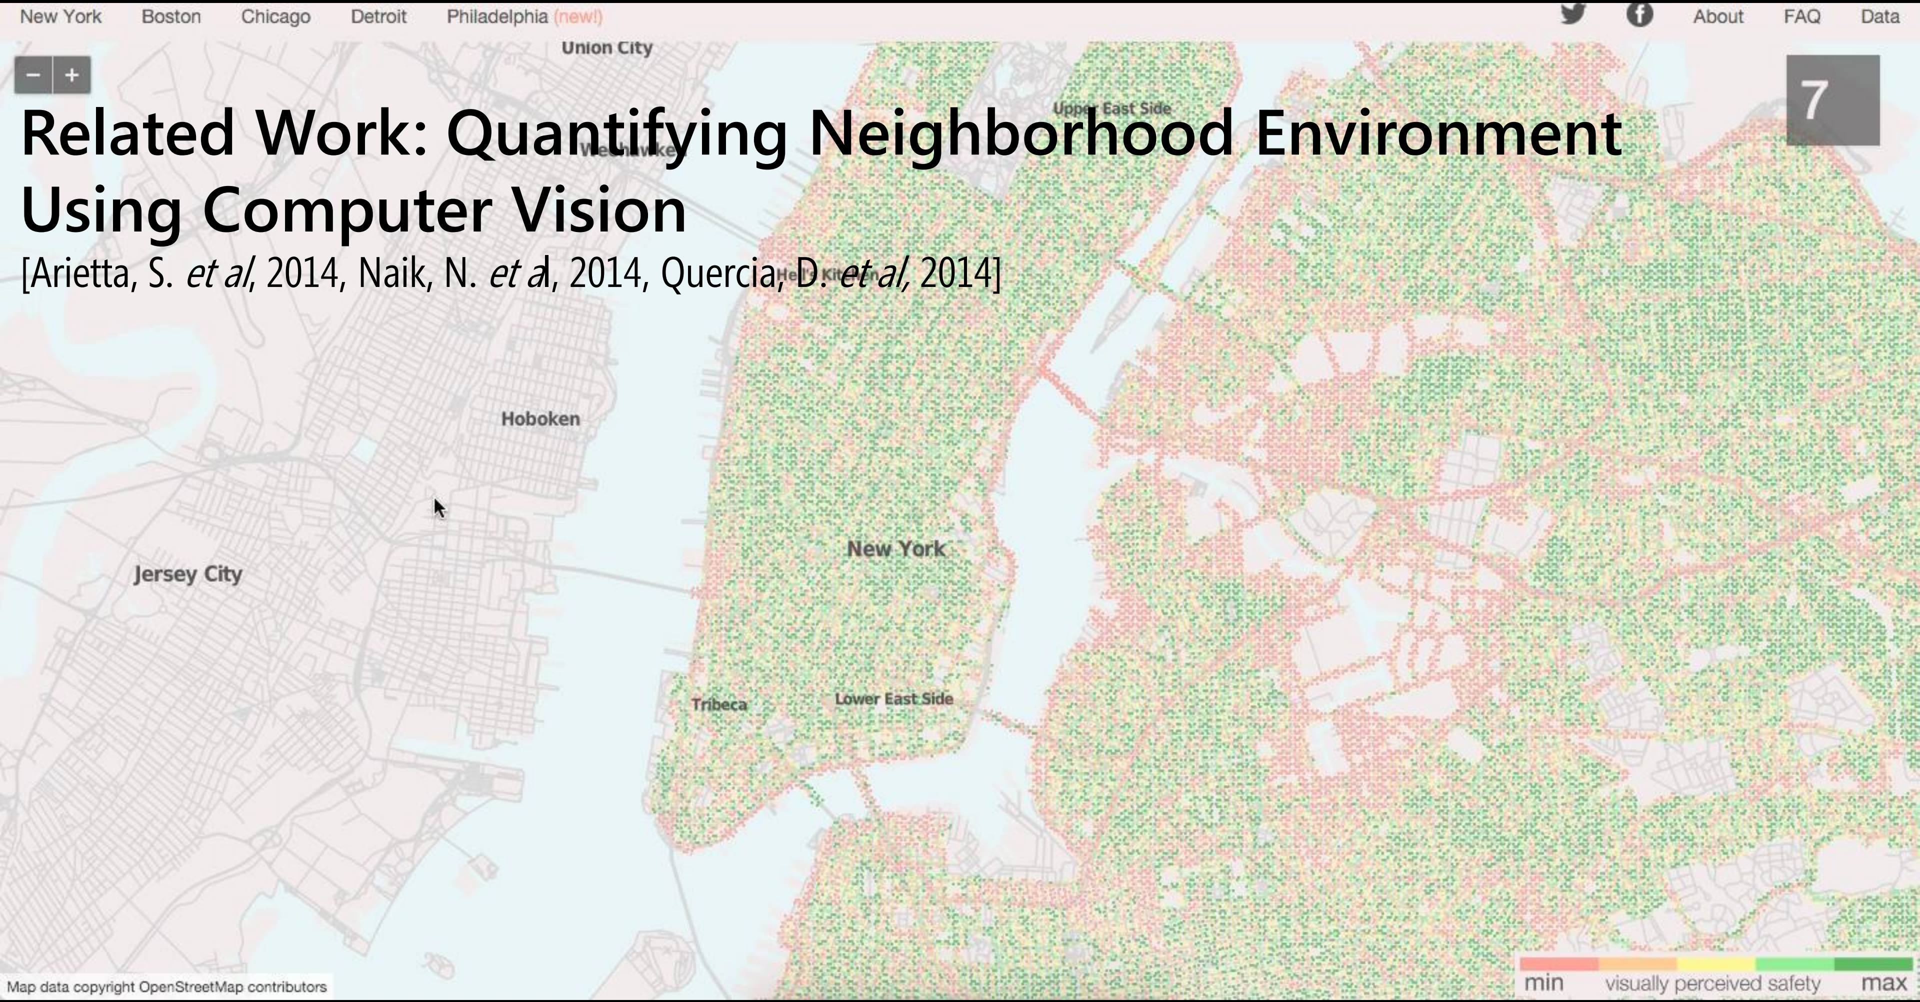

Street Score by Naik, N et al.

Related Work: Quantifying Neighborhood Environment Using Computer Vision

[Arietta, S. *et al*, 2014, Naik, N. *et al*, 2014, Quercia, D. *et al*, 2014]

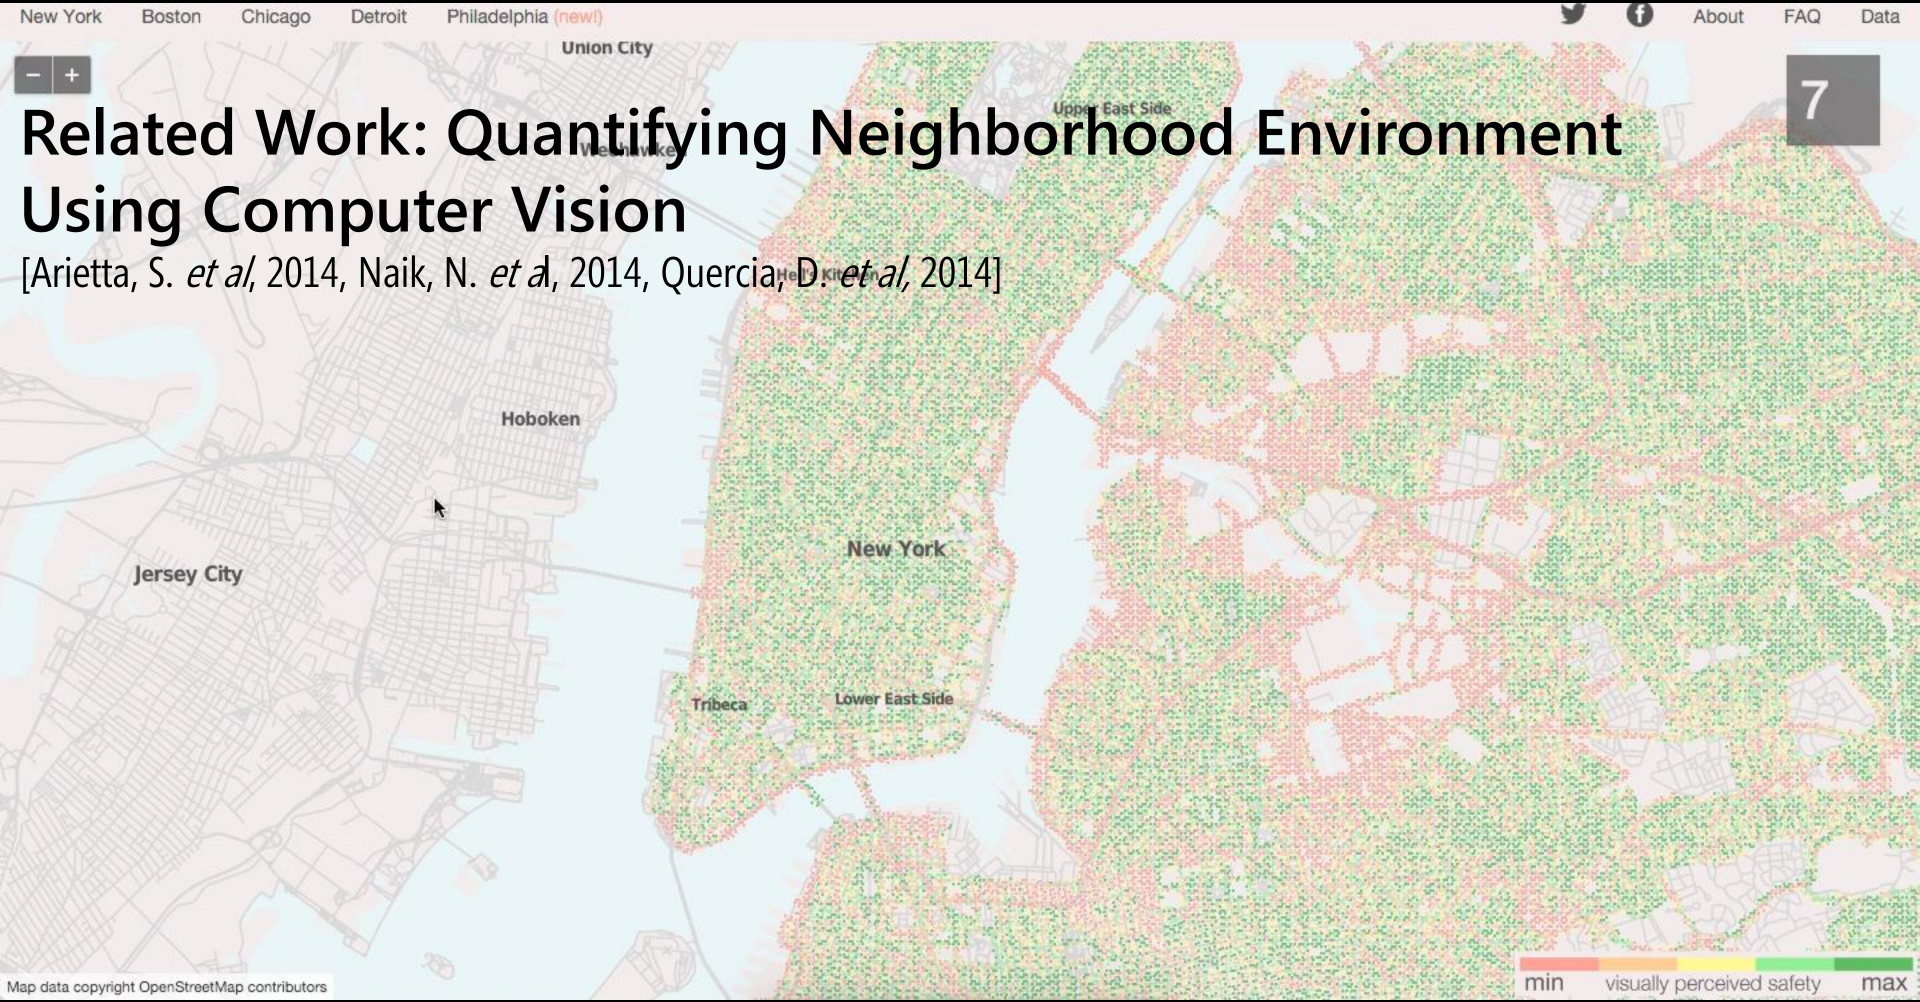

Street Score by Naik, N et al.

Related Work: Quantifying Neighborhood Environment Using Computer Vision

[Arietta, S. *et al*, 2014, Naik, N. *et al*, 2014, Quercia, D. *et al*, 2014]

We need more granular information to understand **which sidewalks are accessible/inaccessible** and **why**

Street Score by Naik, N et al.

Related Work: Object Detection with Human-in-the-Loop

Branson *et al.* 2010, Quinn *et al*. 2010; Su *et al.* 2012

CrowdFlow: Integrating Mac
Mechanical Turk for Speed-C

Alexander J. Quinn[1], Benjamin B. Bedere
University of Mar
Human-Computer Inte
Department of Computer Science[1] ll iSchool[2] ll In
College Park, MD 2
aq@cs.umd.edu, bederson@cs.umd.edu, tomy

Abstract

Humans and machines have competing strengths for tasks such as natural language processing and image understanding. Whereas humans do these things naturally with potentially high accuracy, machines offer greater speed and flexibility. CrowdFlow is our toolkit for a model for blending the two in order to attain tighter control over the inherent tradeoffs in speed, cost and quality. With CrowdFlow, humans and machines work together to do a set of tasks at a user-specified point in the tradeoff space. They work symbiotically, with the humans providing training data to the machine while the machine provides first cut results to the humans to save effort in cases where the machine's answer was already correct. The CrowdFlow toolkit can be considered as a generalization of our other domain-specific efforts aimed at enabling cloud computing services using a variety of computational resources to achieve various tradeoff points.

1.Introduction

There is a large set of problems that can be solved either b human computation or machine learning. These inclu recognizing the faces of missing children in surveillance vide translating documents between languages, or summarizing opinions of blogs relating to a particular topic, and many oth from the realms of natural language processing (NLP).

Generally, humans can solve these problems with hi accuracy than machines alone could do, though human e tend to be costly and time-consuming. Online labor m Amazon Mechanical Turk (AMT) [13] give

Visual Recognition with Humans i

Steve Branson[1], Catherine Wah[1], Florian Schroff[1], Bor
Welinder[2], Pietro Perona[2], and Serge Bel
[1] University of California, San Diego
{sbranson,cwah,gschroff,bbabenko,sjb}@cs.
[2] California Institute of Technology
{welinder,perona}@caltech.edu

Abstract. We present an interactive, hybrid human- for object classification. The method applies to classes recognizable by people with appropriate expertise (*e.g.* airplane model), but not (in general) by people withou can be seen as a visual version of the *20 questions gam* based on simple visual attributes are posed interactiv identify the true class while minimizing the number using the visual content of the image. We introduce a for incorporating almost any off-the-shelf multi-class algorithm into the visual 20 questions game, and pro to account for imperfect user responses and unrelial algorithms. We evaluate our methods on *Birds-200* of 200 tightly-related bird species, and on the *Anim* dataset. Our results demonstrate that incorporating recognition accuracy to levels that are good enough cations, while at the same time, computer vision re human interaction required.

Crowdsourcing Annotations for Visual Object Detection

Hao Su, Jia Deng, Li Fei-Fei
Computer Science Department, Stanford University

Abstract

A large number of images with ground truth object bounding boxes are critical for learning object detectors, which is a fundamental task in compute vision. In this paper, we study strategies to crowd-source bounding box annotations. The core challenge of building such a system is to effectively control the data quality with minimal cost. Our key observation is that drawing a bounding box is significantly more difficult and time consuming than giving answers to multiple choice questions. Thus quality control through additional verification tasks is more cost effective than consensus based algorithms. In particular, we present a system that consists of three simple sub-tasks — a drawing task, a quality verification task and a coverage verification task. Experimental results demonstrate that our system is scalable, accurate, and cost-effective.

1 Introduction

Object detection is one of the fundamental tasks of visual recognition. Given an input image, an object detector outputs a bounding box wherever an object of interest exists. To learn a good detector, it is necessary to have a large number of training images with ground truth annotations in the form of bounding boxes, i.e. tight rectangles around the object of interest. Indeed, state of the art detection systems (Viola and Jones 2004; Felzenszwalb et al. 2010) have relied on accurate bounding box annotations. Although it is possible to use weaker supervision, e.g. binary labels of object presence, it substantially increases the difficulty of learning.

In this paper, we study strategies to crowd-source bounding box annotations. Our goal is to build a system that is fully automated, highly accurate, and cost-effective. Given a collection of images where the object of interest has been verified to exist, for each image the system collects a tight bounding box for every instance of the object. Specifically, we have the following two requirements.

- Quality. Each bounding box n

Figure 1: An example of bounding box annotations for the "bottle" category.

- Coverage. Every object instance needs to have a bounding box. This is important for detection because it tells the learning algorithms with certainty what is *not* the object.

Figure 1 shows examples of bounding box annotations that meet both the quality and coverage requirements.

The core challenge of building such a system is how to achieve both high quality and complete coverage in a cost-effective way, i.e. minimizing cost while guaranteeing quality. A basic quality control strategy is majority voting—collecting answers from multiple human subjects and taking the consensus. This approach has been successfully applied to image annotation tasks such as verifying the presence of objects or attributes (Deng et al. 2009; Sorokin and Forsyth 2008). However, drawing bounding box is significantly more time consuming tha

Related Work: Object Detection with Human-in-the-Loop

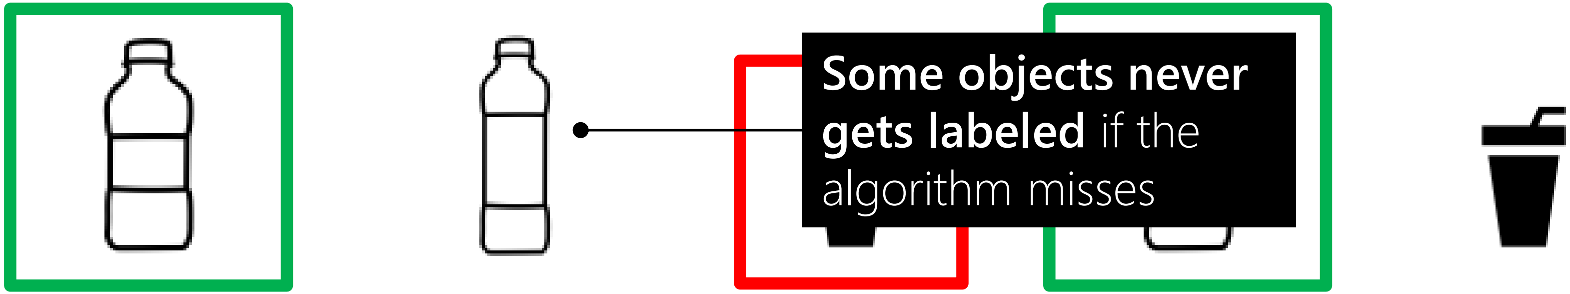

Task: Detect all the bottles

Step 1: An object detection algorithm detects bottles

Step 2: Humans verify the object detection outputs

More accurate compared to computer vision alone and **cheaper than human labeling**

Adaptive Workflow for Optimizing Efficiency

Varying the number of workers to recruit depending on task difficulty
[Kamar *et al.* 2012; Welinder and Perona 2010]

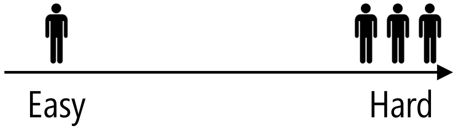

Assigning stronger workers to harder tasks
[Dai *et al.* 2011]

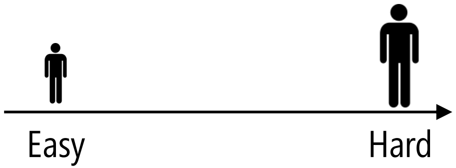

Reducing the tasks that require human work
[Deng *et al.* 2014; Jain, Grauman, and Betke 2016]

Task ~~Task~~ Task Task ~~Task~~

Changing task interface based on worker characteristics
[Jain and Grauman 2013, Lin *et al.* 2012, Russakovsky *et al.* 2015]

Interface 1 or Interface 2

Tohme

遠目·*Remote Eye*

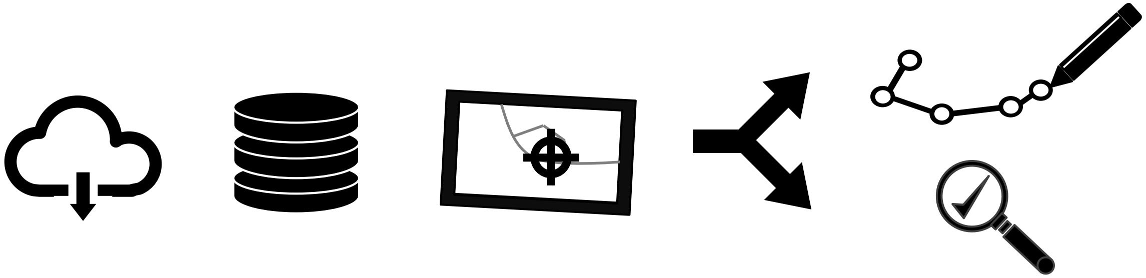

Hara K., Sun J., Chazan J., Jacobs D., Froehlich J.E. [HCOMP 2013]; Hara K., Sun J., Moore R., Jacobs D., Froehlich J.E. [UIST 2014]

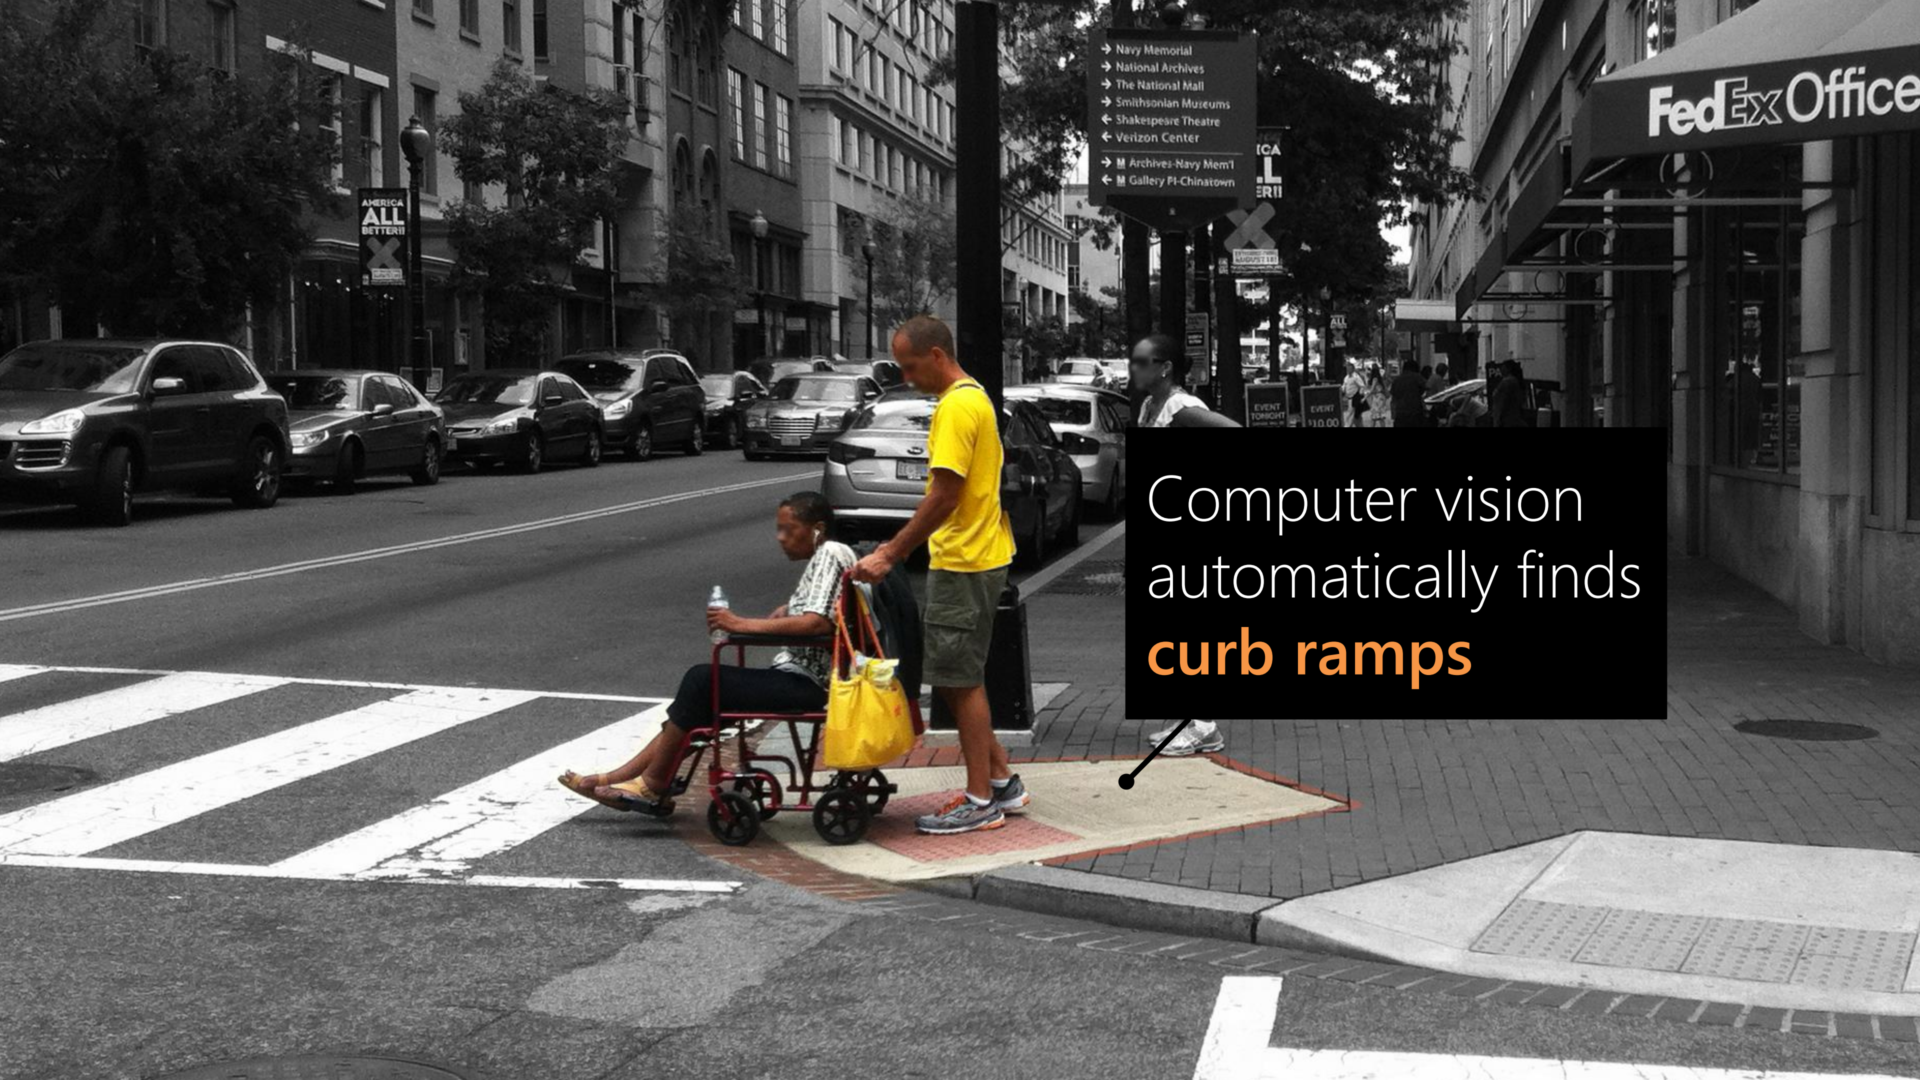

Navy Memorial
National Archives
The National Mall
Smithsonian Museums
Shakespeare Theatre
Verizon Center
M Archives-Navy Mem'l
M Gallery Pl-Chinatown
AMERICA ALL BETTER!!
FedEx Office
Computer vision automatically finds **curb ramps**

Without curb cuts, **people with ambulatory disabilities simply cannot navigate the city.**

Kinney v. Yerusalim, 1993

3rd Circuit, Court of Appeals

Curb Ramps are Visually Salient

Tohme
遠目·*Remote Eye*

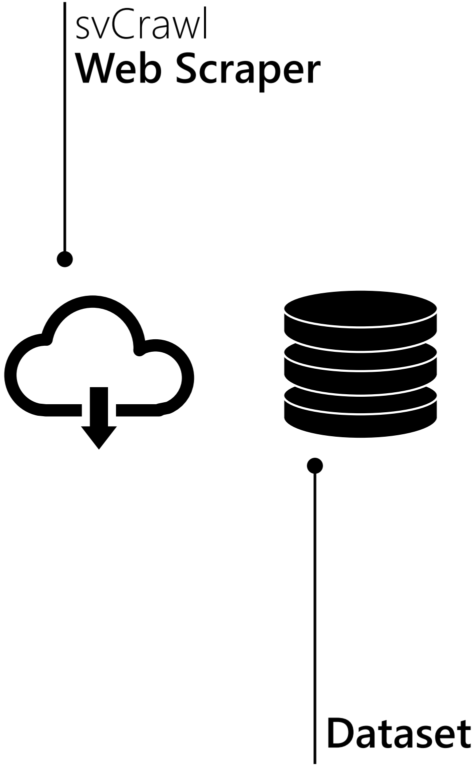

Tohme
遠目· *Remote Eye*

svCrawl
Web Scraper

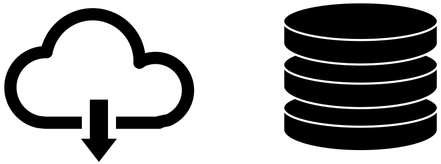

3D Depth Map
GIS Metadata (*e.g.*, topological data)
Top down map images
Street View images
Dataset

Tohme
遠目·*Remote Eye*

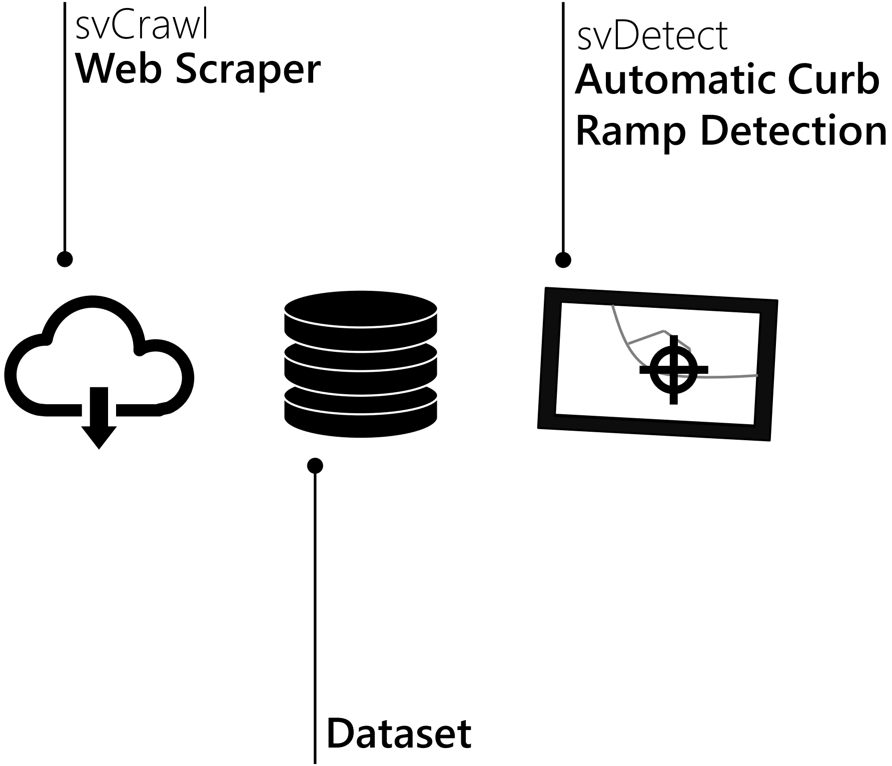

Tohme
遠目·*Remote Eye*

svCrawl
Web Scraper

svDetect
Automatic Curb Ramp Detection

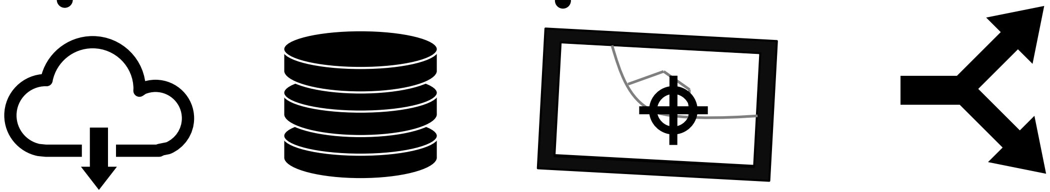

Did our computer vision algorithm perform well?

Dataset

svControl
Automatic Task Allocation

Tohme
遠目· *Remote Eye*

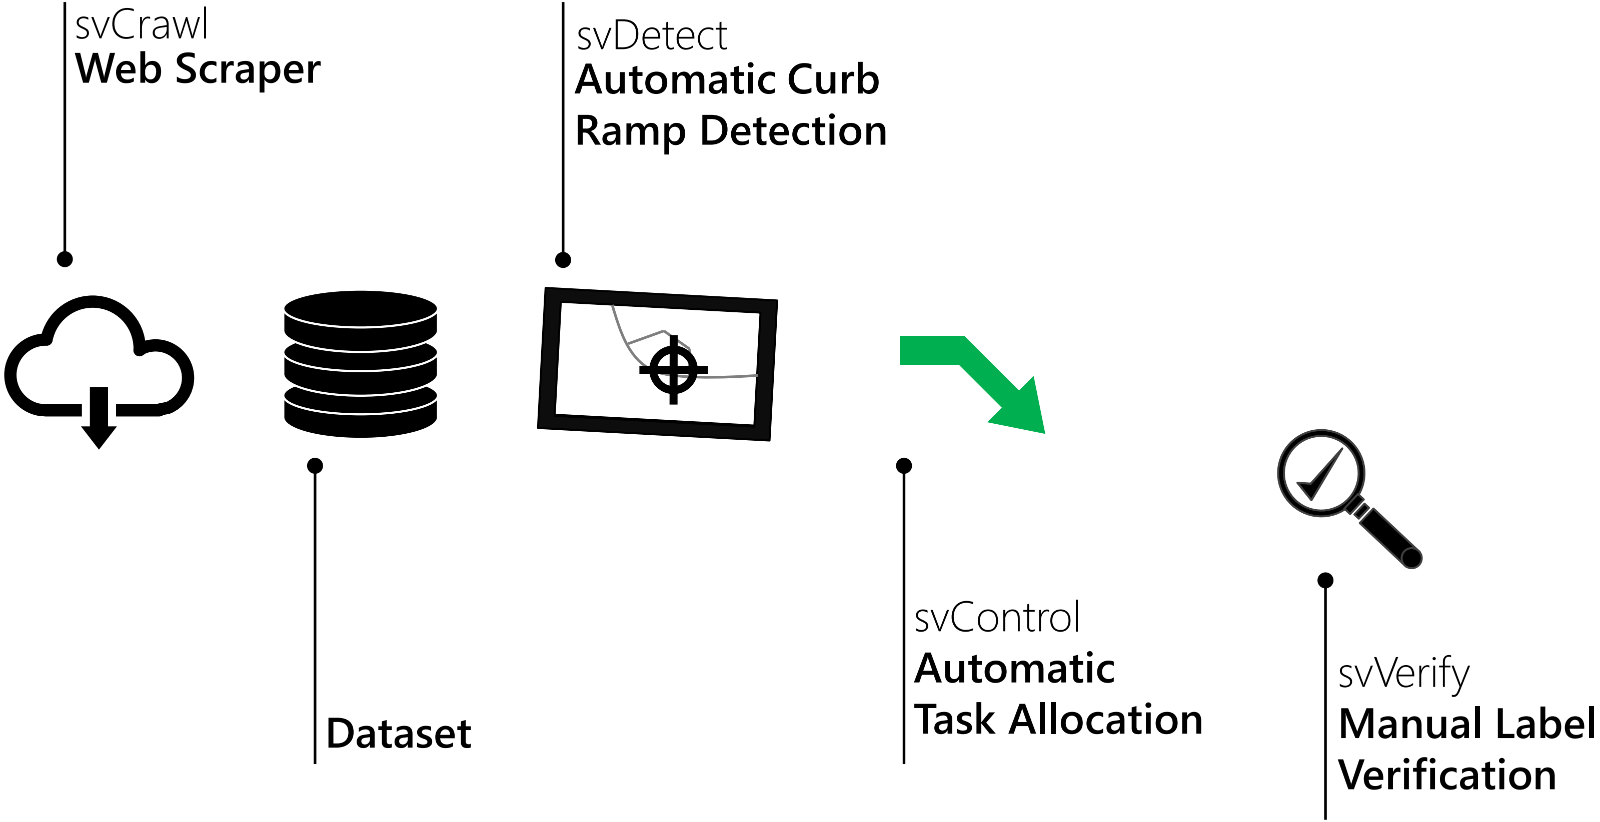

Tohme
遠目·Remote Eye

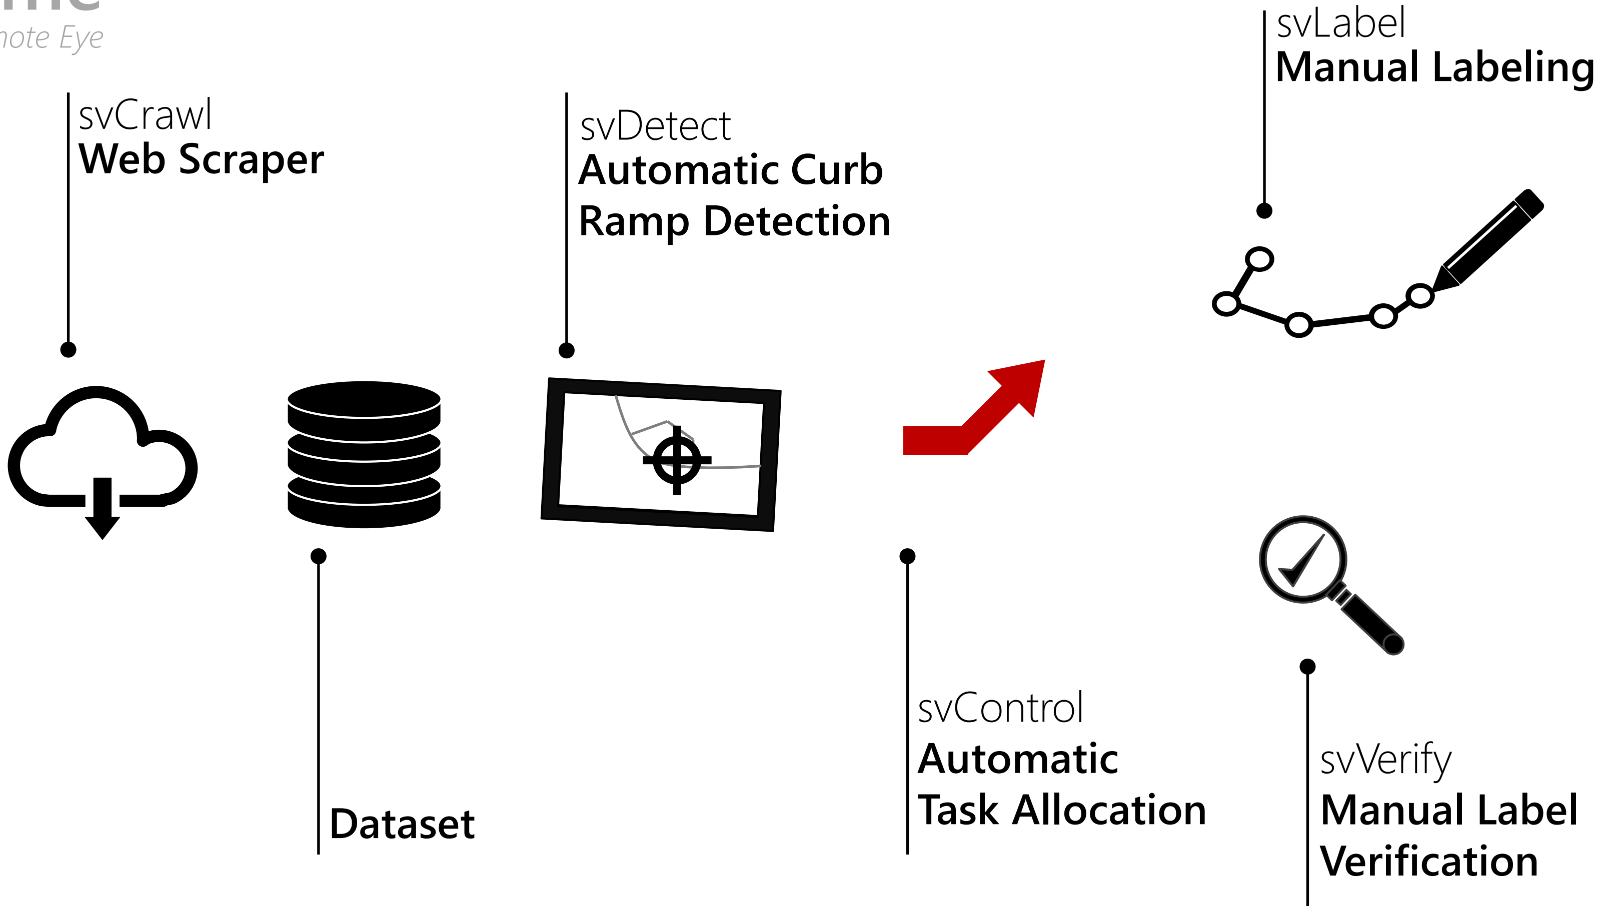

How do we define computer vision **failure?**

Correct detection

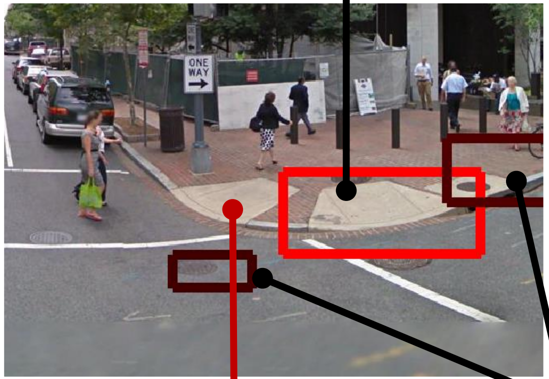

False positive detections

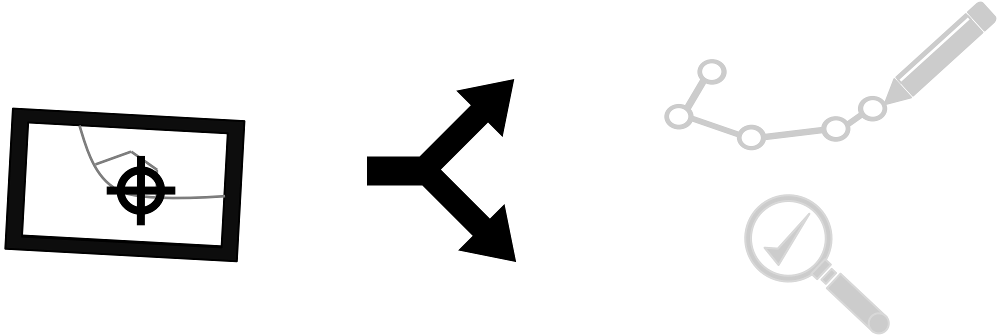

False Negative Error = Computer Vision Failure
Because asking humans to label missed curb ramps is much more expensive than asking to verify

Tohme
遠目·*Remote Eye*

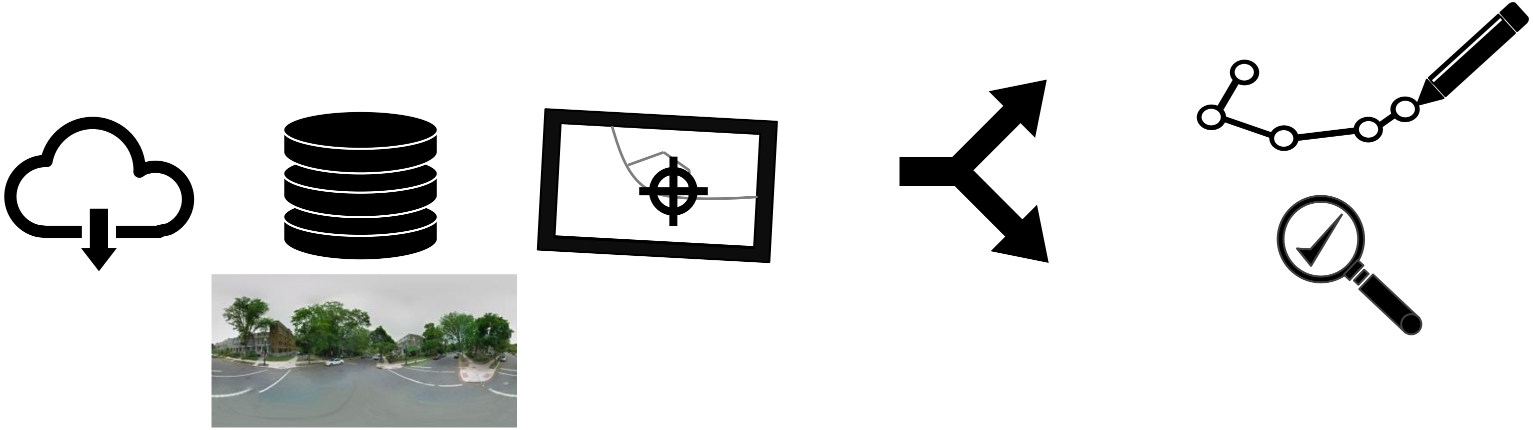

Tohme
遠目·*Remote Eye*

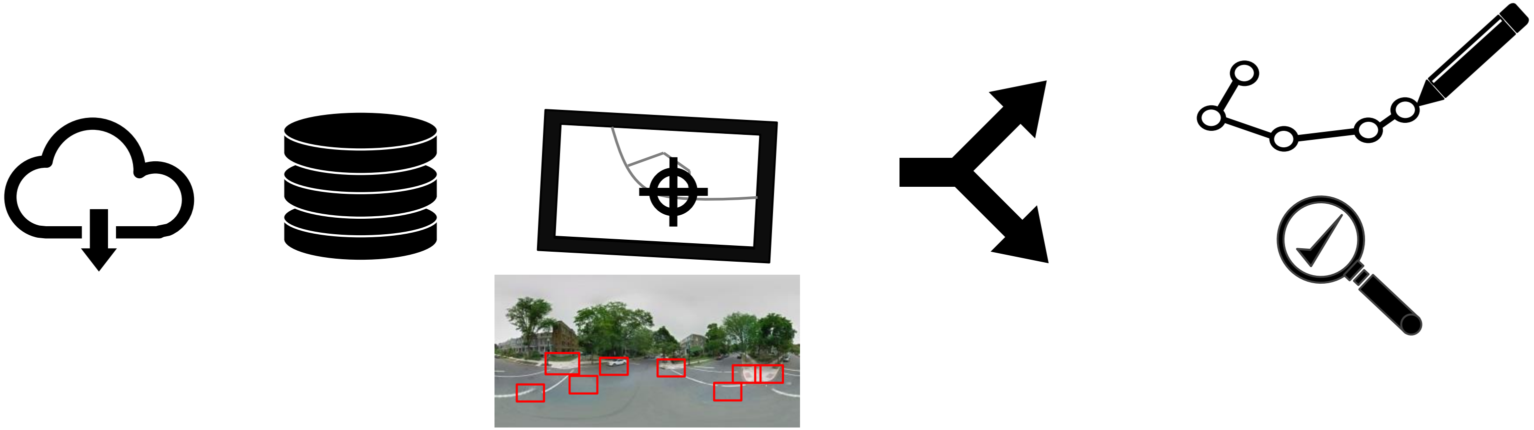

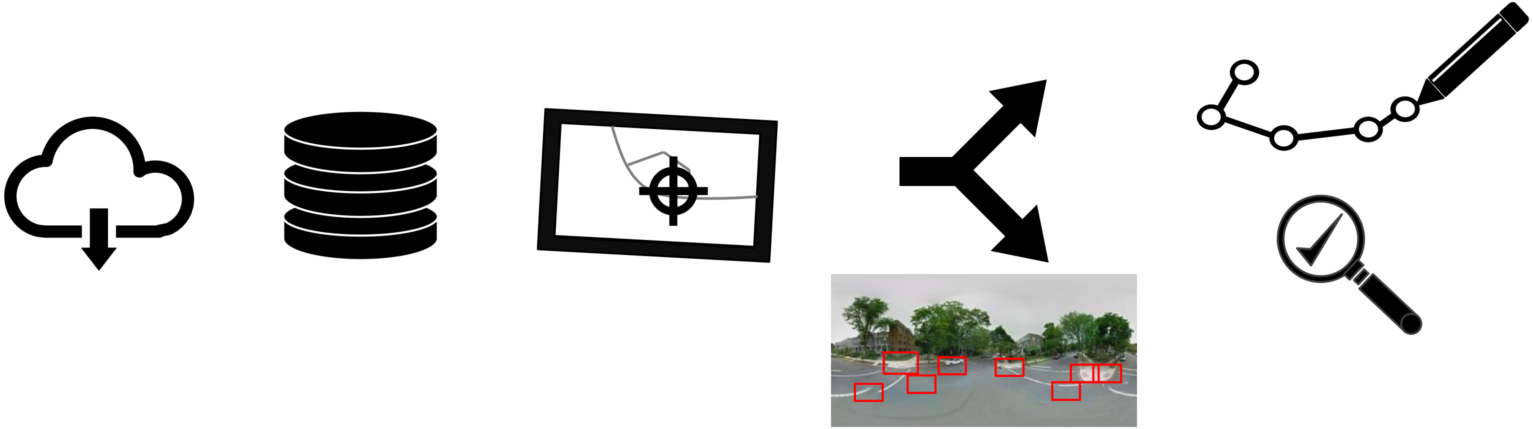

A feature vector used to in the workflow controller

Complexity:	0.14
Cardinality:	0.33
Depth:	0.21
CV:	0.22

Tohme
遠目·*Remote Eye*

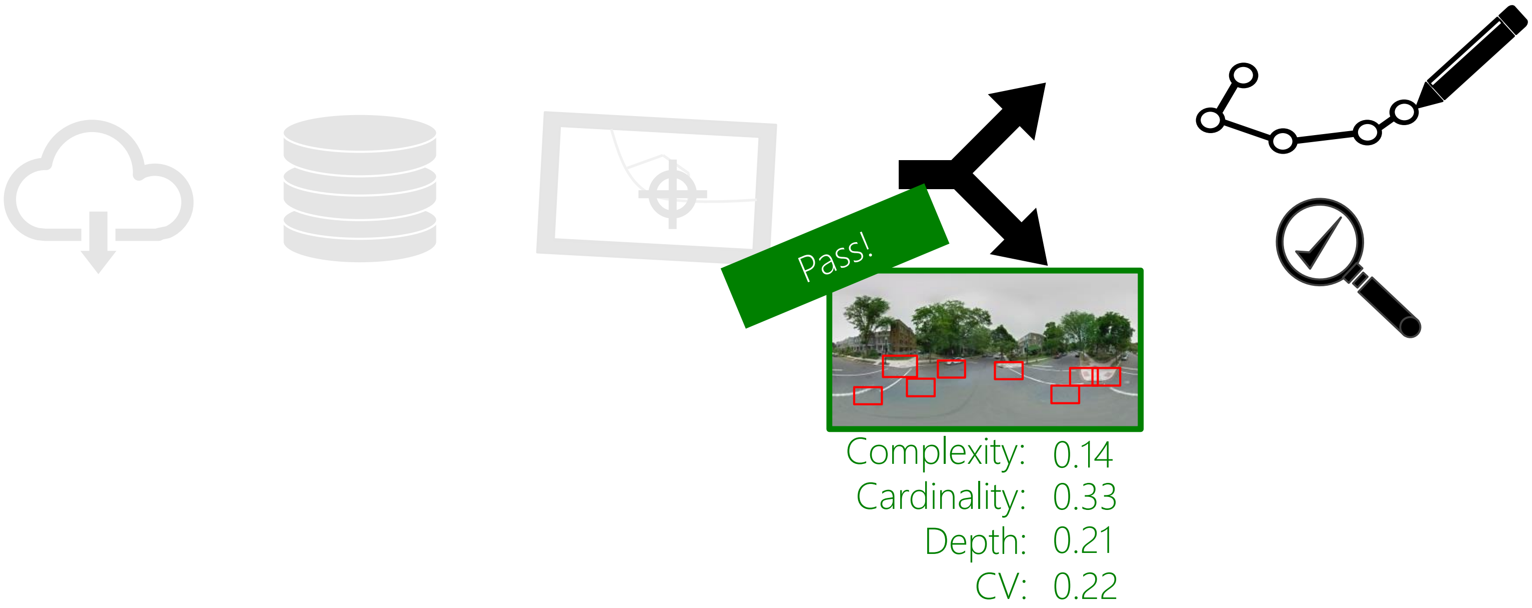

Tohme
遠目· *Remote Eye*

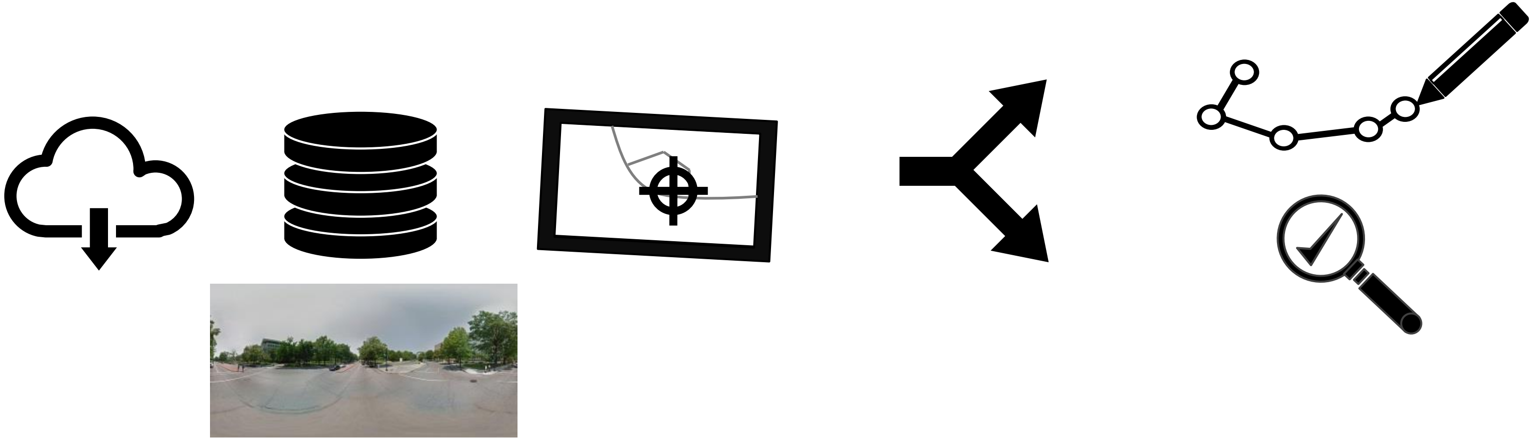

Tohme

遠目·*Remote Eye*

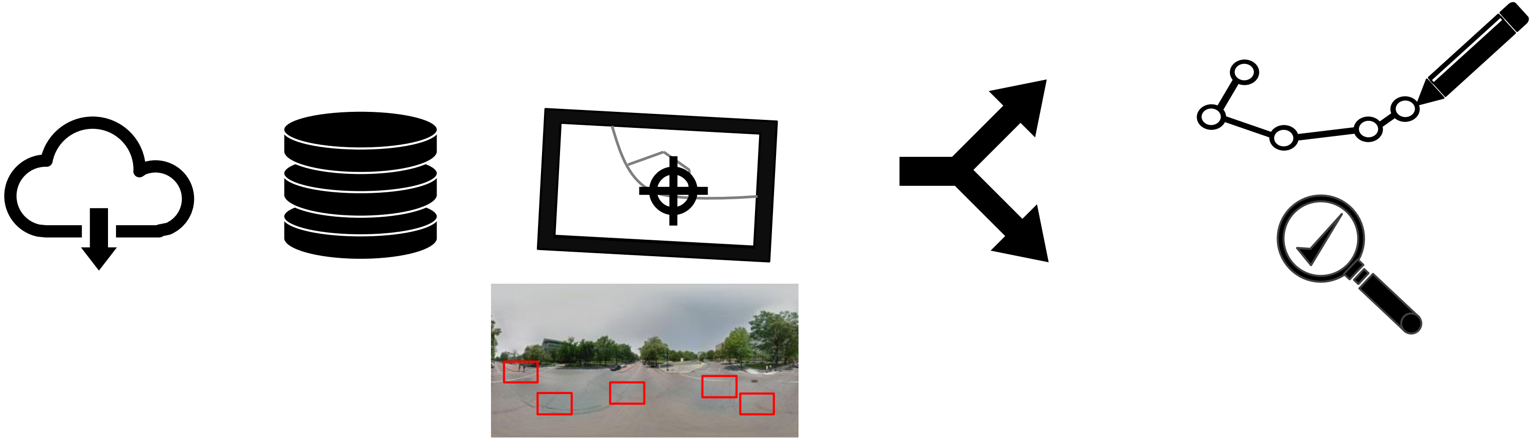

Tohme

遠目·*Remote Eye*

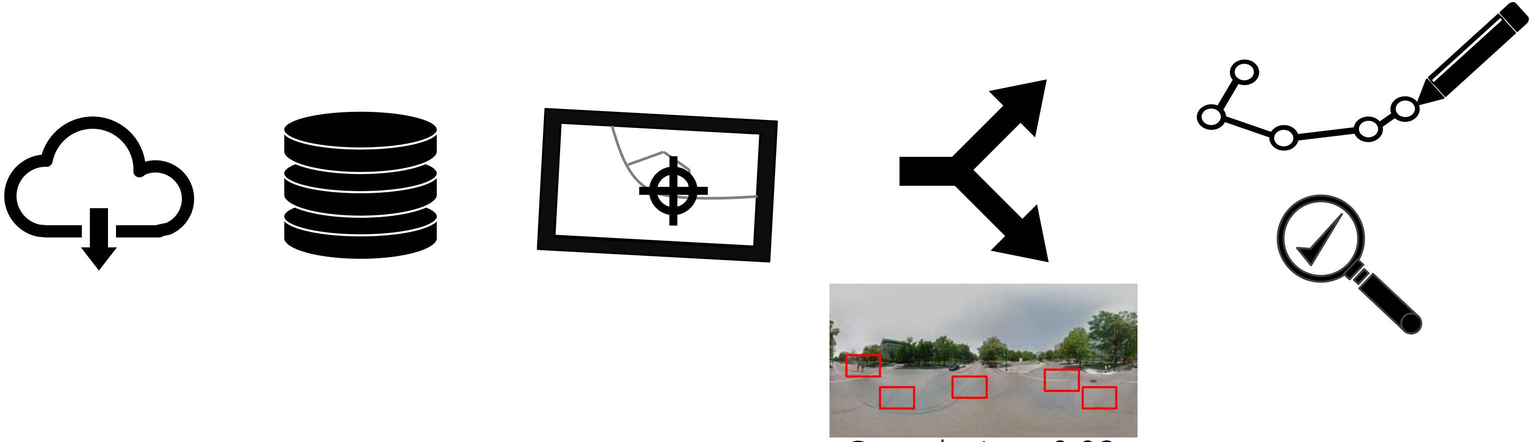

Complexity: 0.82
Cardinality: 0.25
Depth: 0.96
CV: 0.54

Tohme
遠目·*Remote Eye*

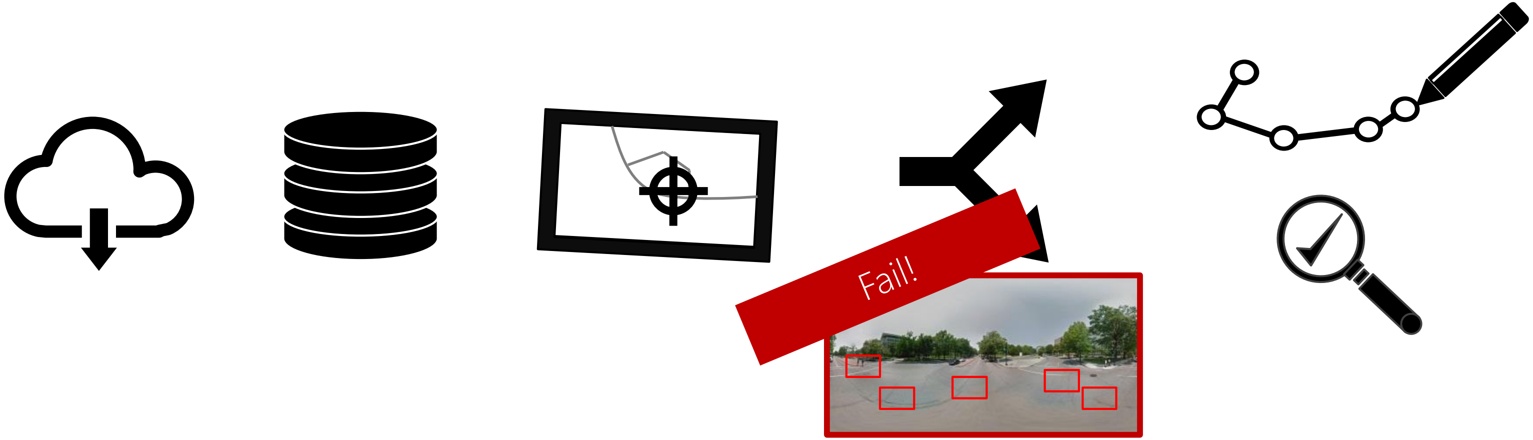

Complexity:	0.82
Cardinality:	0.25
Depth:	0.96
CV:	0.54

Tohme

遠目· *Remote Eye*

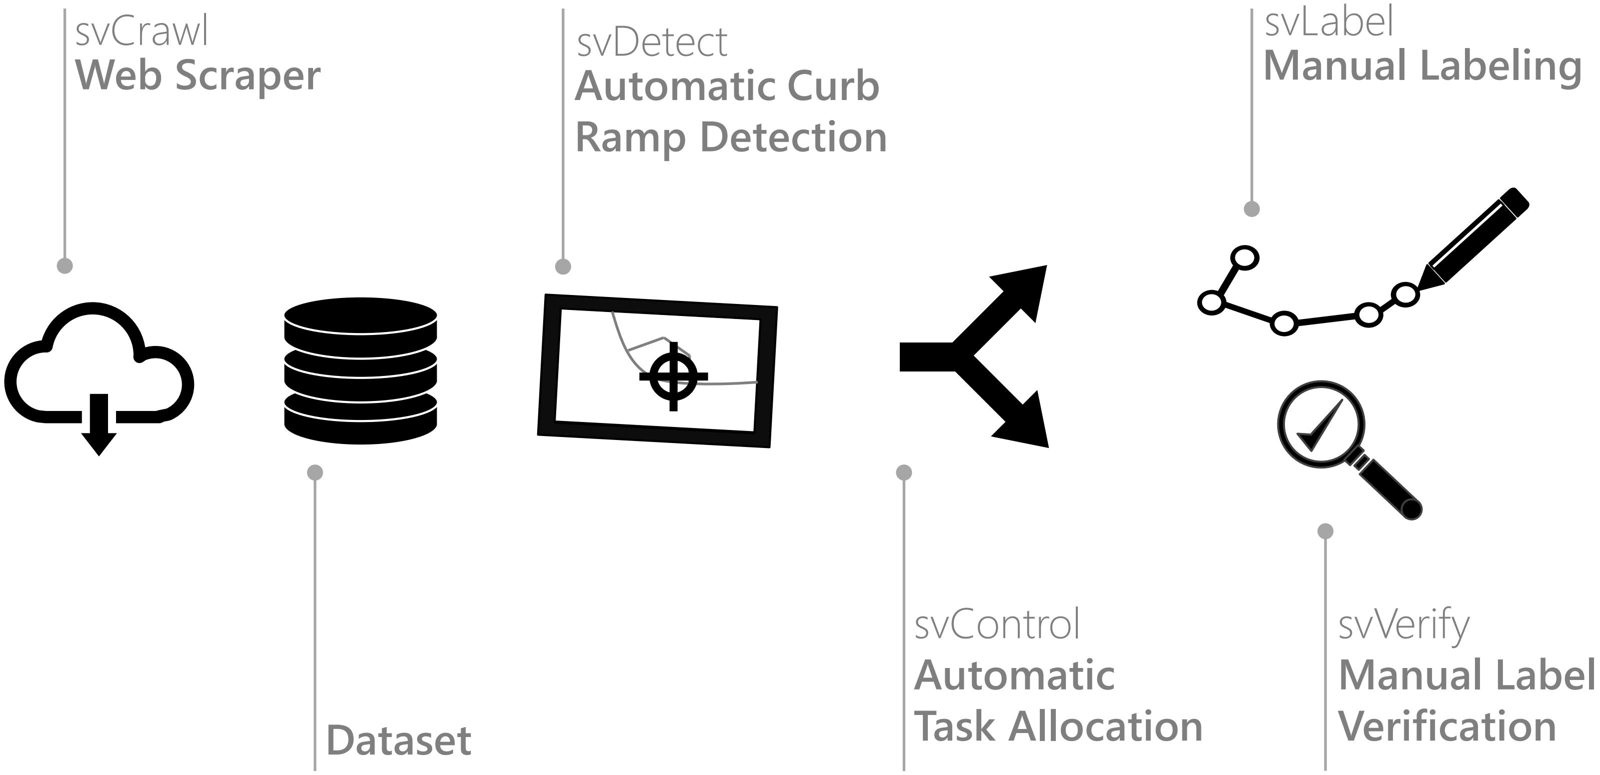

Tohme

遠目·*Remote Eye*

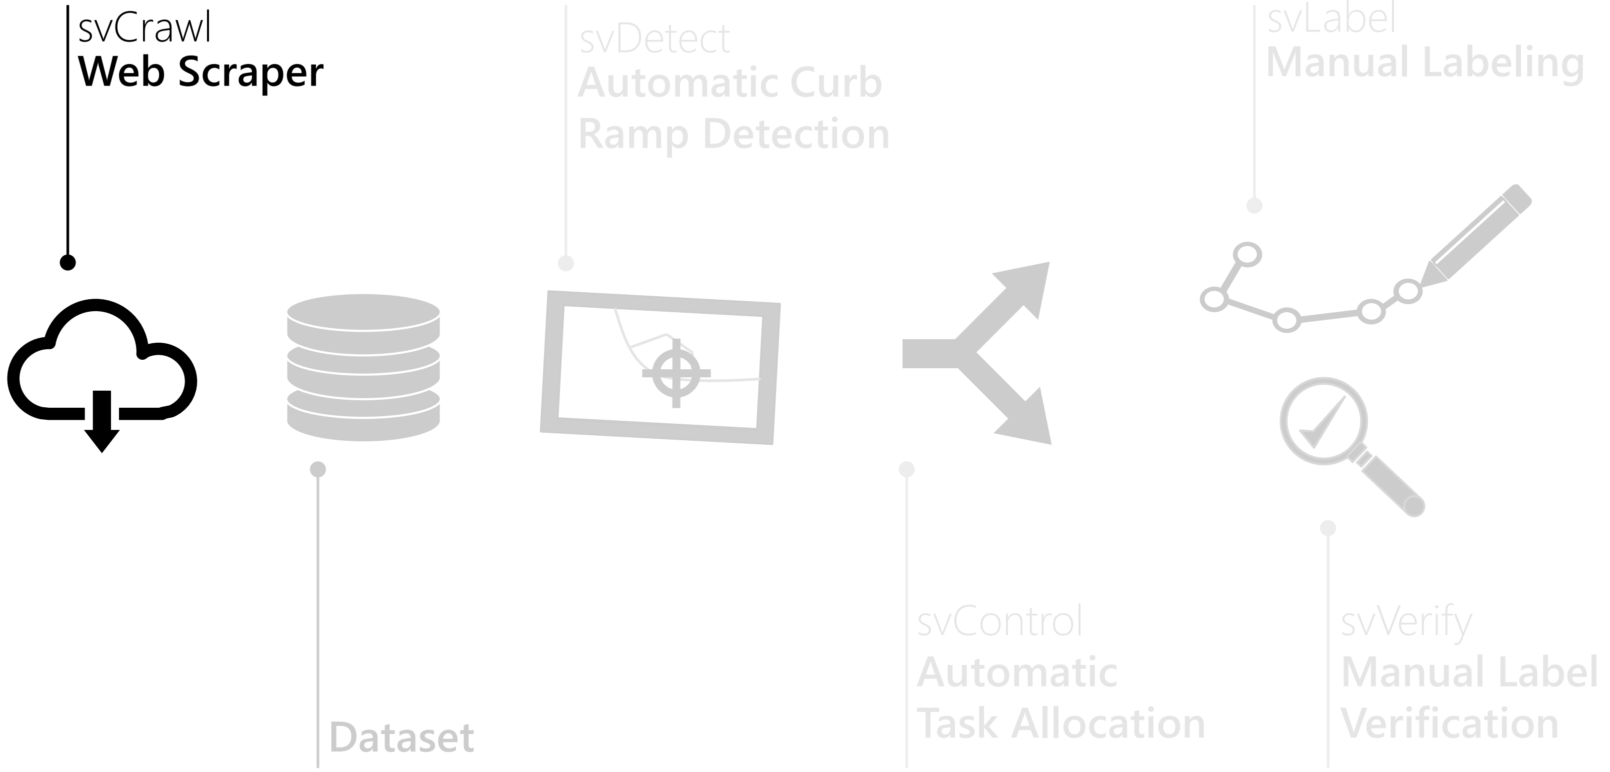

Scraper

Street View Images

We collected images from intersections because that's where we find curb ramps

Scraper

Street View Images

Point-cloud Data

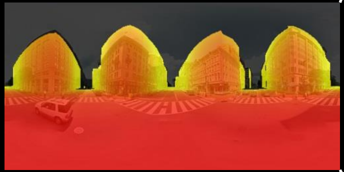

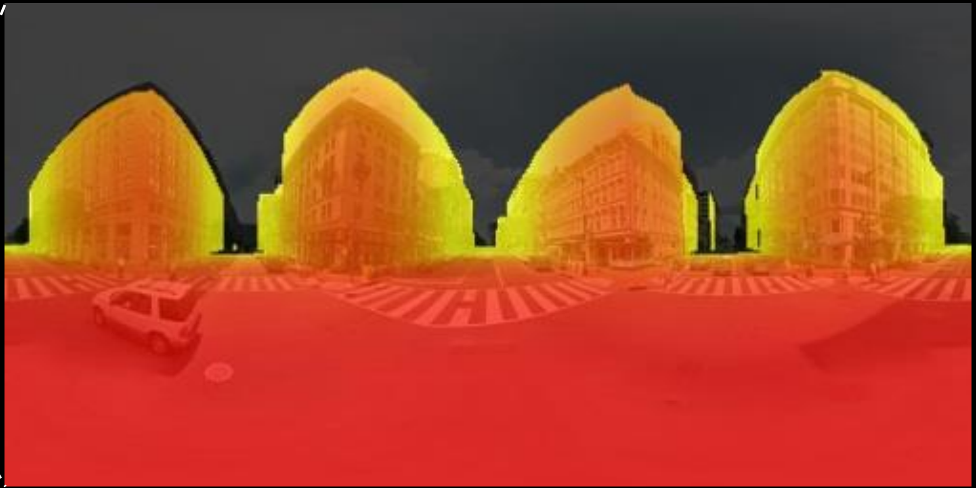

Scraper

Street View Images

Point-cloud Data

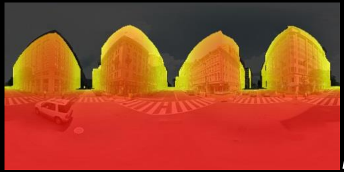

Metadata
(*e.g.,* street topology)


```
{
  - Data: {
        image_width: "13312",
        image_height: "6656",
        tile_width: "512",
        tile_height: "512",
        image_date: "2014-07",
        imagery_type: 1,
        copyright: "© 2015 Google"
    },
  - Projection: {
        projection_type: "spherical",
```



```
  - Links: [
      - {
            yawDeg: "118.97",
            panoId: "WDH0V_F6s9QEAXFMMwktOg",
            road_argb: "0x80fdf872",
            description: "Morse St NE"
        },
      - {
            yawDeg: "299.71",
            panoId: "UCZmw_4Q1SrGAiJoEa9fng",
            road_argb: "0x80fdf872",
            description: "Morse St NE"
```

Scraper

Street View Images

Point-cloud Data

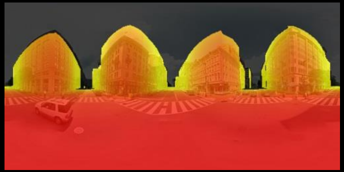

Metadata
(*e.g.*, street topology)


```
{
 - Data: {
      image_width: "13312",
      image_height: "6656",
      tile_width: "512",
      tile_height: "512",
      image_date: "2014-07",
      imagery_type: 1,
      copyright: "© 2015 Google"
   },
 - Projection: {
      projection_type: "spherical",
```


Top-down Google
Maps Imagery

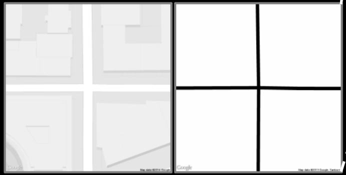

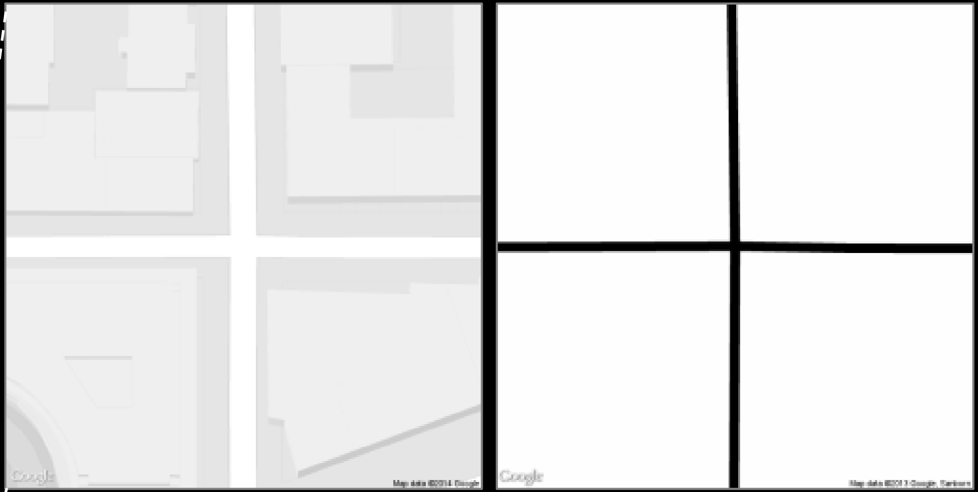

Scraper

Street View Images

Point-cloud Data

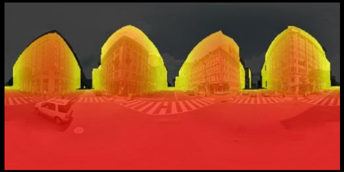

Metadata
(*e.g.*, street topology)

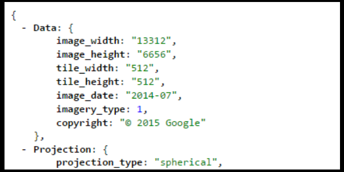

```
{
 - Data: {
       image_width: "13312",
       image_height: "6656",
       tile_width: "512",
       tile_height: "512",
       image_date: "2014-07",
       imagery_type: 1,
       copyright: "© 2015 Google"
   },
 - Projection: {
       projection_type: "spherical",
```


Top-down Google
Maps Imagery

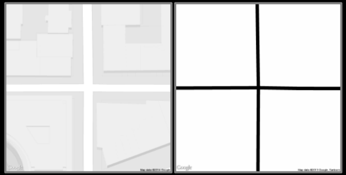

Used to train curb ramp detector and workflow controller

Scraper

Street View Images

Point-cloud Data

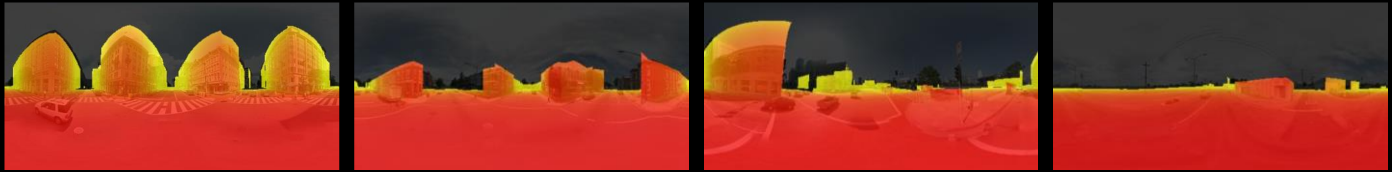

Metadata
(*e.g.*, street topology)

Top-down Google
Maps Imagery

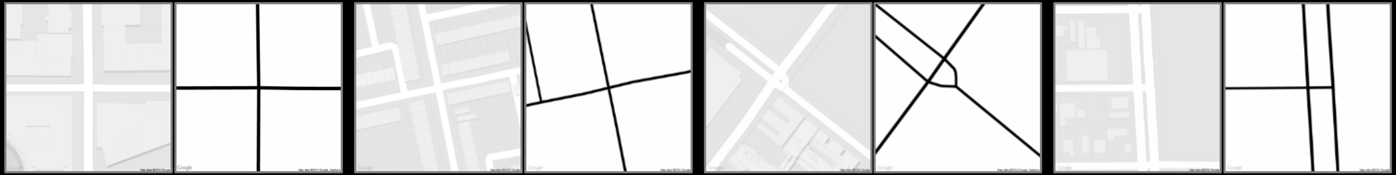

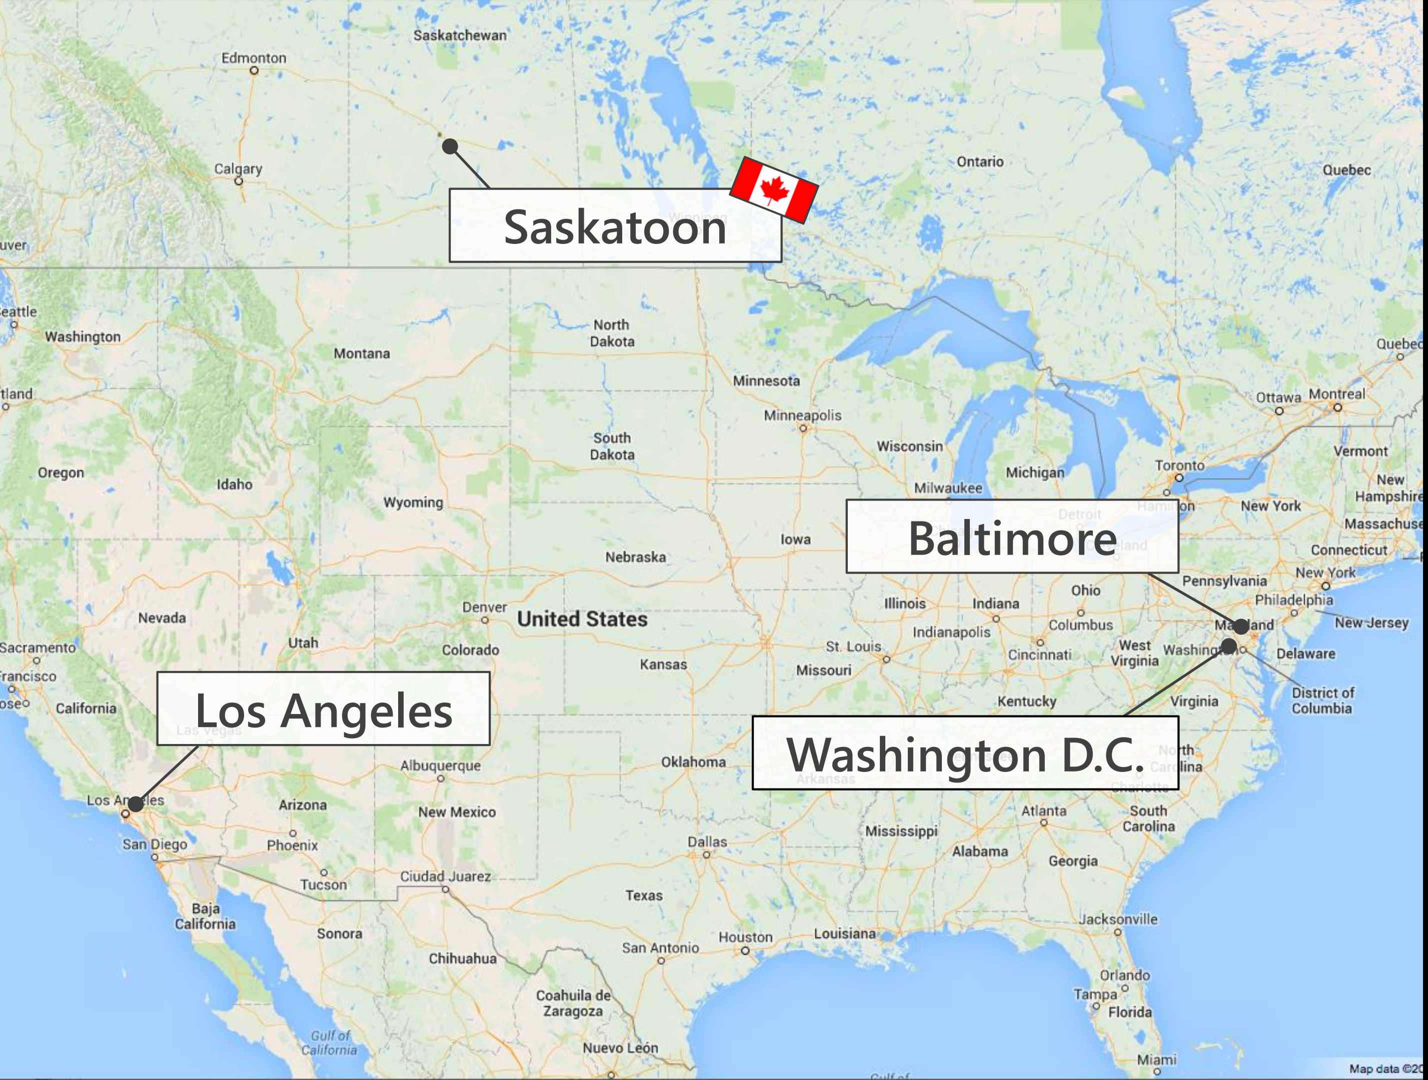
Saskatoon
Baltimore
Los Angeles
Washington D.C.
United States

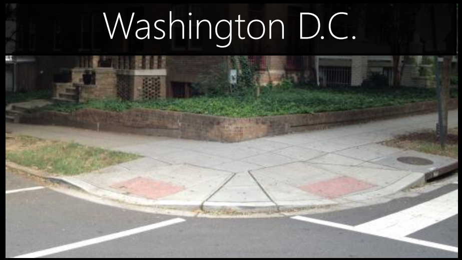
Washington D.C.

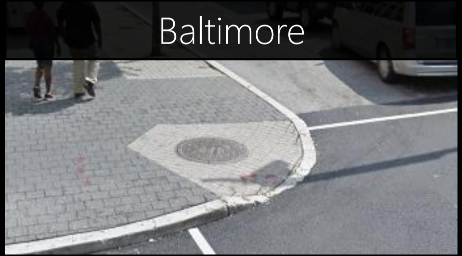
Baltimore

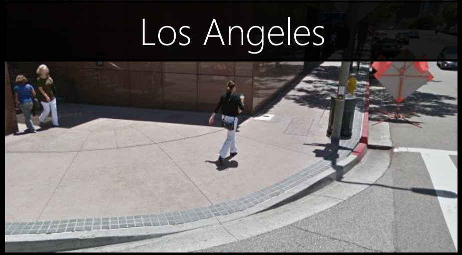
Los Angeles


Saskatoon

Scraper | Areas of Study

Scraper

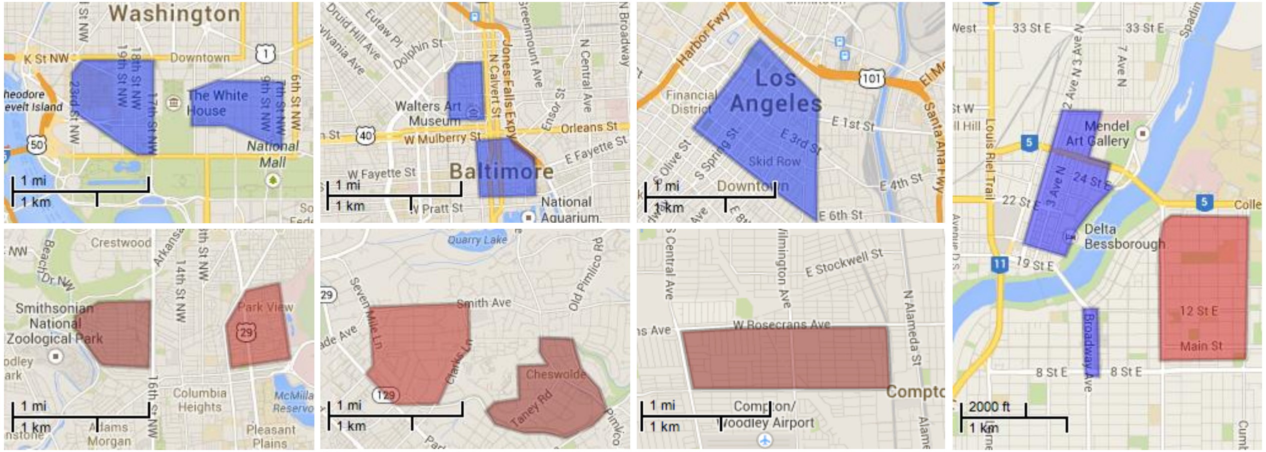

Washington D.C. Baltimore Los Angeles Saskatoon

Total Area: 11.3 km^2

Intersections: 1,086

Curb Ramps: 2,877

Missing Curb Ramps: 647

Avg. GSV Data Age: 2.2 yr*

* At the time of downloading data in summer 2013

Tohme

遠目· *Remote Eye*

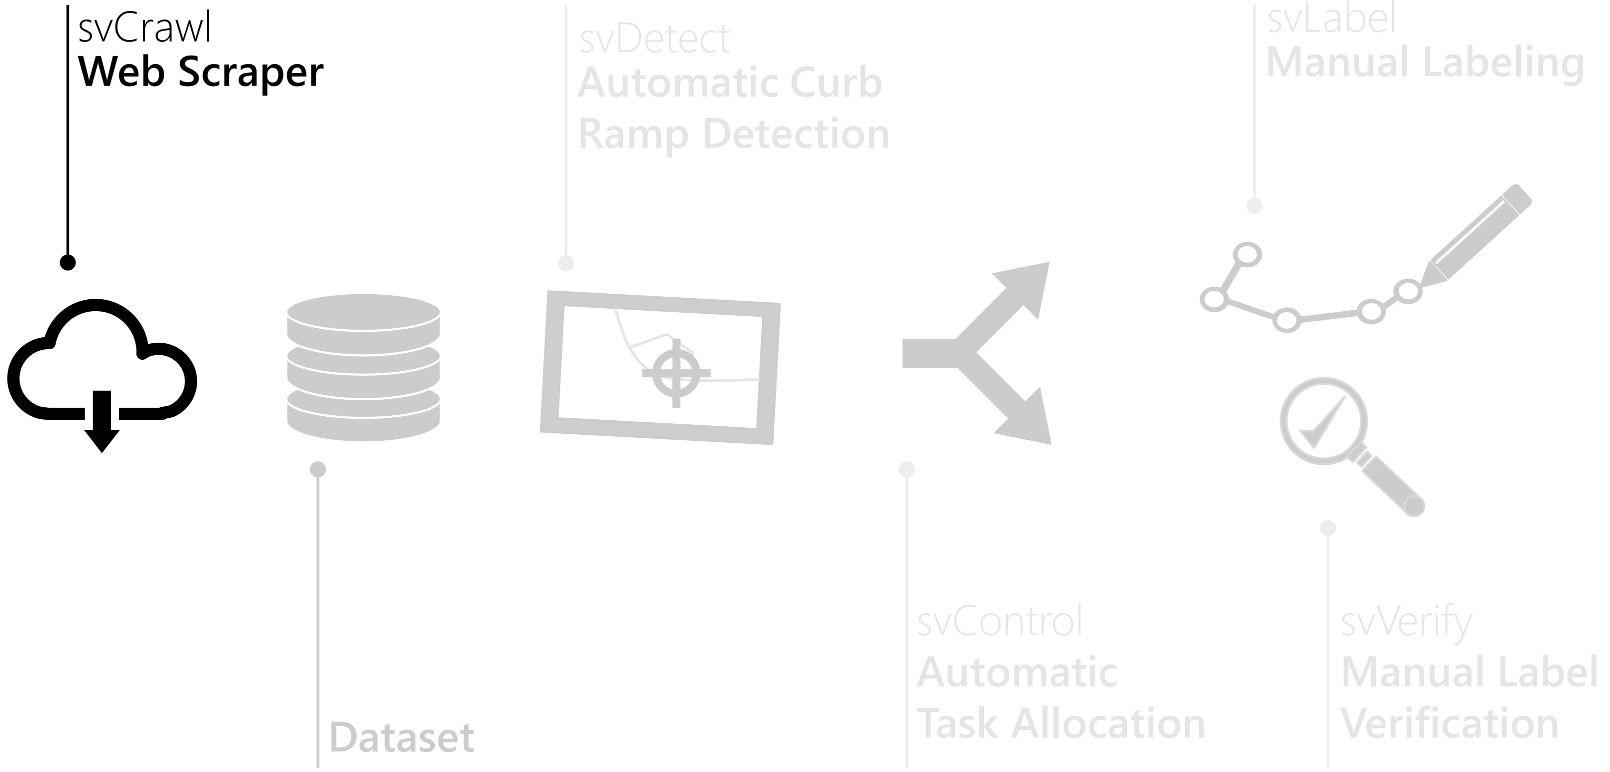

Tohme

遠目·*Remote Eye*

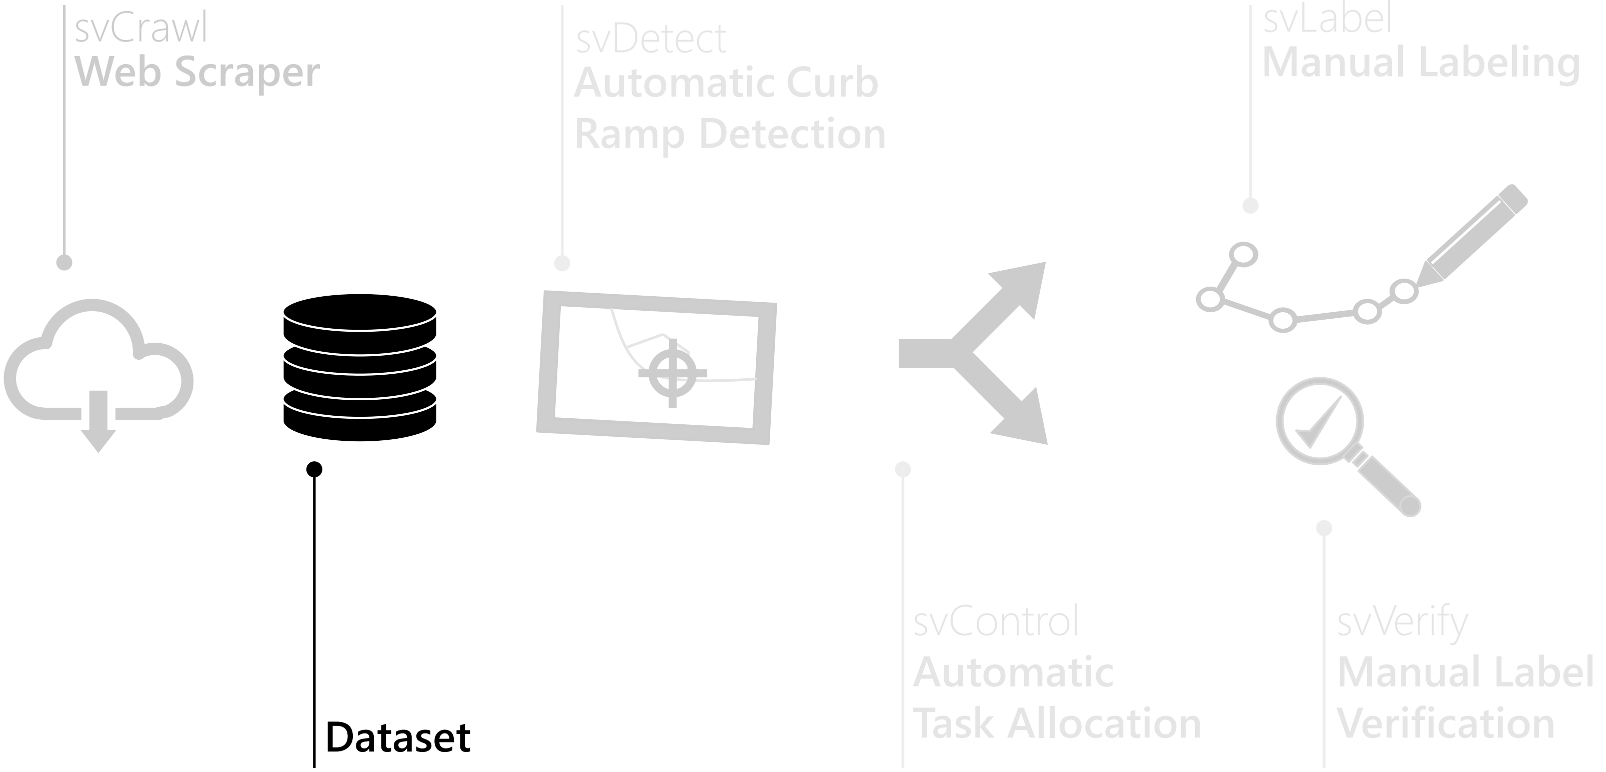

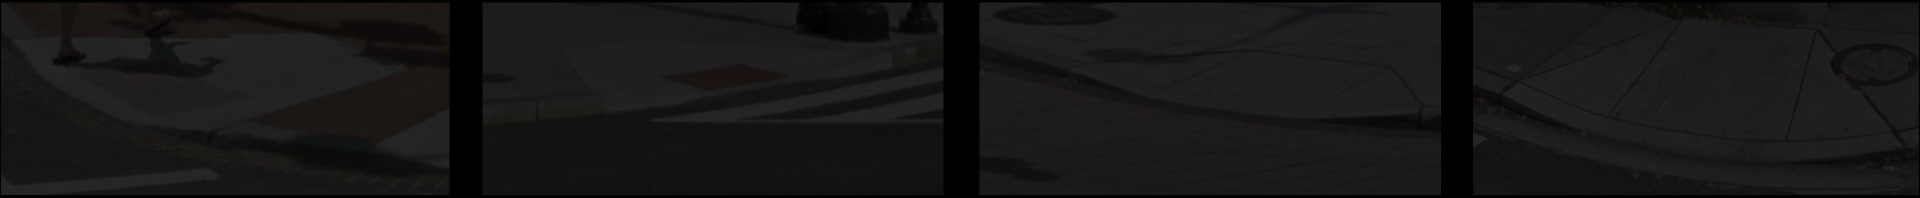

Ground Truth Curb Ramp Dataset

2 researchers labeled curb ramps in our dataset

2,877 curb ramp labels (*Avg.*=2.6 per intersection)

Dataset

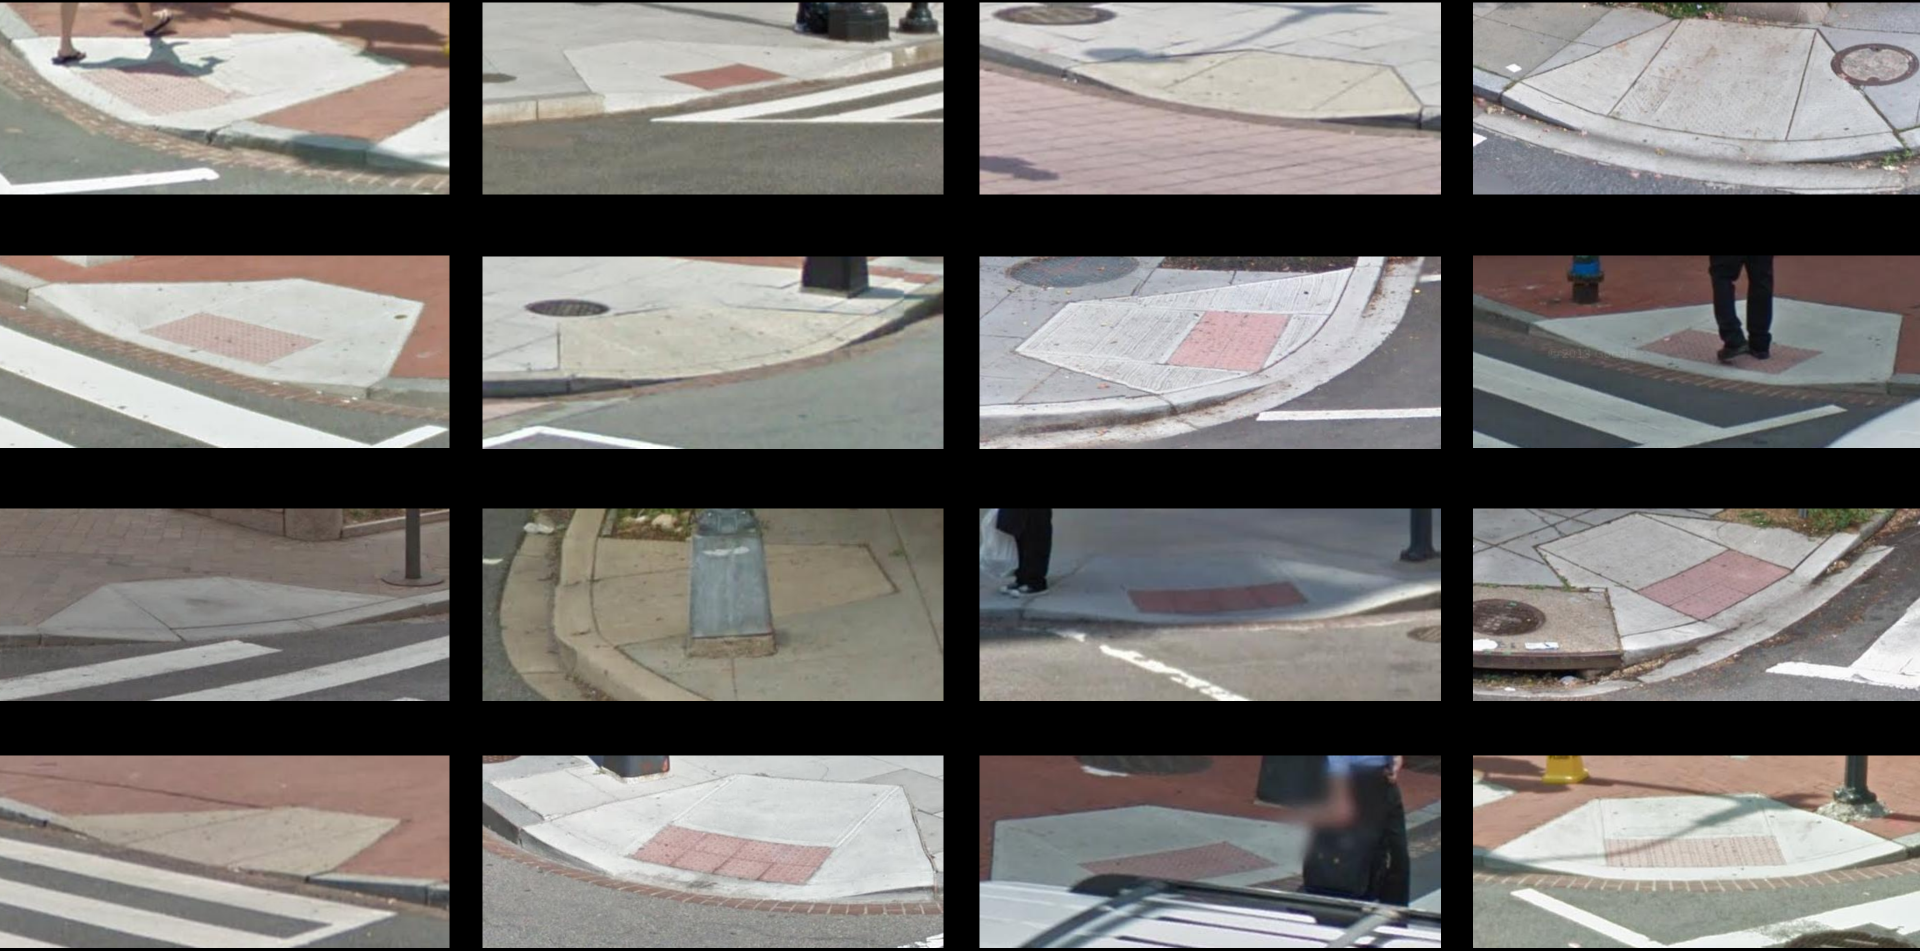

Tohme

遠目·*Remote Eye*

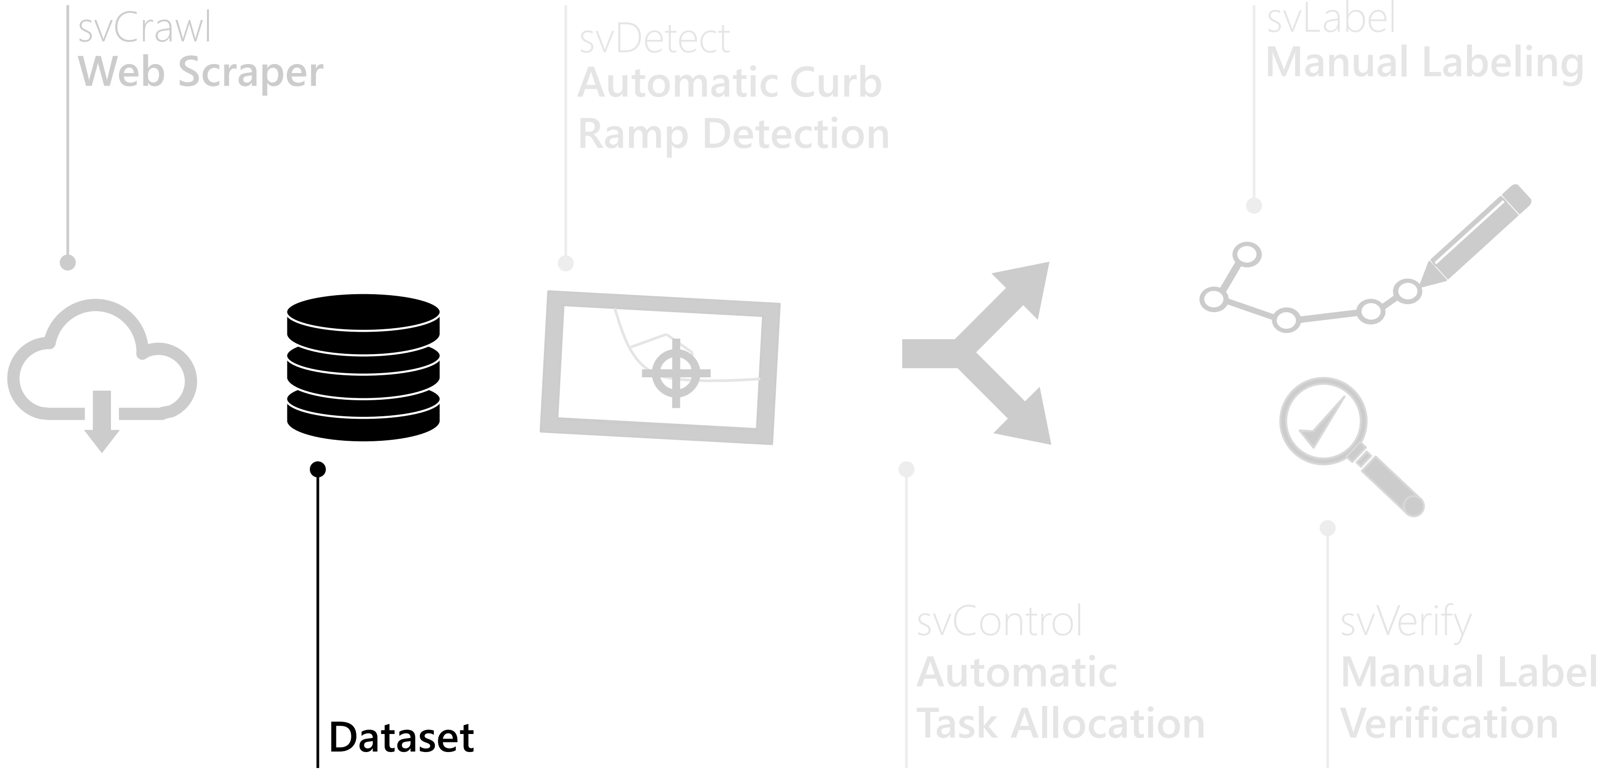

Tohme

遠目·*Remote Eye*

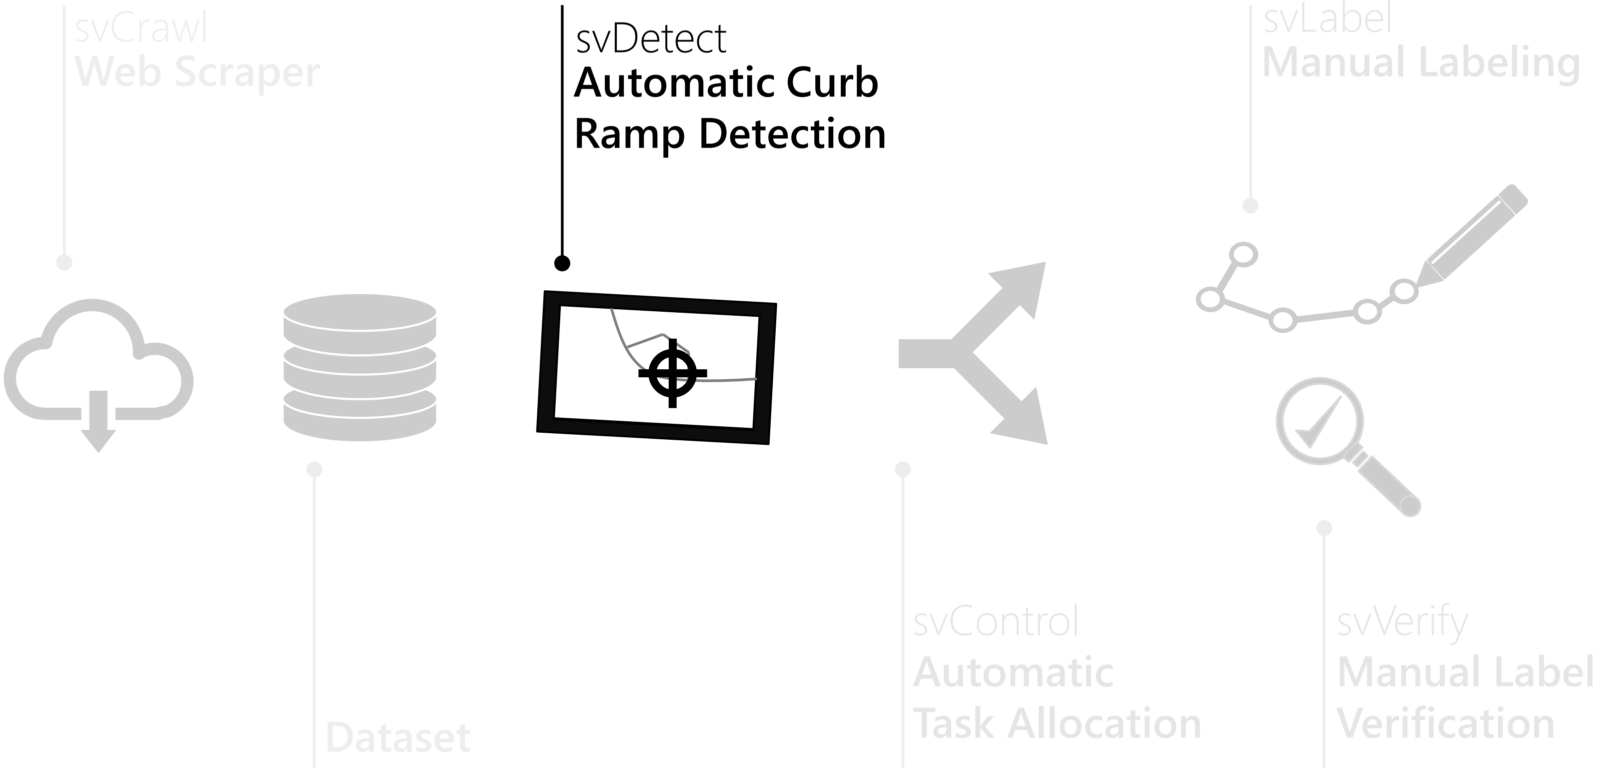

Automatic Curb Ramp Detection

1. Curb ramp detection with Deformable Part Model

2. Post-processing to filter out errors

3. SVM-based classification for output refinement

Deformable Part Models

Felzenszwalb *et al.* 2008

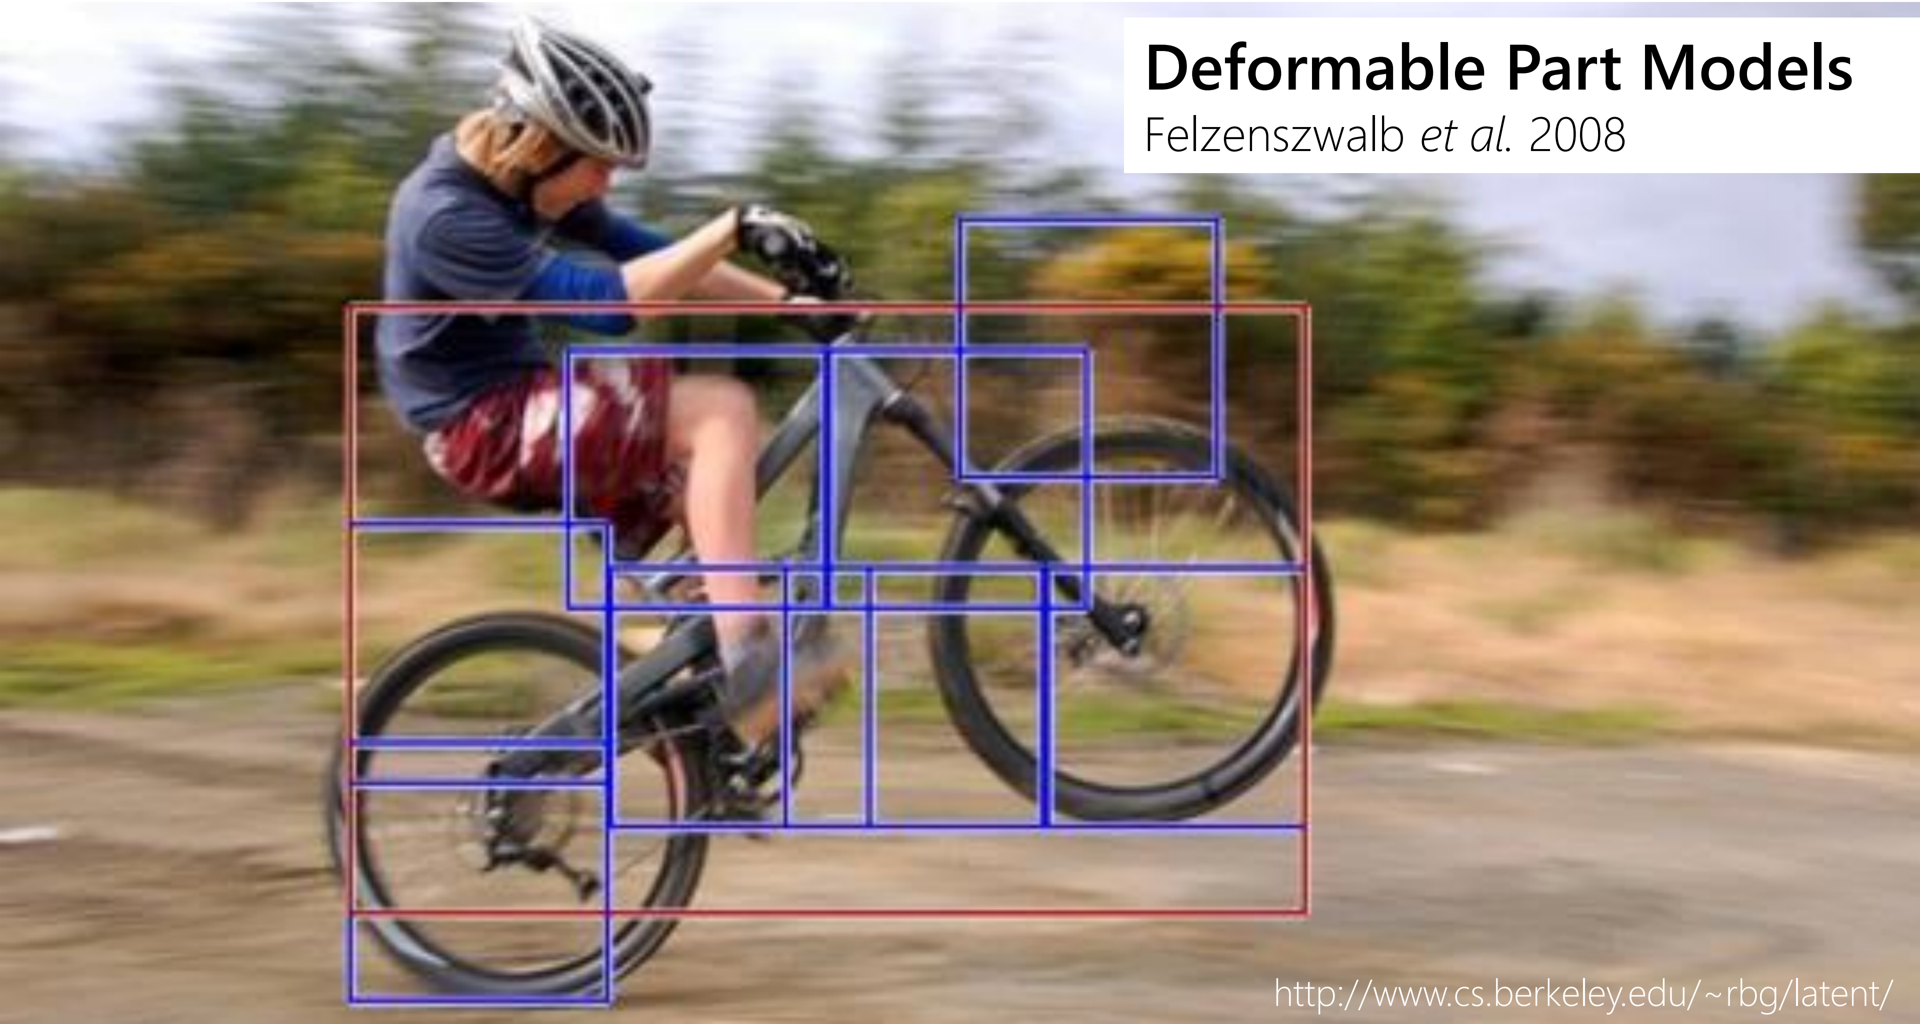

http://www.cs.berkeley.edu/~rbg/latent/

Deformable Part Models

Felzenszwalb *et al.* 2008

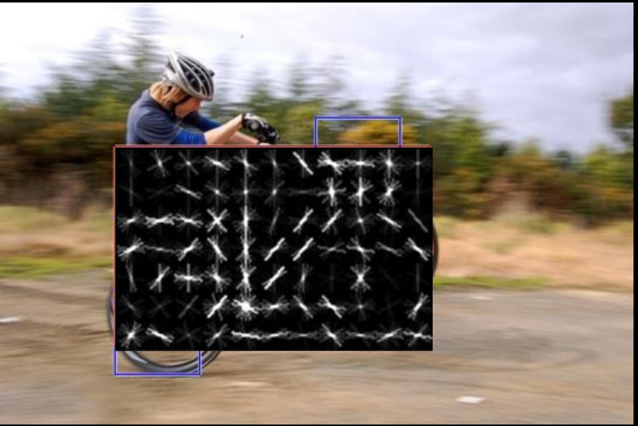

Root filter

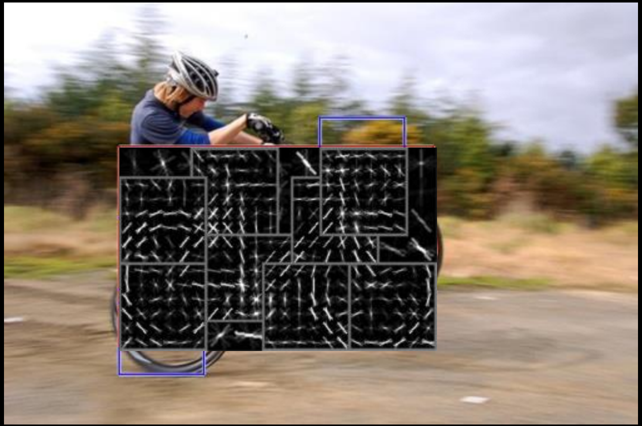

Parts filter

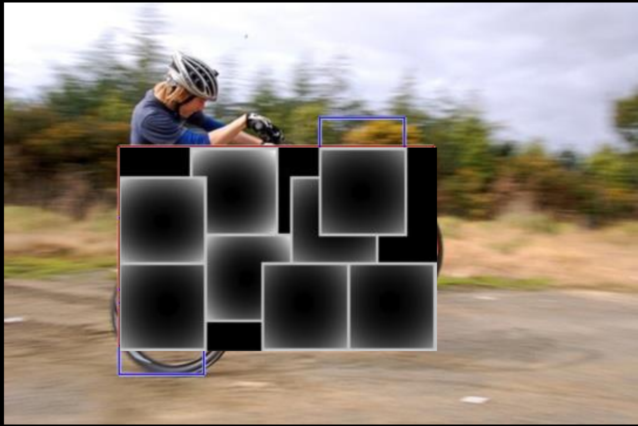

Displacement cost

http://www.cs.berkeley.edu/~rbg/latent/

Automatic Curb Ramp Detection

Detected Labels

Stage 1: Deformable Part Model

Sliding window detection with deformable part model

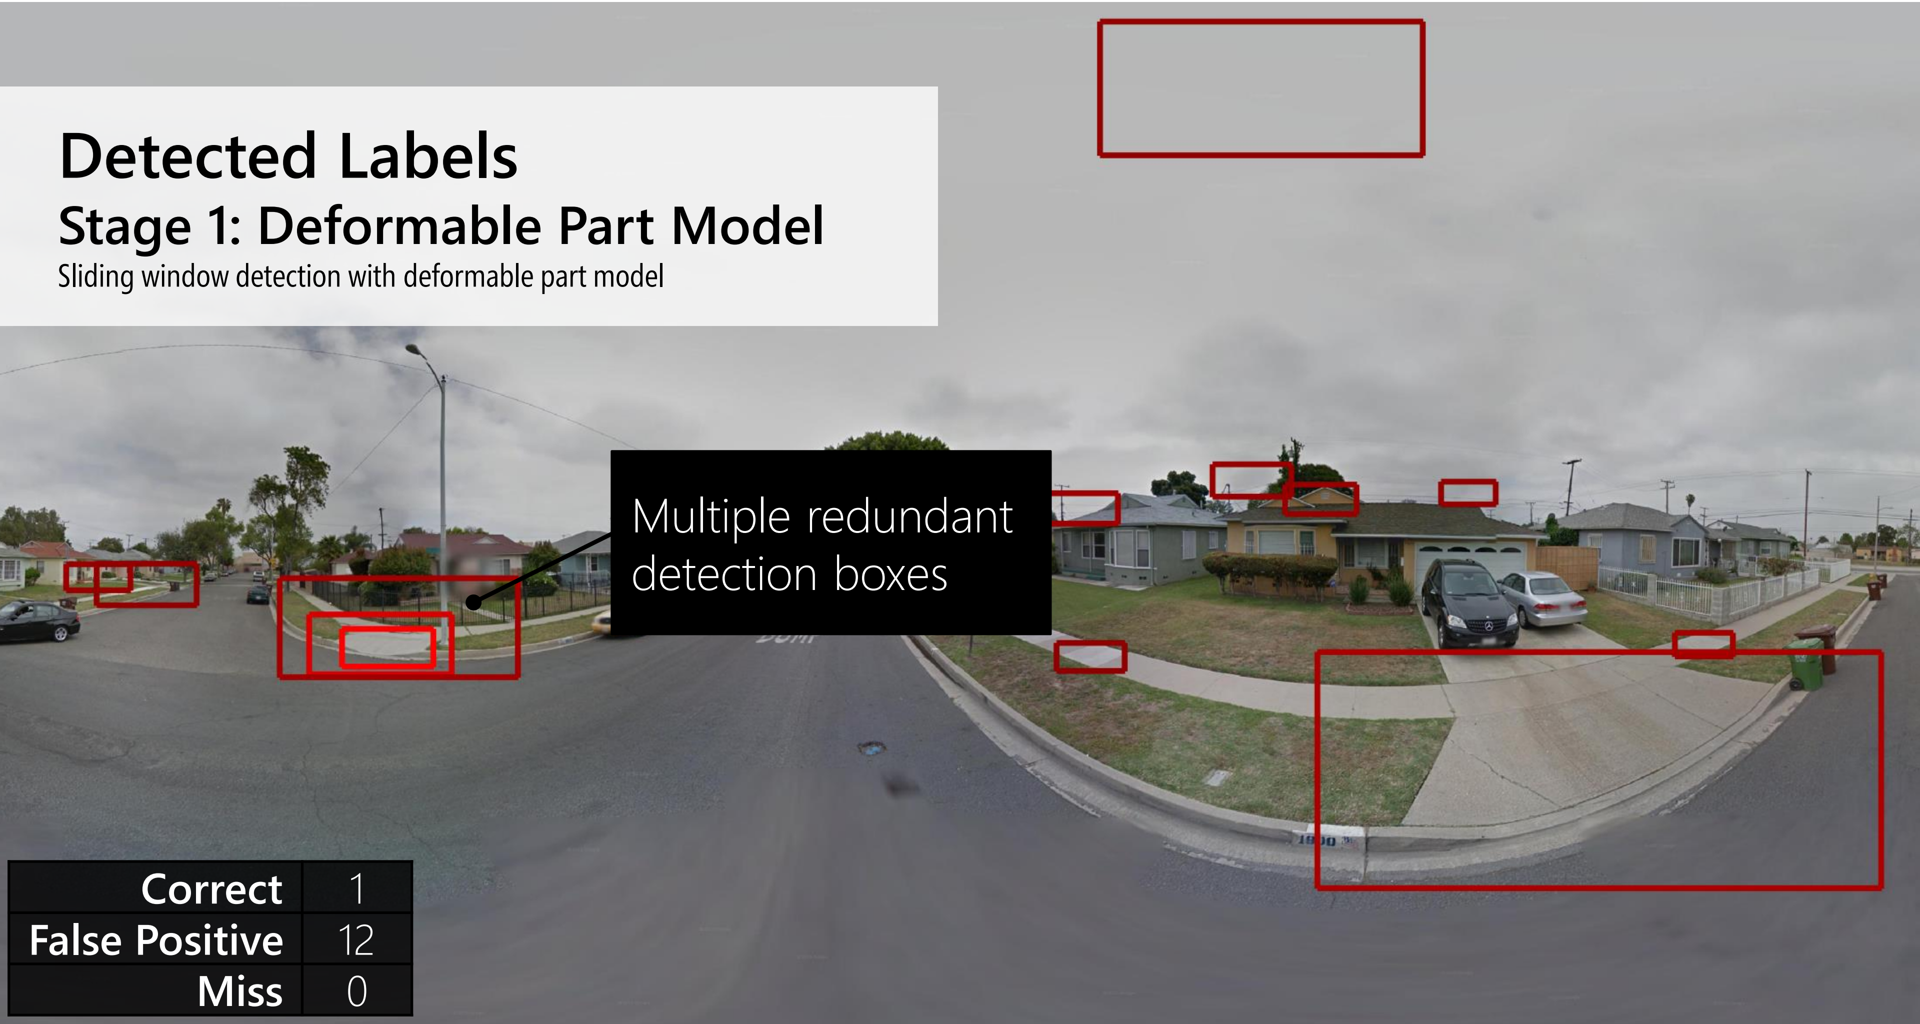

Correct	1
False Positive	12
Miss	0

Detected Labels

Stage 1: Deformable Part Model

Sliding window detection with deformable part model

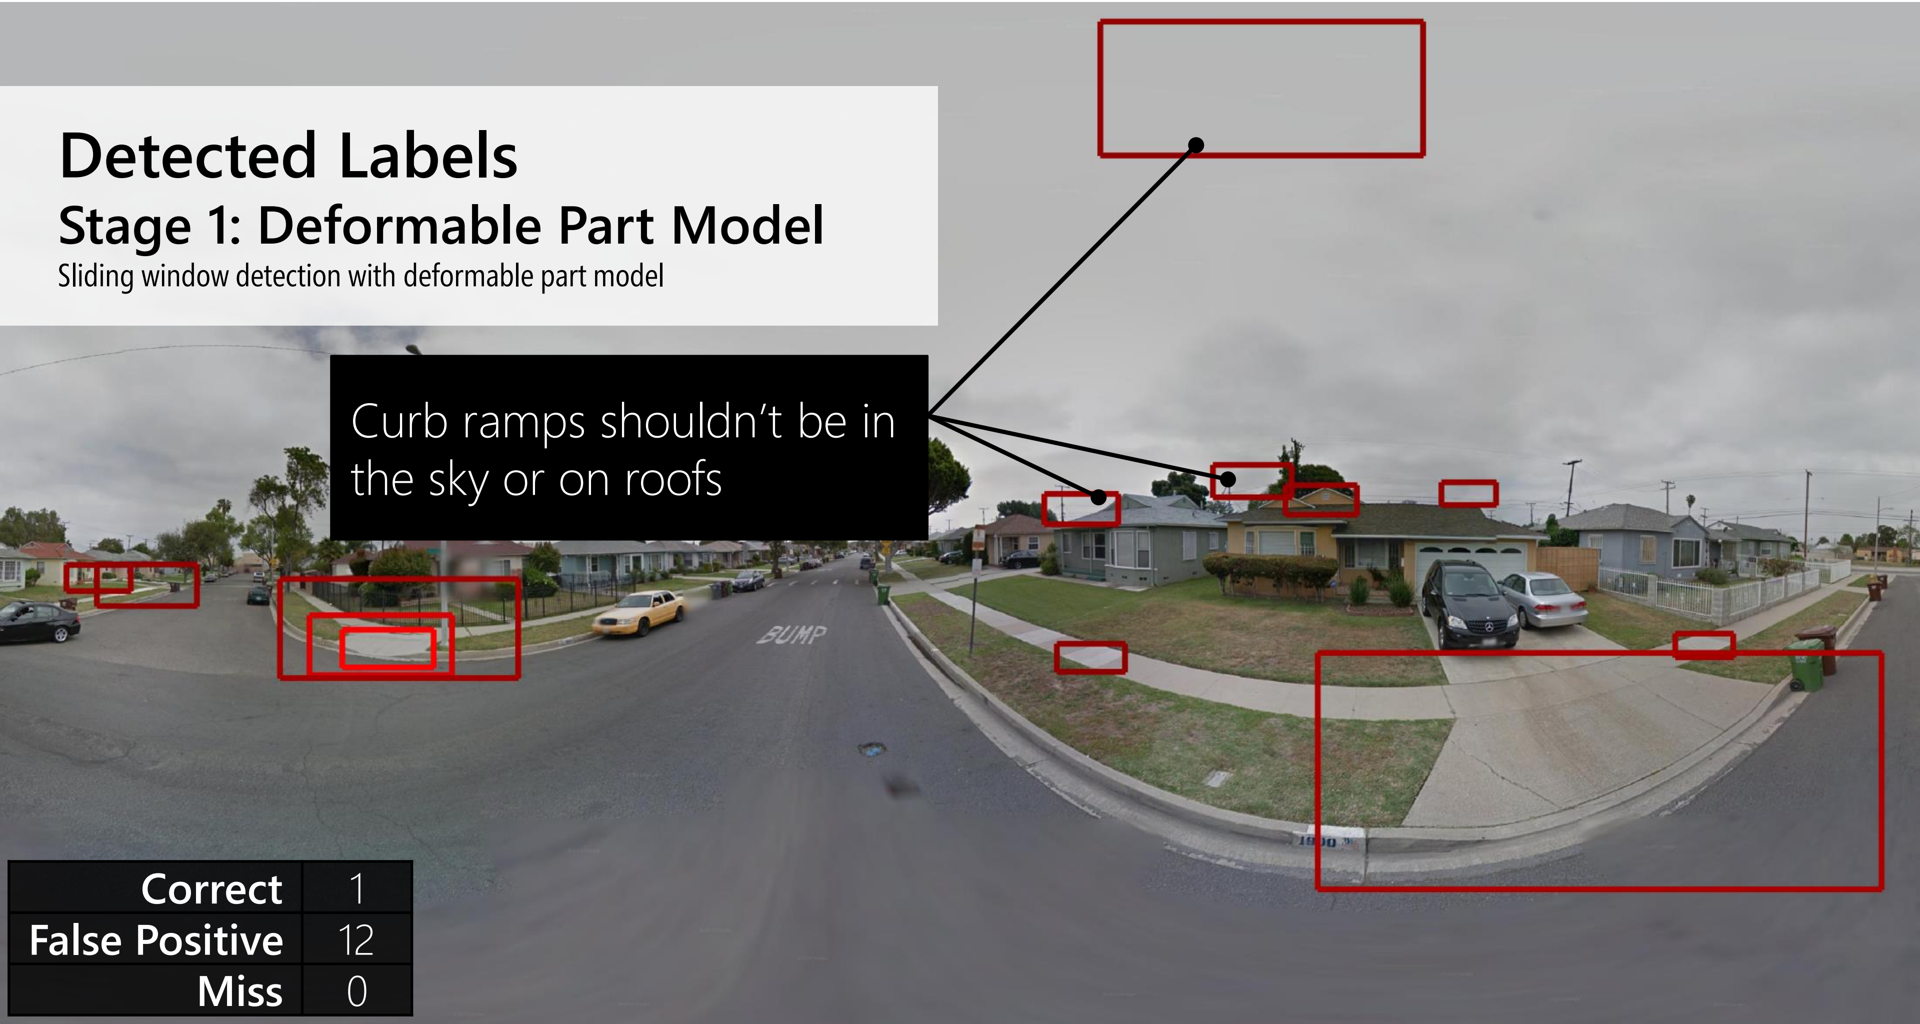

Curb ramps shouldn't be in the sky or on roofs

Correct	1
False Positive	12
Miss	0

Detected Labels

Stage 2: Post-processing

Rejects errors using 3D data and applies non-maxima suppression

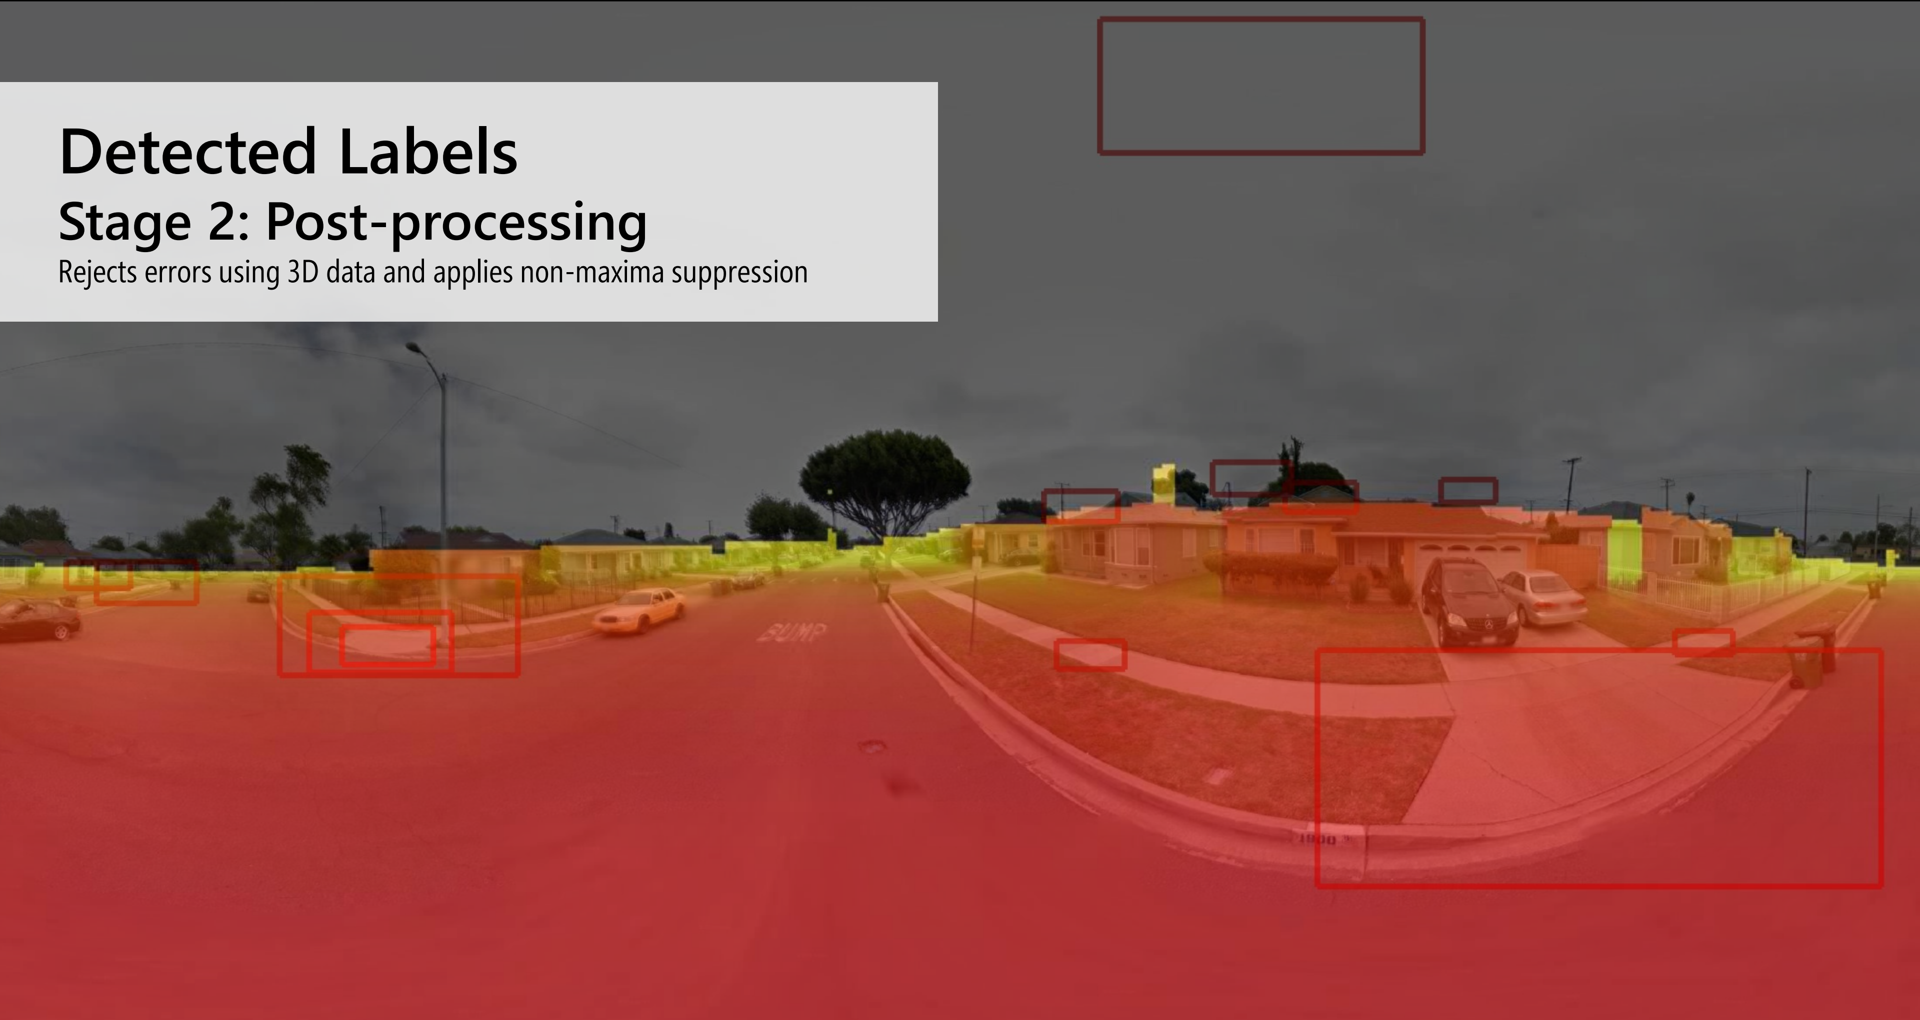

Detected Labels

Stage 2: Post-processing

Rejects errors using 3D data and applies non-maxima suppression

Correct	1
False Positive	5
Miss	0

Detected Labels

Stage 3: SVM-based Refinement

Takes size, color, and position and further filters out false detections

Correct	1
False Positive	5
Miss	0

Detected Labels

Stage 3: SVM-based Refinement

Takes size, color, and position and further filters out false detections

Correct	1
False Positive	3
Miss	0

Automatic Curb Ramp Detection

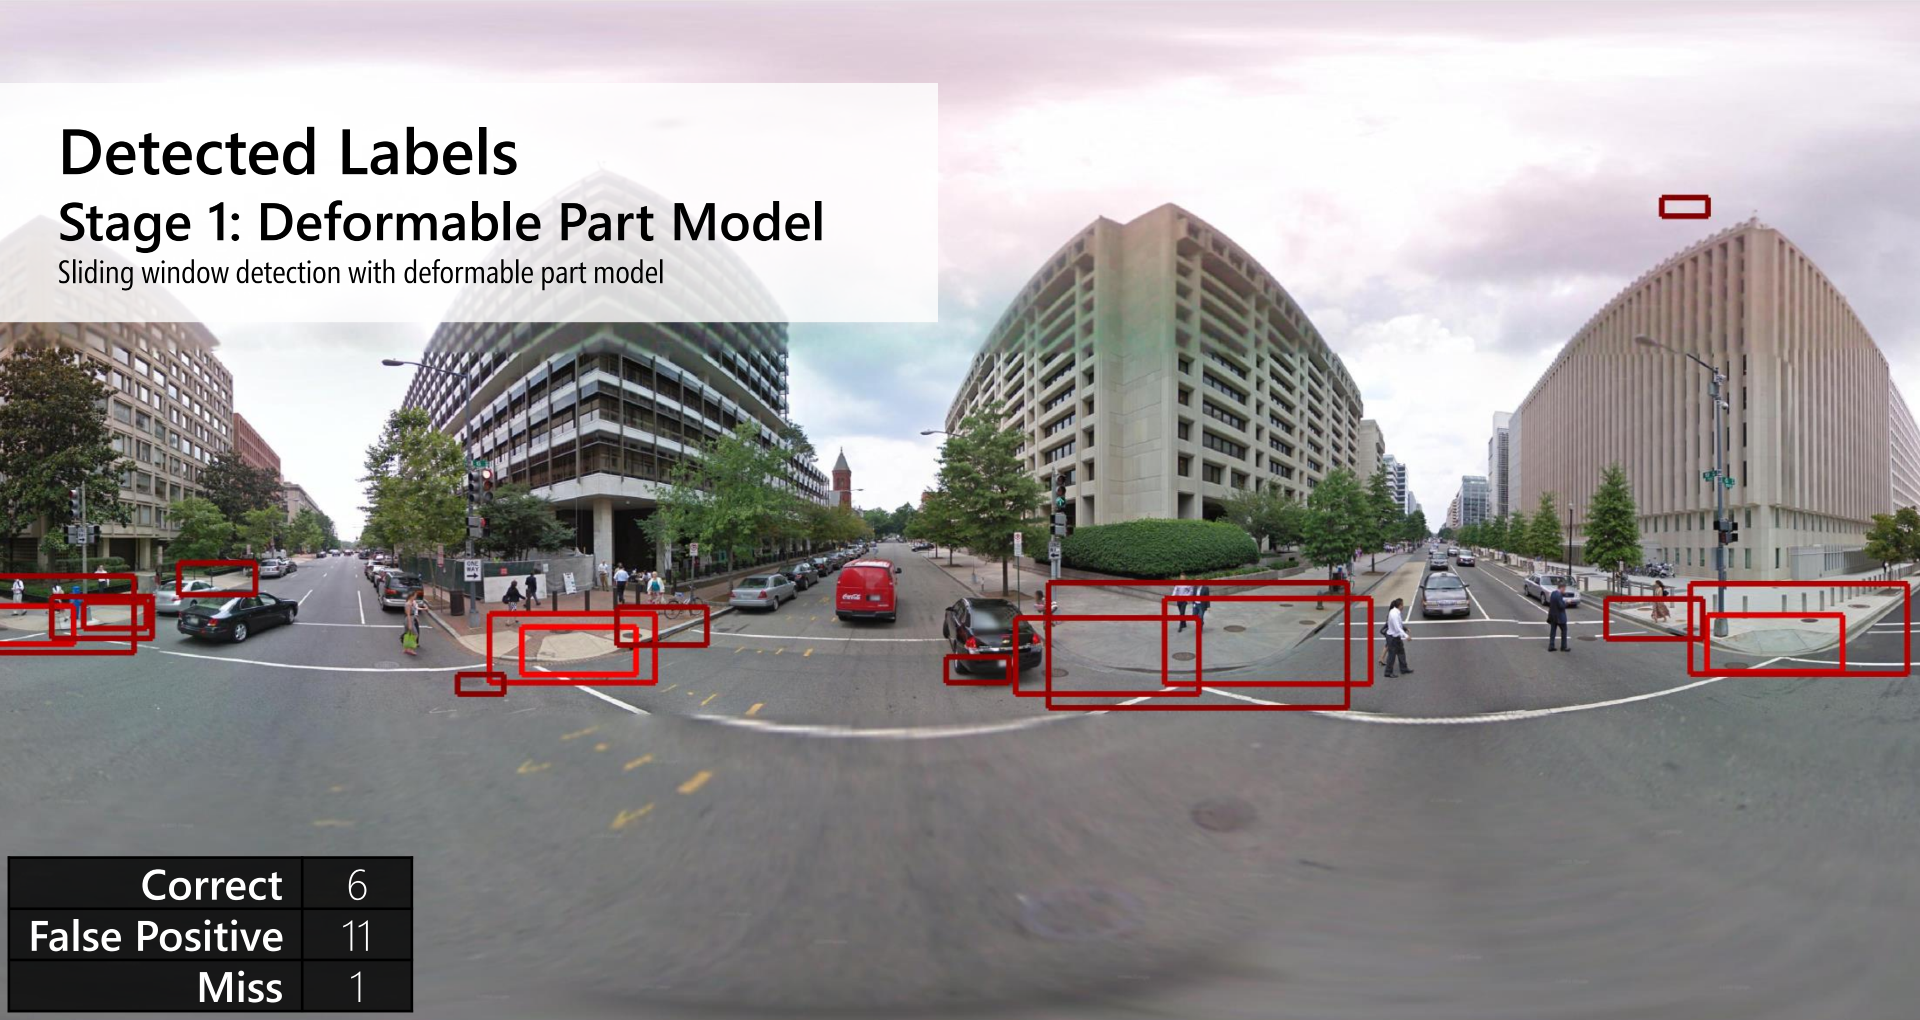

Detected Labels

Stage 1: Deformable Part Model

Sliding window detection with deformable part model

Correct	6
False Positive	11
Miss	1

Detected Labels

Stage 2: Post-processing

Rejects errors using 3D data and applies non-maxima suppression

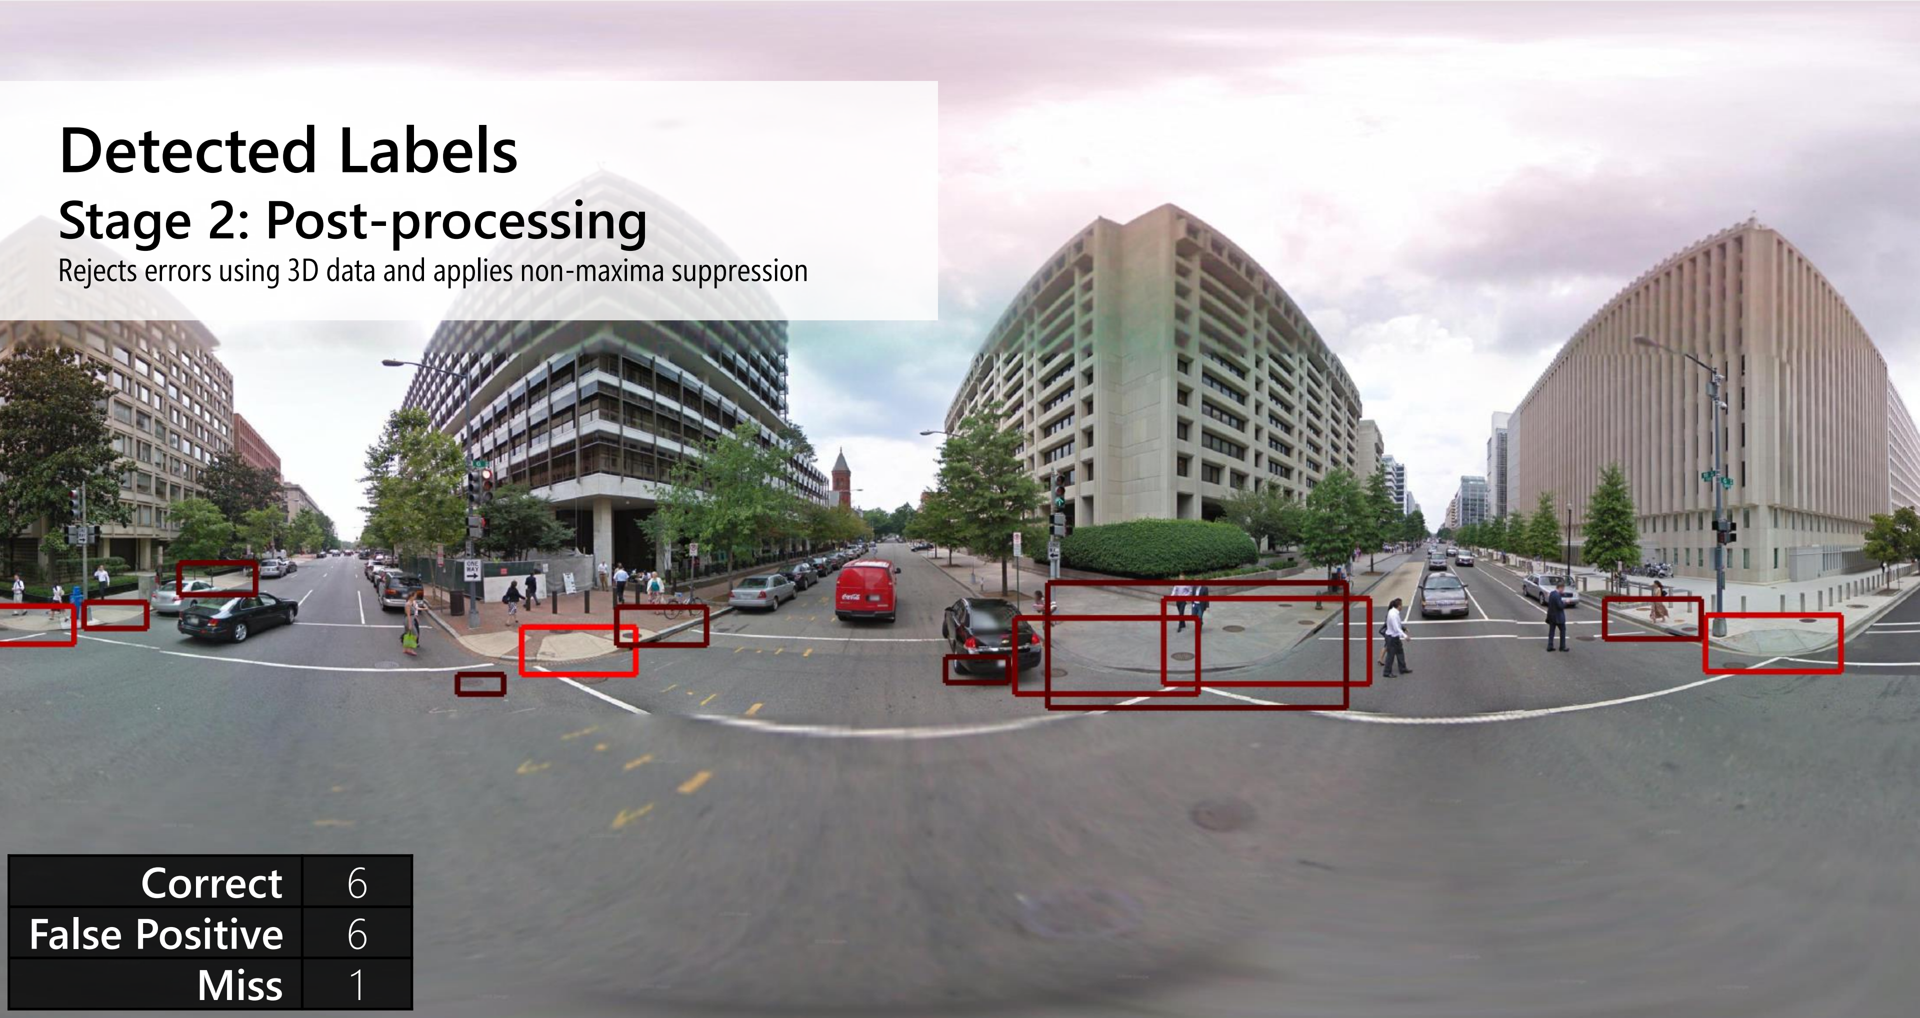

Correct	6
False Positive	6
Miss	1

Detected Labels

Stage 3: SVM-based Refinement

Takes size, color, and position and further filters out false detections

Correct	6
False Positive	4
Miss	1

Automatic Curb Ramp Detection Accuracy

Used two-fold cross validation to evaluate CV sub-system

Computer Vision Sub-System Results

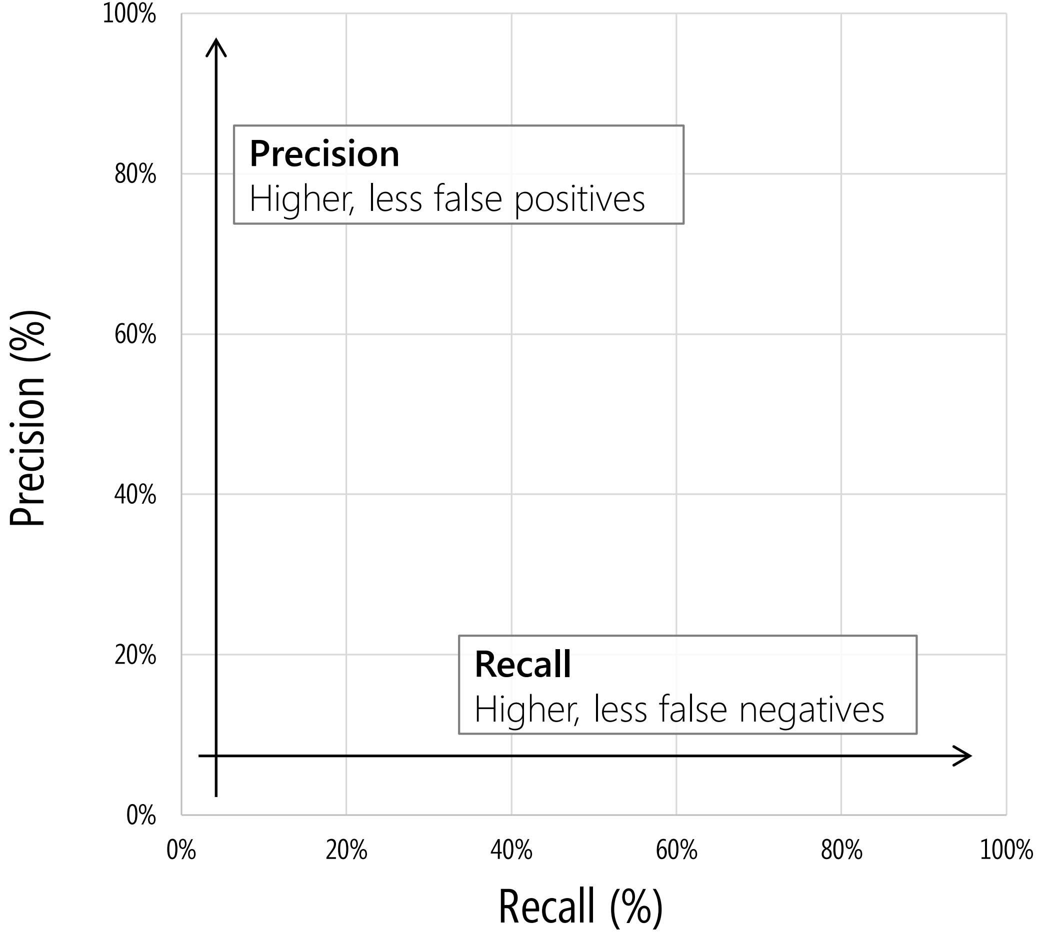

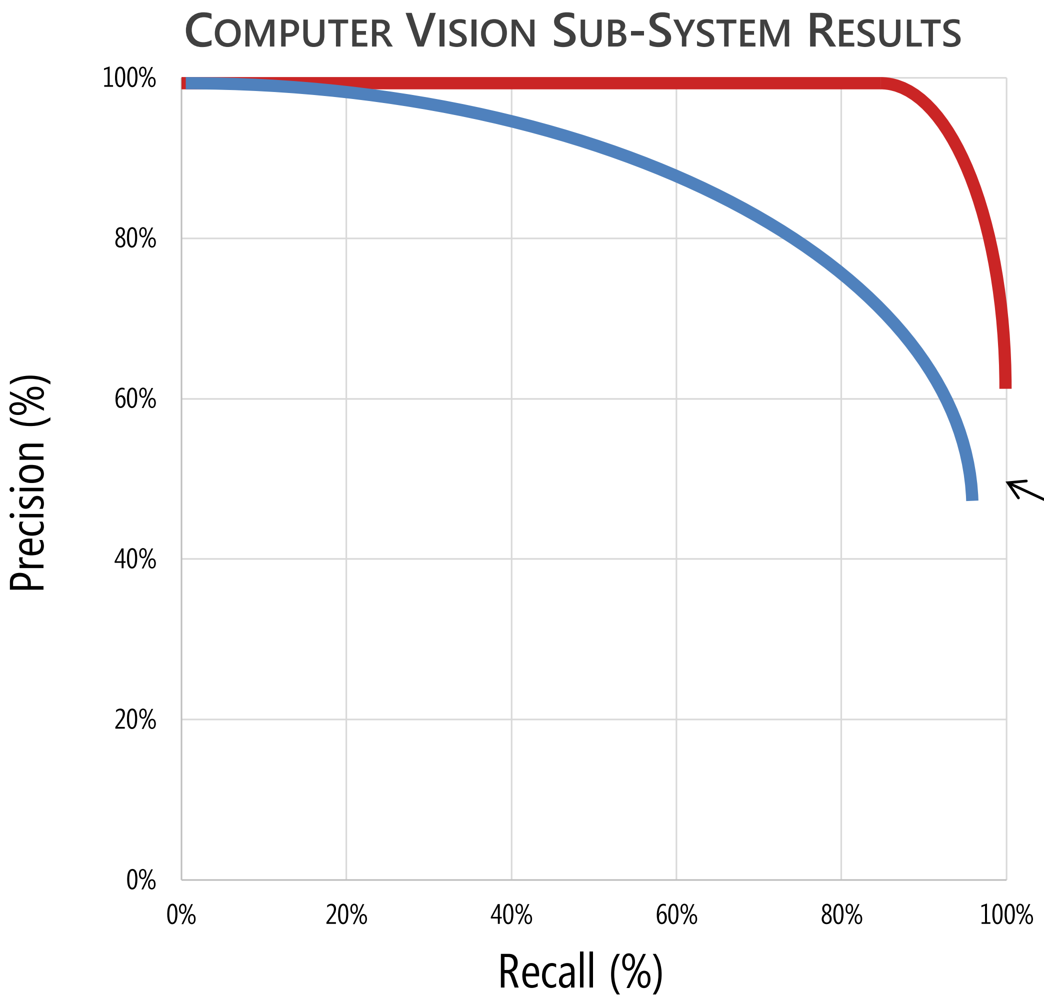

Goal: maximize area under curve

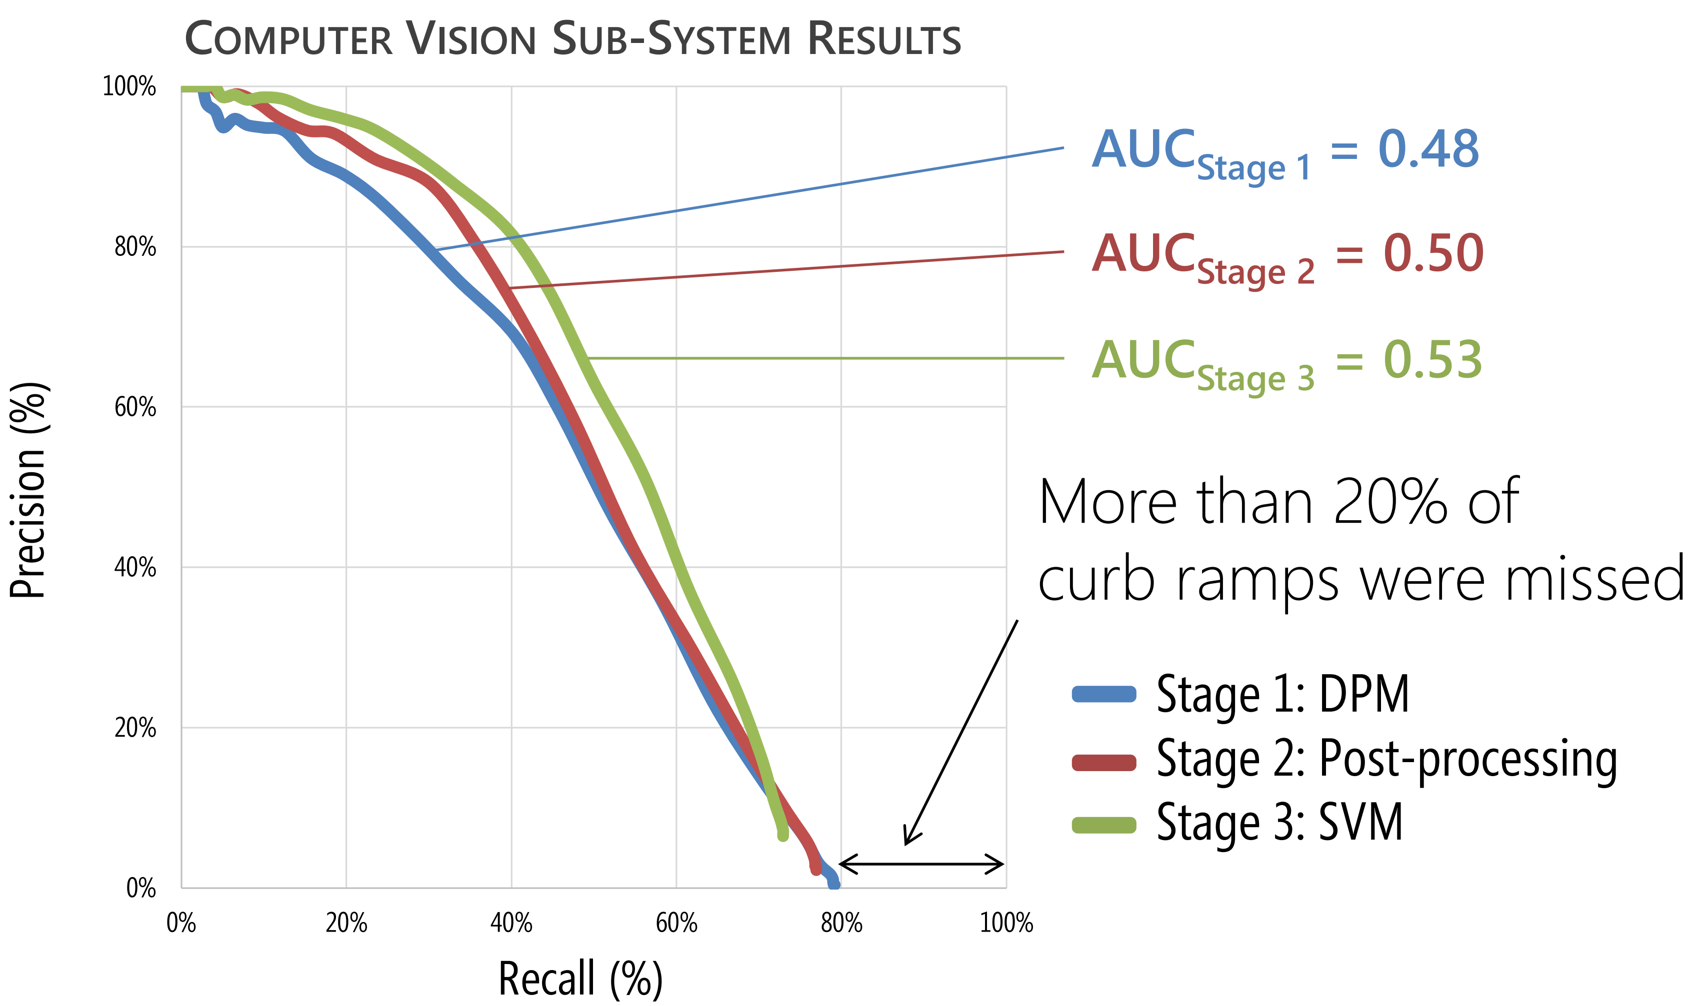

Computer Vision Sub-System Results
100%
80%
60%
40%
20%
0%
Precision (%)
0%
20%
40%
60%
80%
100%
Recall (%)
AUC$_{Stage\ 1}$ = 0.48
AUC$_{Stage\ 2}$ = 0.50
AUC$_{Stage\ 3}$ = 0.53
More than 20% of curb ramps were missed
Stage 1: DPM
Stage 2: Post-processing
Stage 3: SVM

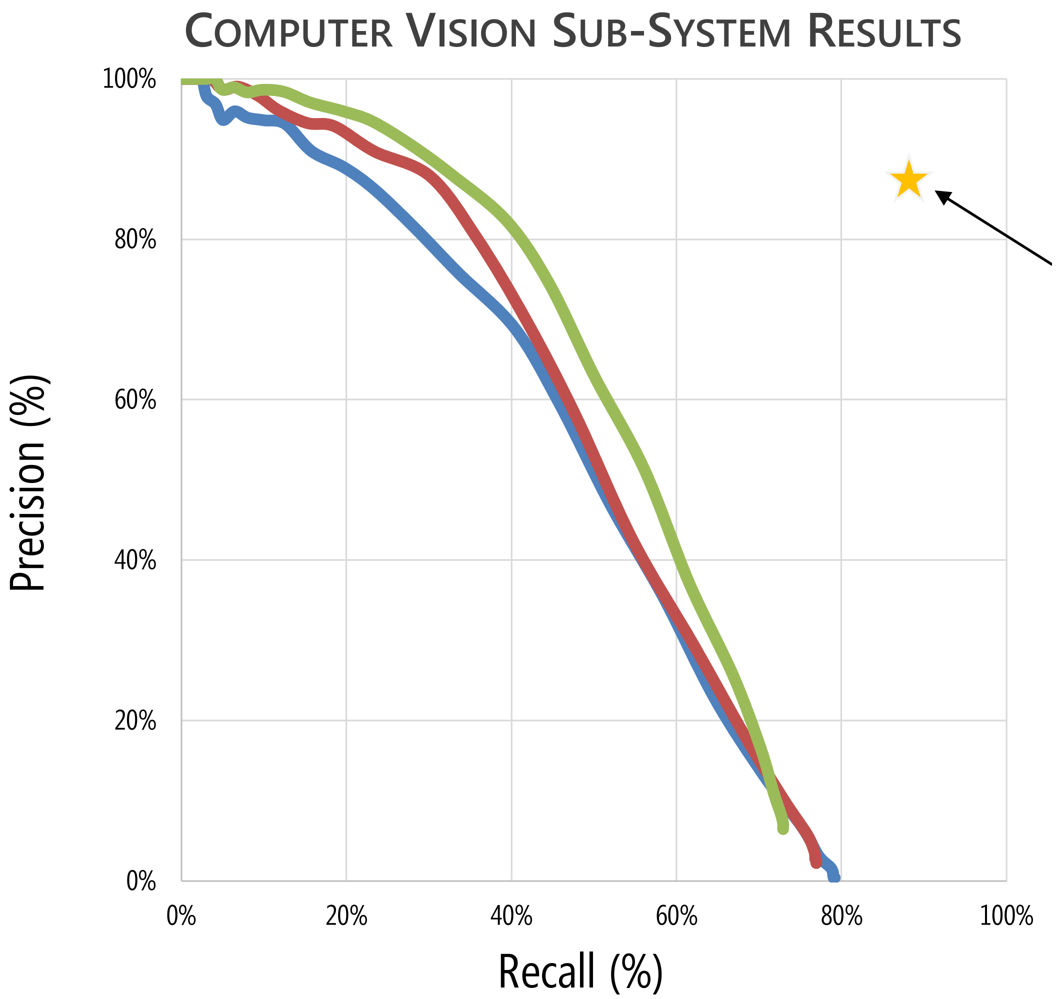

Computer vision alone is not sufficient to accurately find curb ramps

Curb Ramp Detection is a Hard Problem

Occlusion

Occlusion

Curb Ramp Detection is a Hard Problem

Occlusion

Illumination

Illumination

Curb Ramp Detection is a Hard Problem

Occlusion

Scale

Illumination

Scale

Curb Ramp Detection is a Hard Problem

Occlusion

Illumination

Scale

Viewpoint Variation

View Point Variation

Curb Ramp Detection is a Hard Problem

Occlusion

Scale

Structures Similar to Curb Ramps

Illumination

Viewpoint Variation

Structures Similar to Curb Ramps

Curb Ramp Detection is a Hard Problem

Occlusion

Illumination

Scale

Viewpoint Variation

Structures Similar to Curb Ramps

Curb Ramp Design Variation

Curb Ramp Design Variation

Tohme

遠目·*Remote Eye*

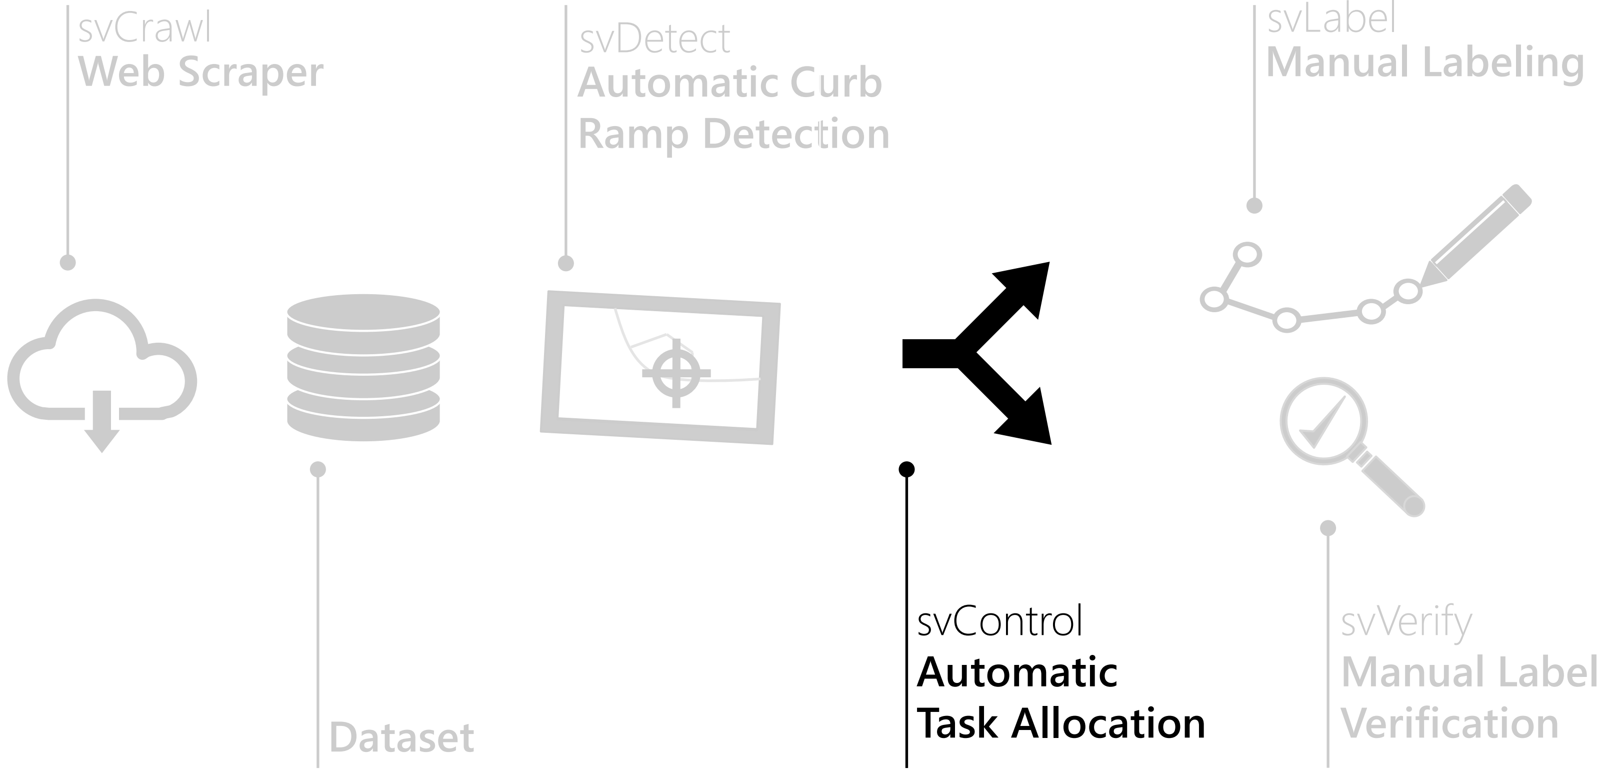

Automatic Task Allocation | Features to Assess Scene Difficulty for CV

A number of streets from metadata

Depth information for a road width and variance in distance

Top-down images to assess complexity of an intersection

A number of detections and confidence values

Automatic Task Allocation | Features to Assess Scene Difficulty for CV

A number of street from metadata

Depth information for a road width and variance in distance

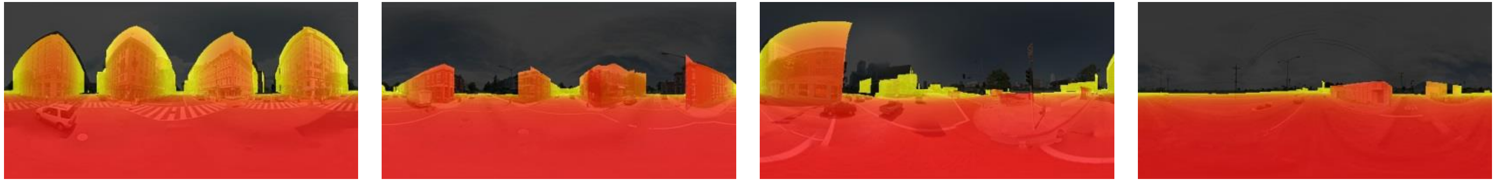

Top-down images to assess complexity of an intersection

A number of detections and confidence values

Depth information for a road width and variance in distance

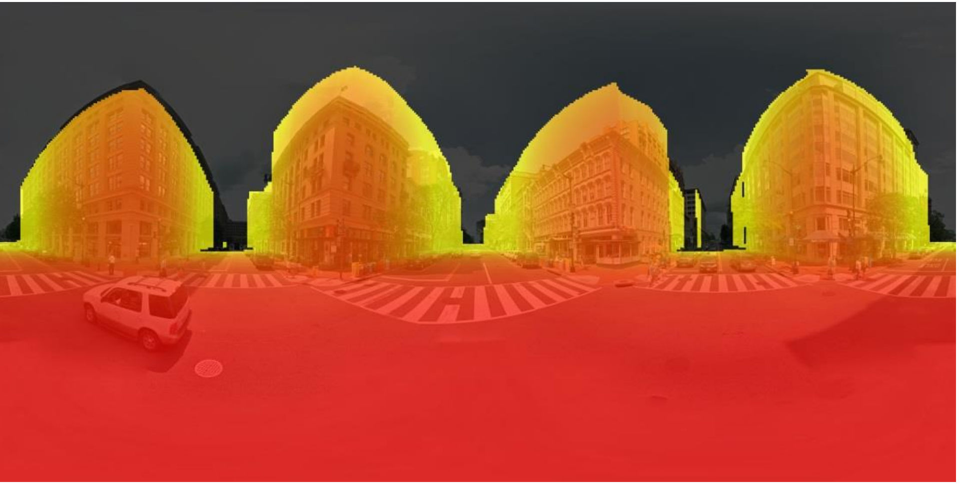

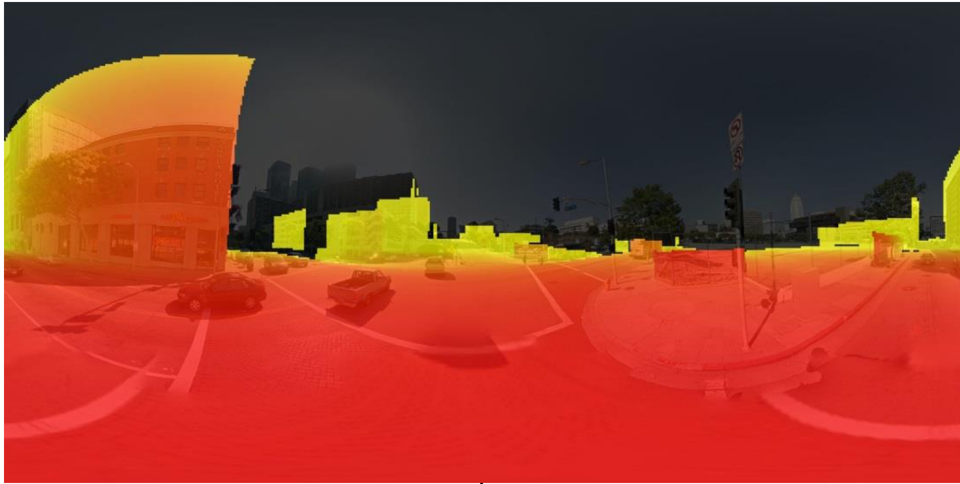

Finding curb ramps on distant sidewalks is difficult; they look smaller (i.e., scaling issue)

A number of streets from metadata

Depth information for a road width and variance in distance

Top-down images to assess complexity of an intersection

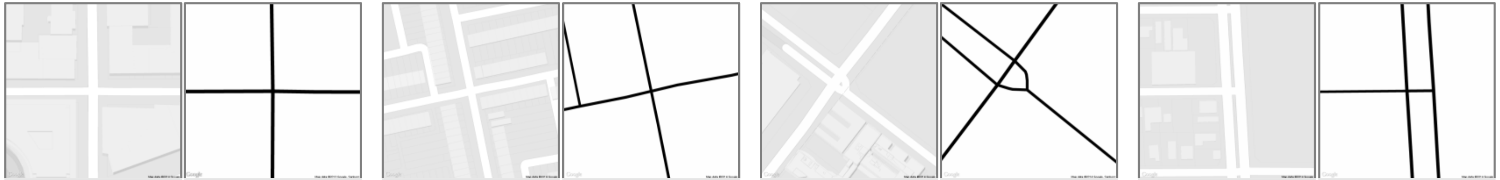

A number of detections and confidence values

As a proxy for intersection complexity, we count the number of black :
more black pixels = more complex intersection
(*i.e.,* more viewpoint variation)

A number of streets from metadata

Depth information for a road width and variance in distance

Top-down images to assess complexity of an intersection

CV Output: A number of detections and confidence values

Binary Classification

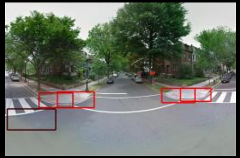

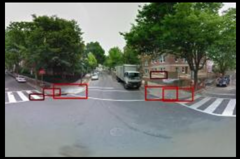

Binary classifier to detect false-negatives

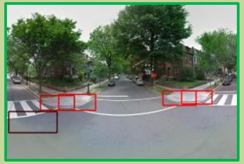

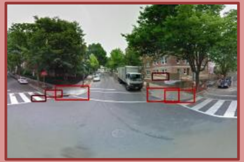

Pass

Fail

Tohme
遠目· *Remote Eye*

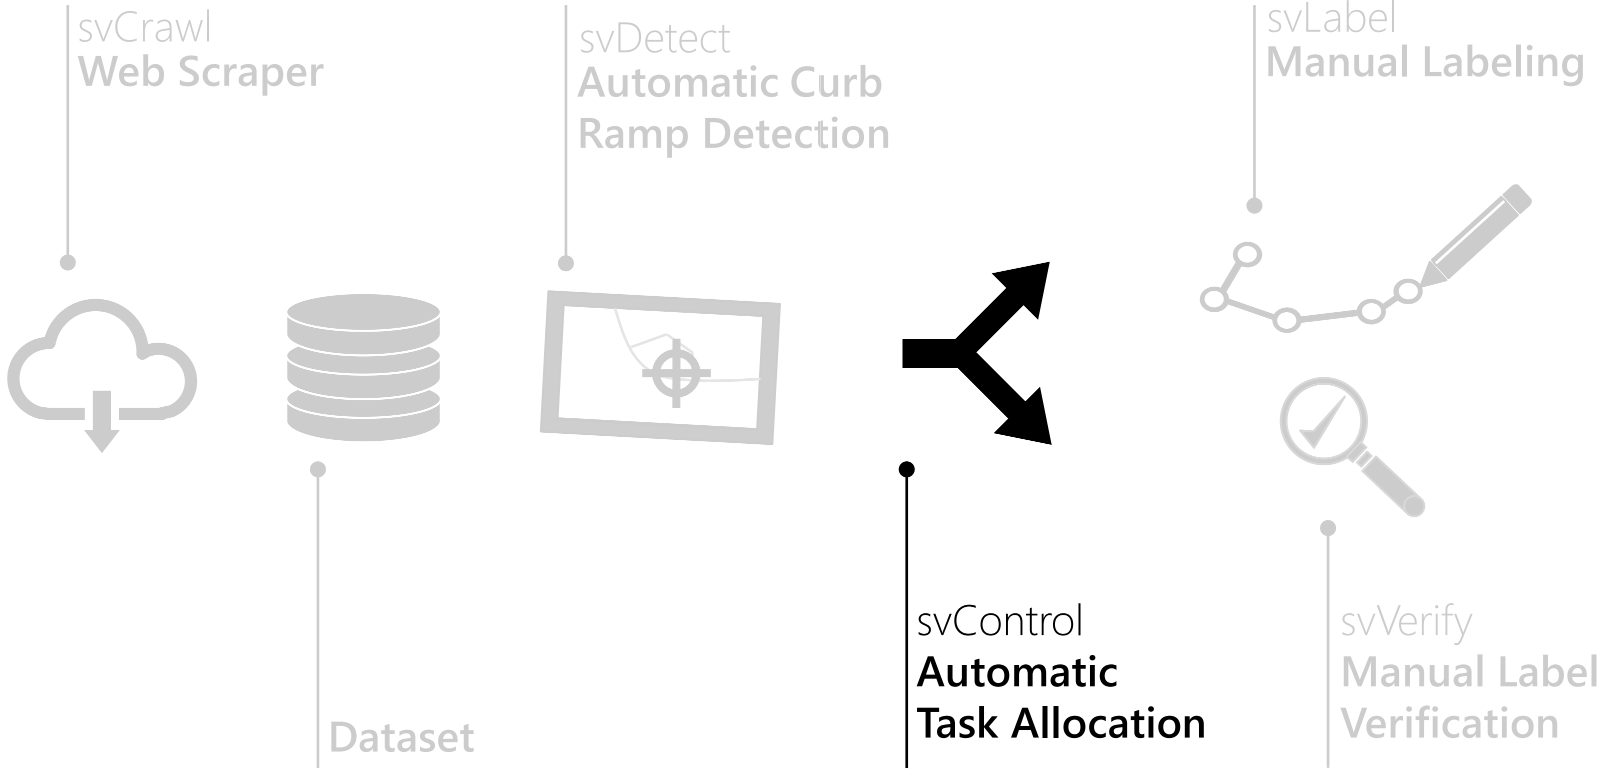

Tohme

遠目·*Remote Eye*

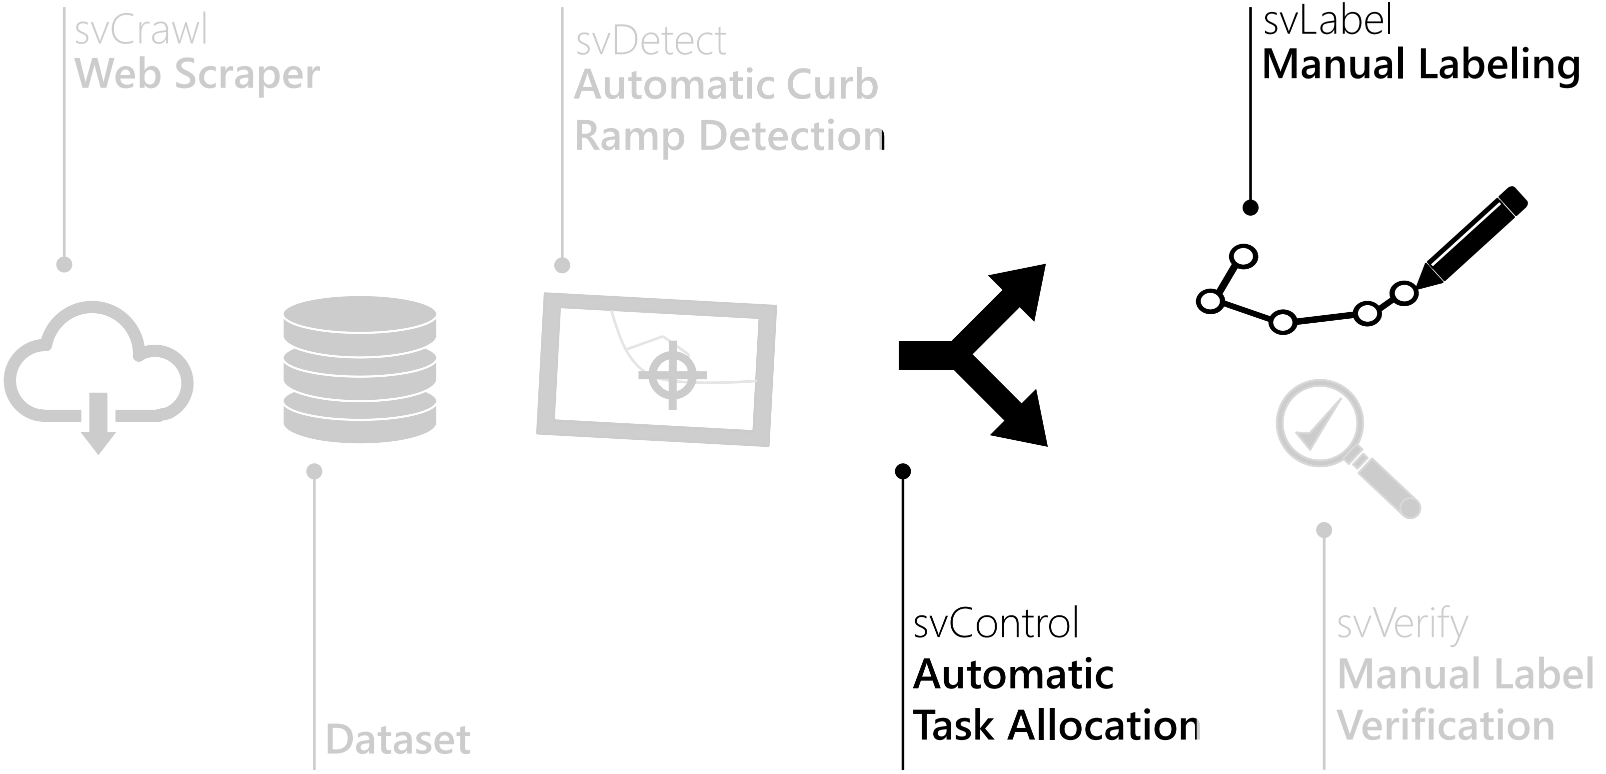

Manual Labeling | Labeling Interface

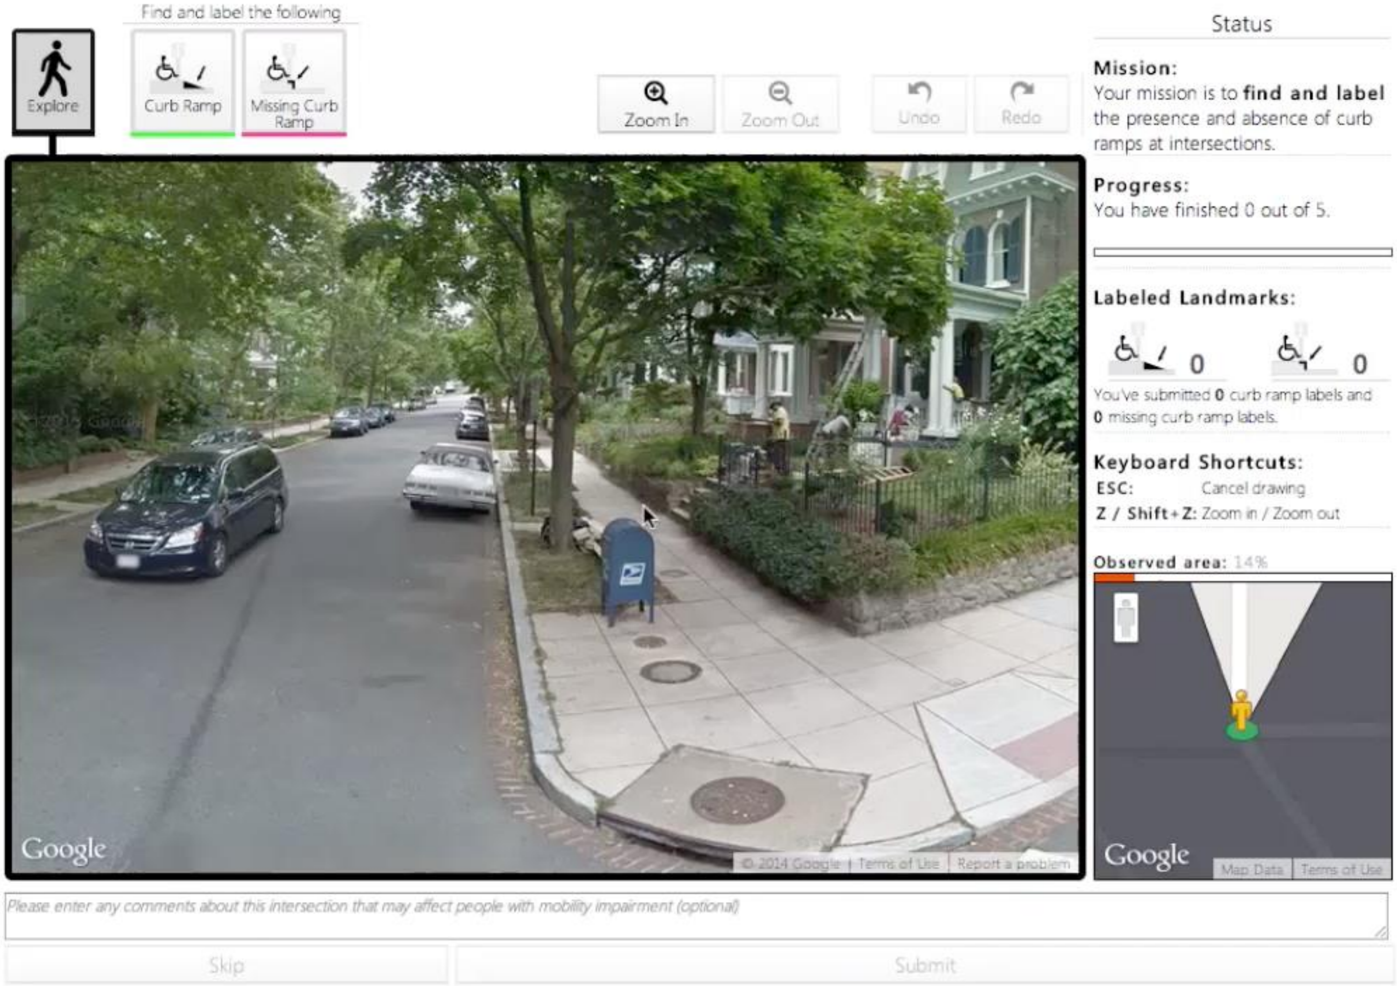

3x

Manual Labeling | Interactive Tutorial

Help Us Improve Street Accessibility

Hi, we're exploring new ways to find accessibility problems in cities, and we need your help! In this task, **your mission is to label curb ramps** and **missing curb ramps** in Google Street View. Curb ramps are very important--without them, people in wheelchairs cannot move about the city.

An image of curb ramps at an intersection.

A lack of a curb ramp at this corner obstruct wheelchair users from getting on and off the sidewalk.

We'll **begin with a short, interactive tutorial** to get you started! Thanks for your help in improving the accessibility of cities.

Next

4x

Manual Labeling | Golden Insertion

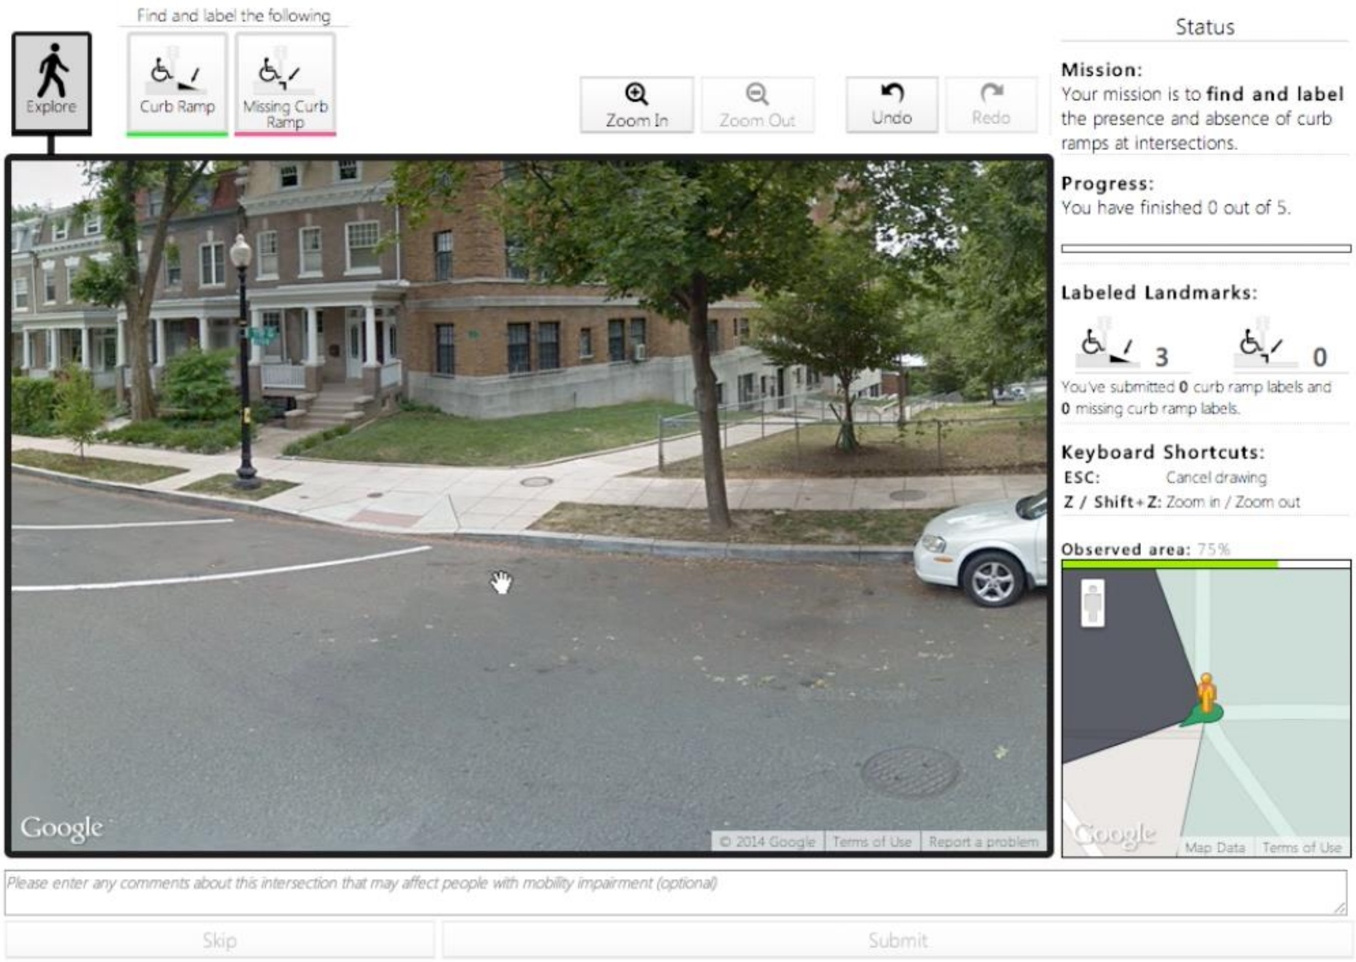

2x

Tohme
遠目·*Remote Eye*

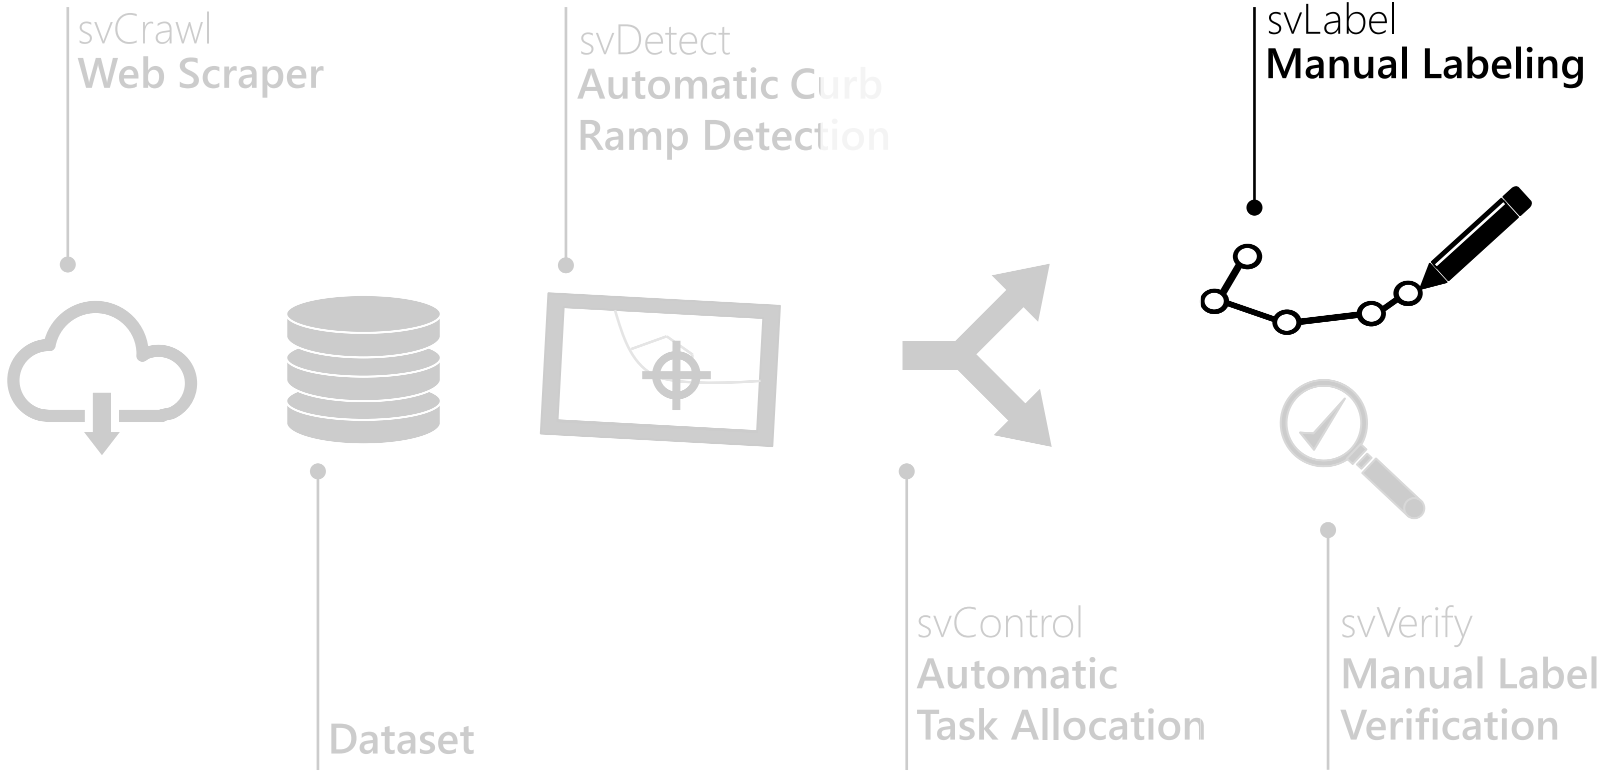

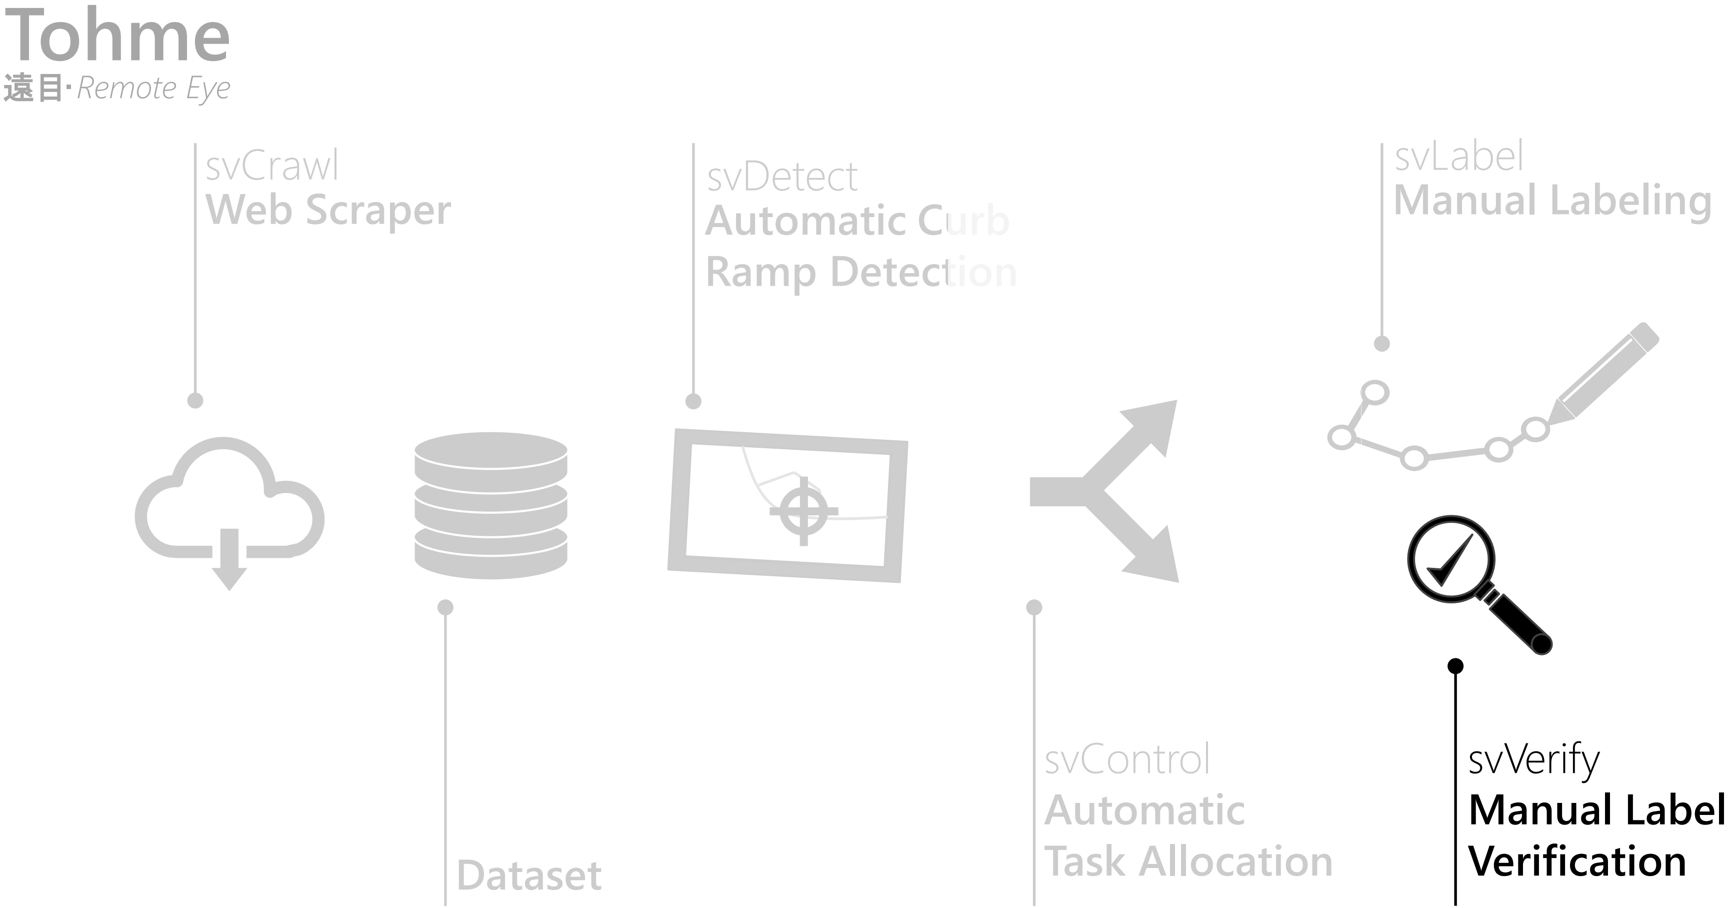
Tohme
遠目· Remote Eye
svCrawl
Web Scraper
svDetect
Automatic Curb
Ramp Detection
svLabel
Manual Labeling
Dataset
svControl
Automatic
Task Allocation
svVerify
Manual Label
Verification

Manual Label Verification

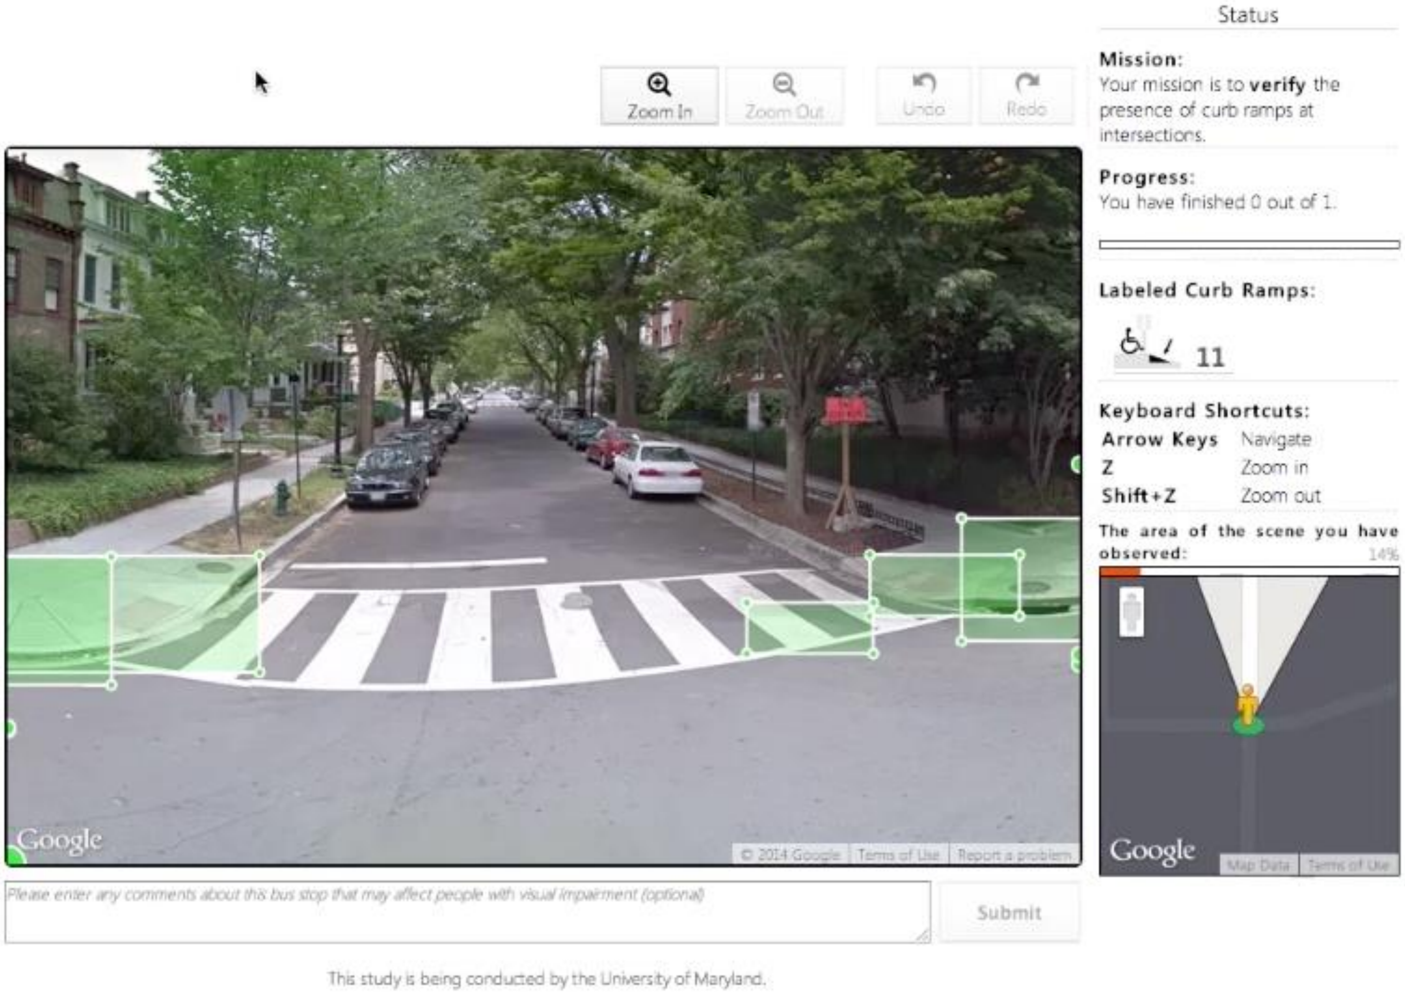

2x

Manual Label Verification

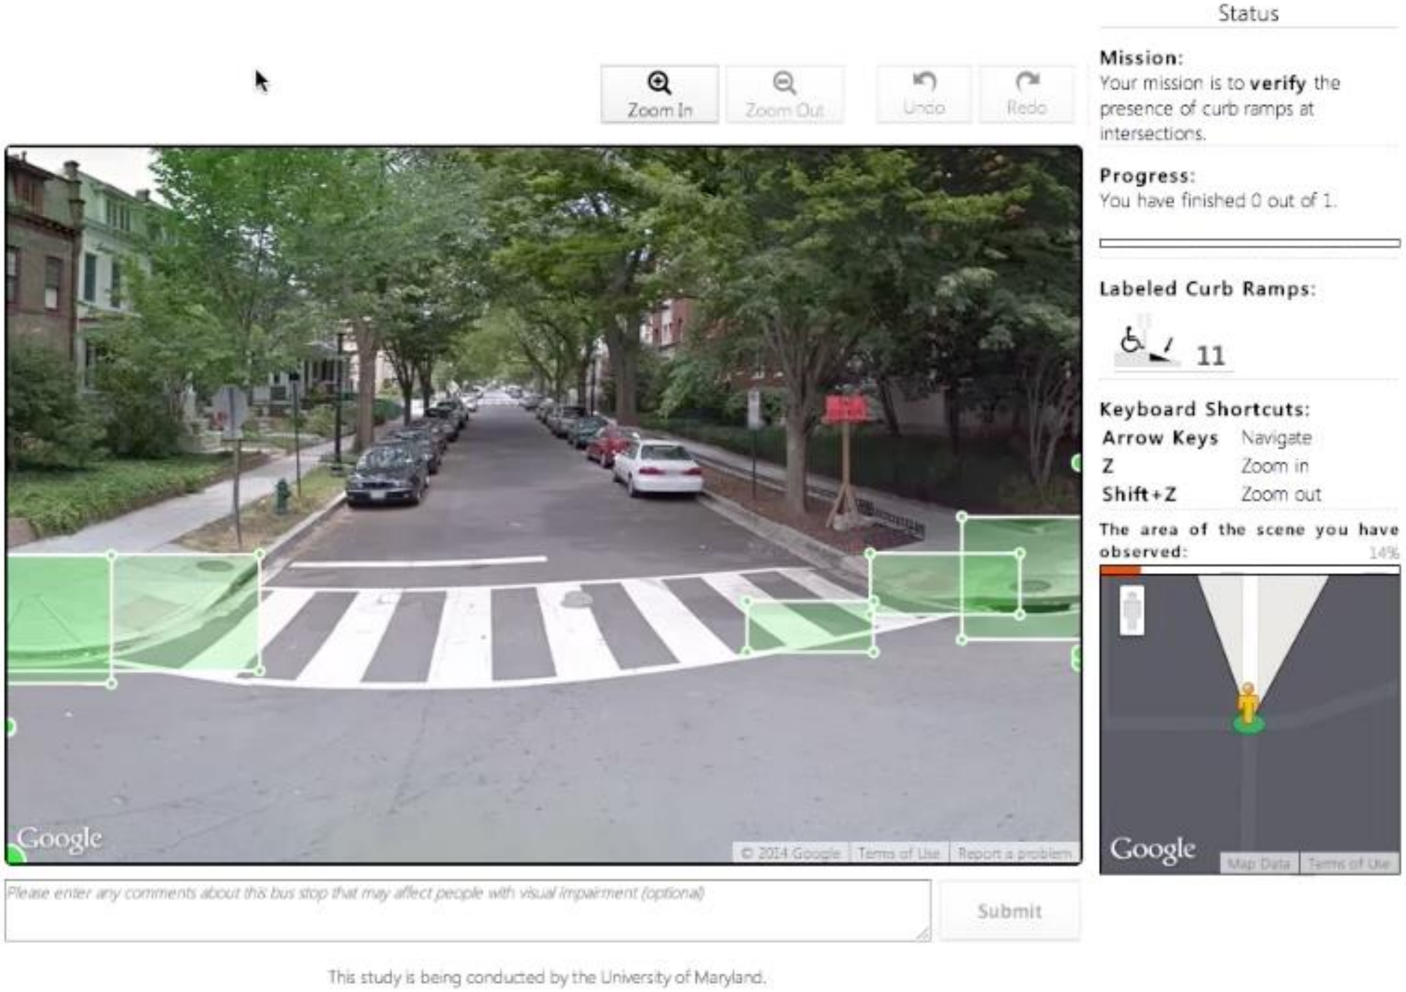

2x

Study Method: Conditions

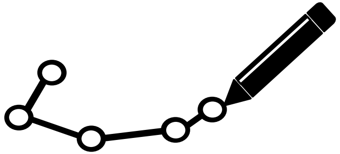

vs.

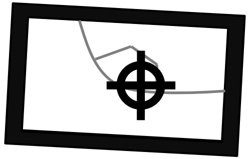

&

vs.

Tohme

遠目·*Remote Eye*

Manual labeling without smart task allocation

CV + Verification without smart task allocation

Study Method: Measures

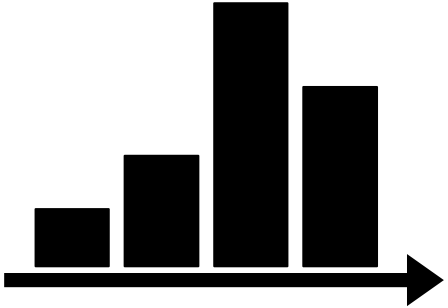

Accuracy

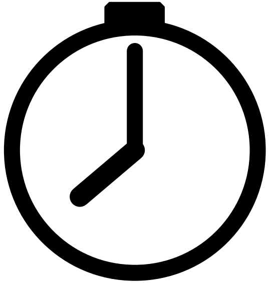

Task Completion Time

Study Method: Approach

We recruited workers from Amazon Mechanical Turk to work on labeling tasks and verification tasks

We used 1,046 GSV images

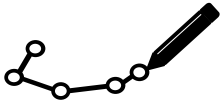

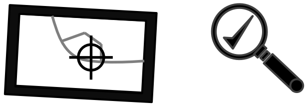

	Labeling Tasks	**Verification Tasks**
# of distinct turkers:	242	161
# of tasks completed:	6,350	4,820

We evaluated the result with Monte Carlo simulation

Evaluation | Labeling Accuracy and Time Cost

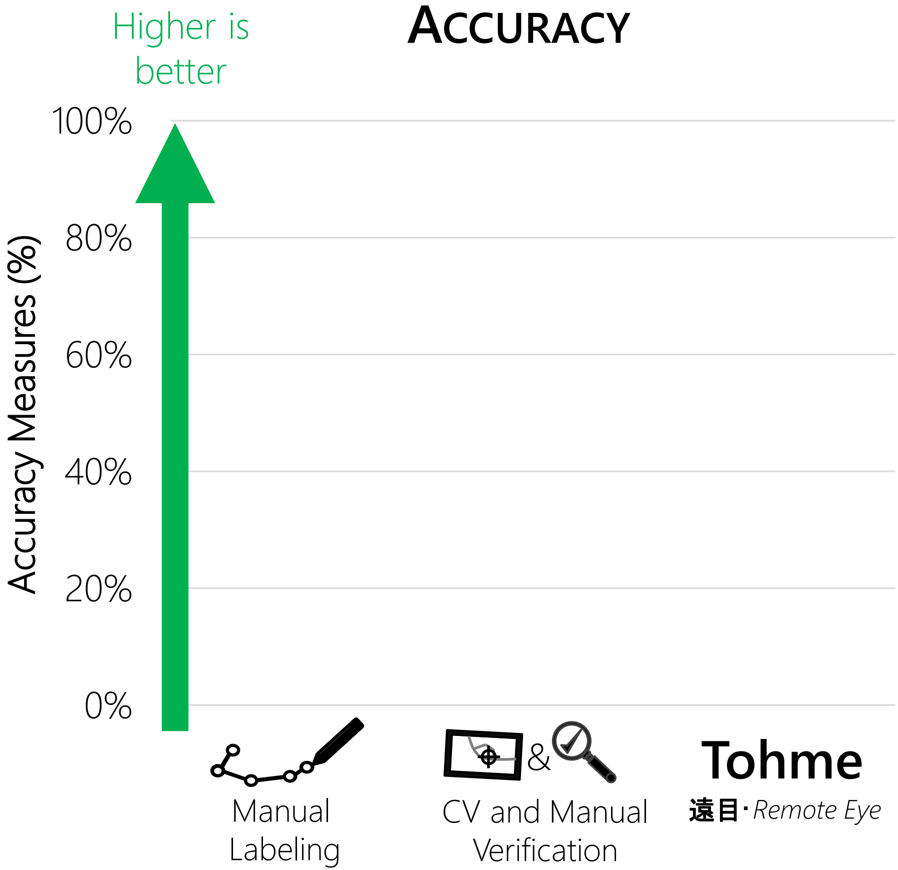

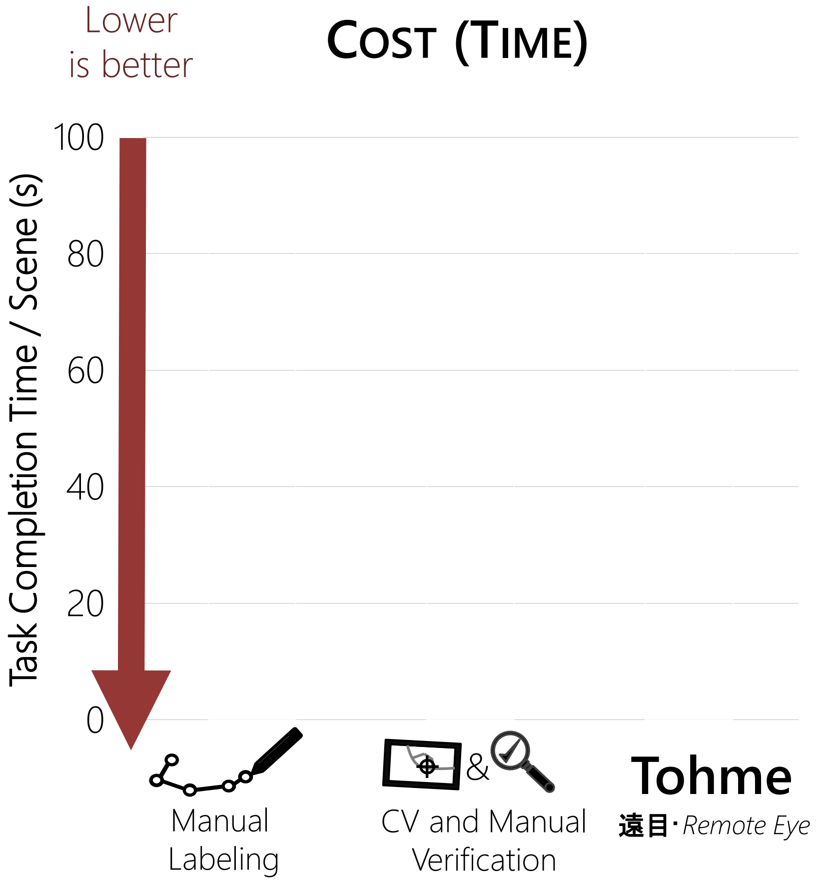

Error bars are standard deviations.

Evaluation | Labeling Accuracy and Time Cost

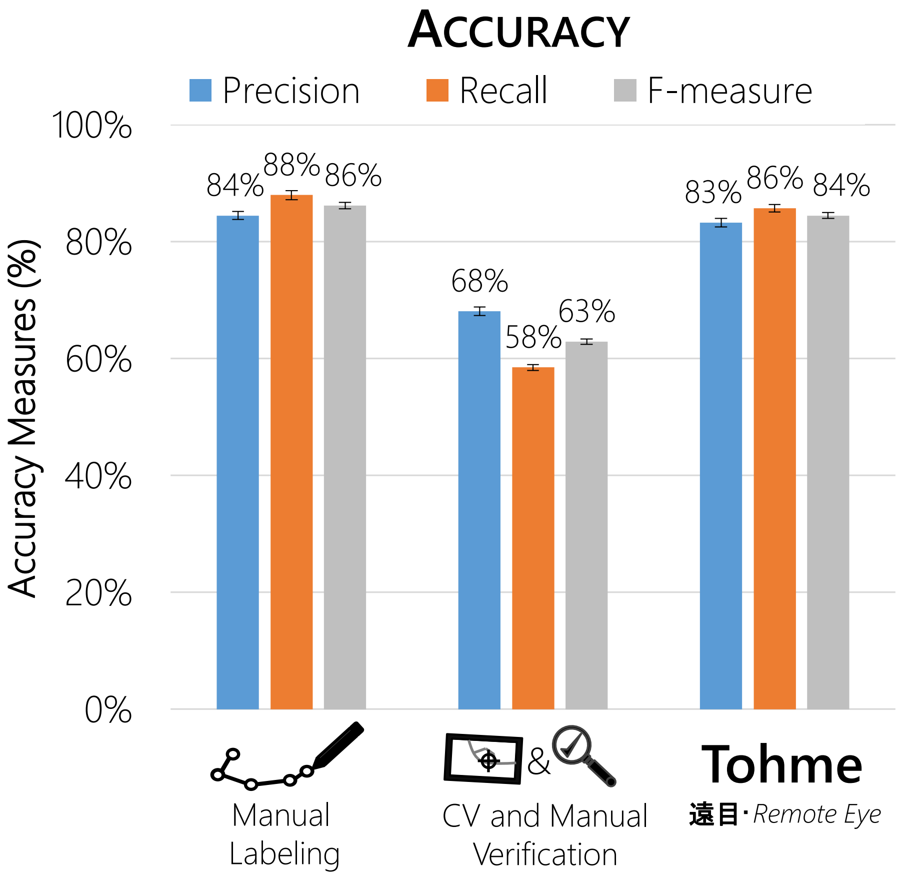

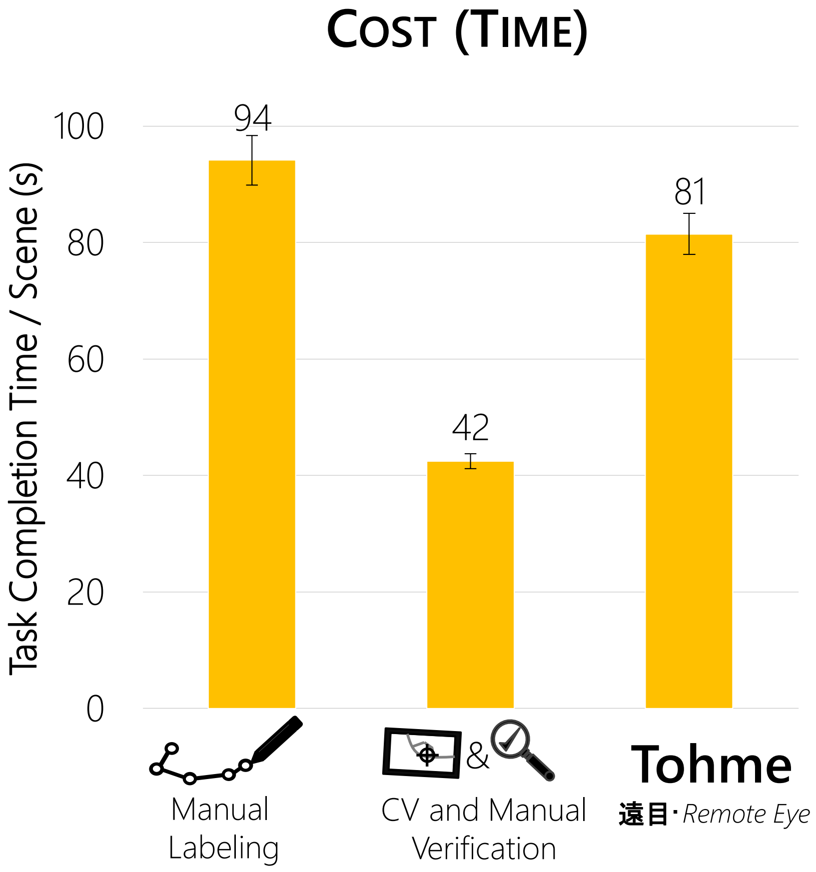

Error bars are standard deviations.

Evaluation | Labeling Accuracy and Time Cost

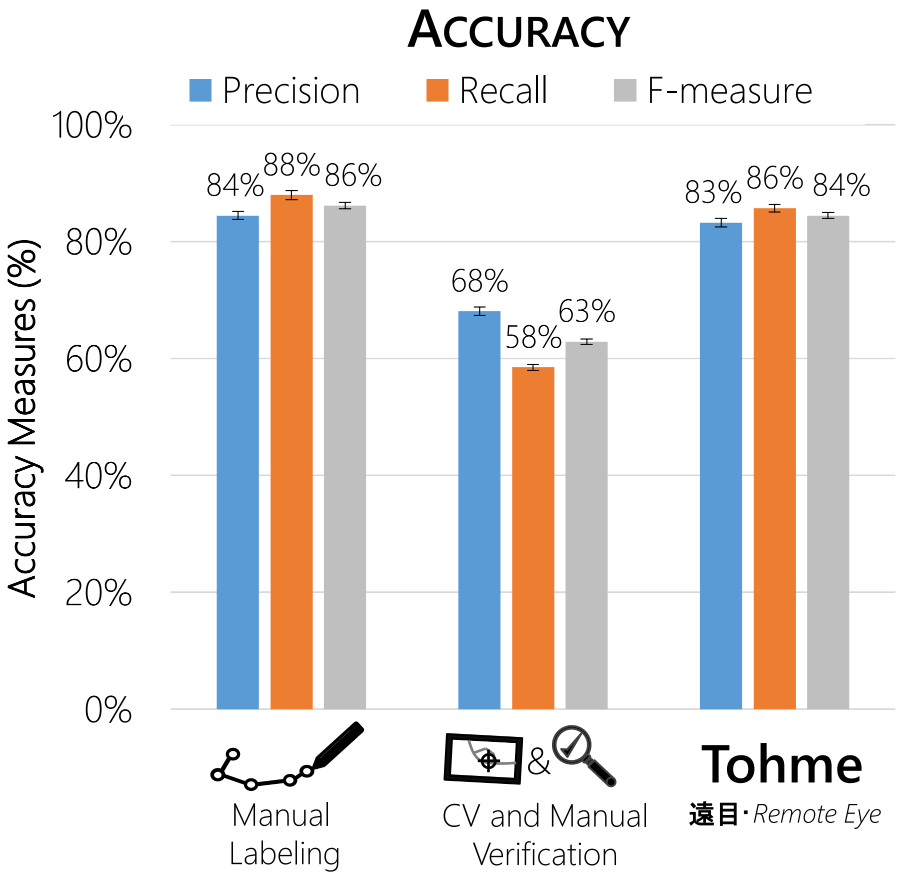

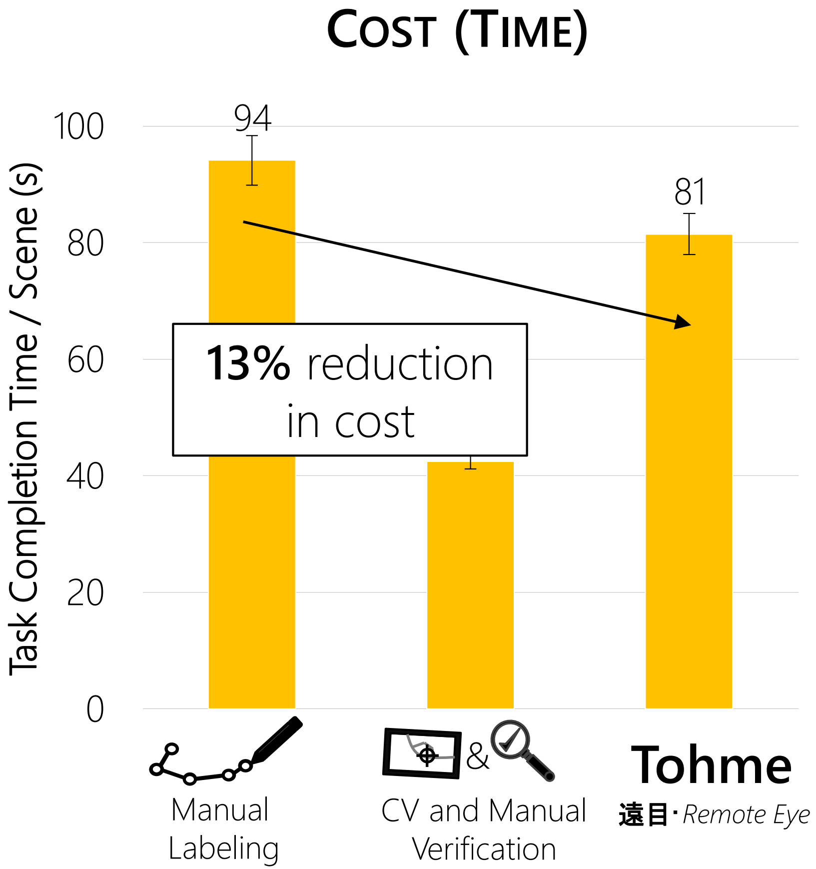

Error bars are standard deviations.

Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

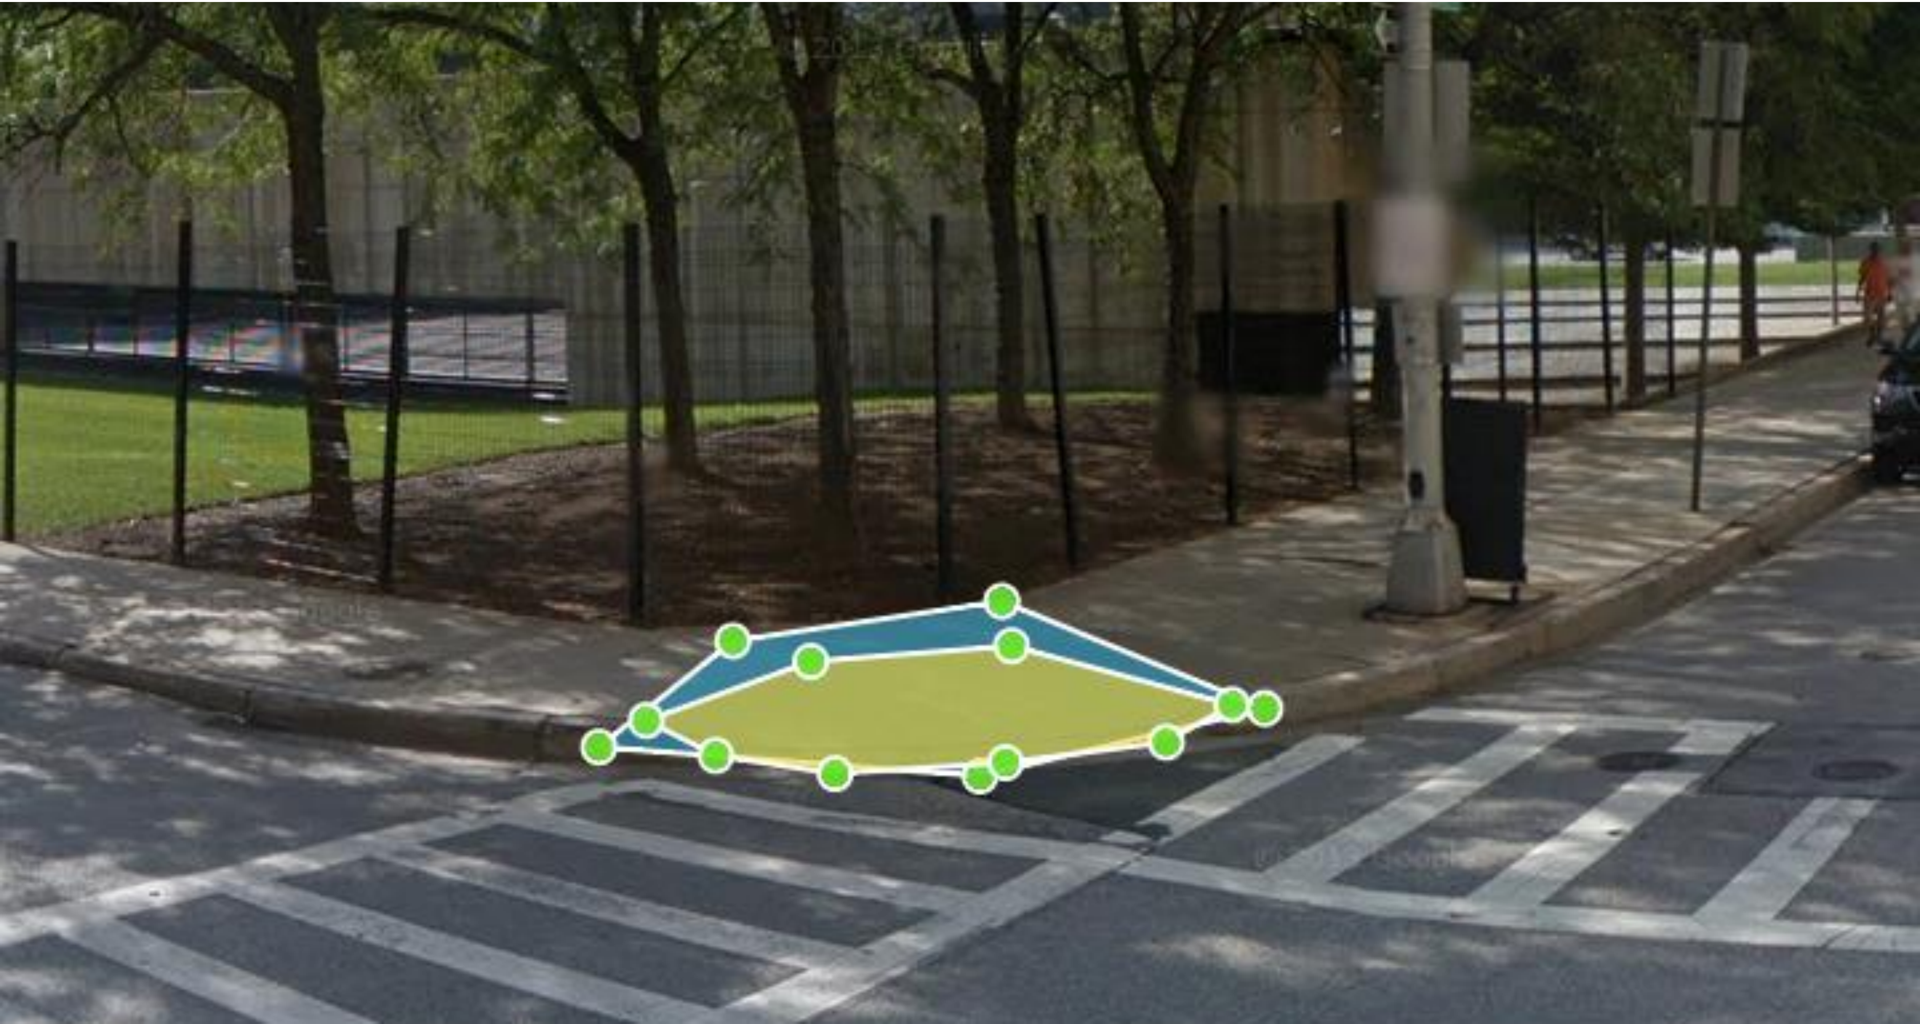

Evaluation | Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

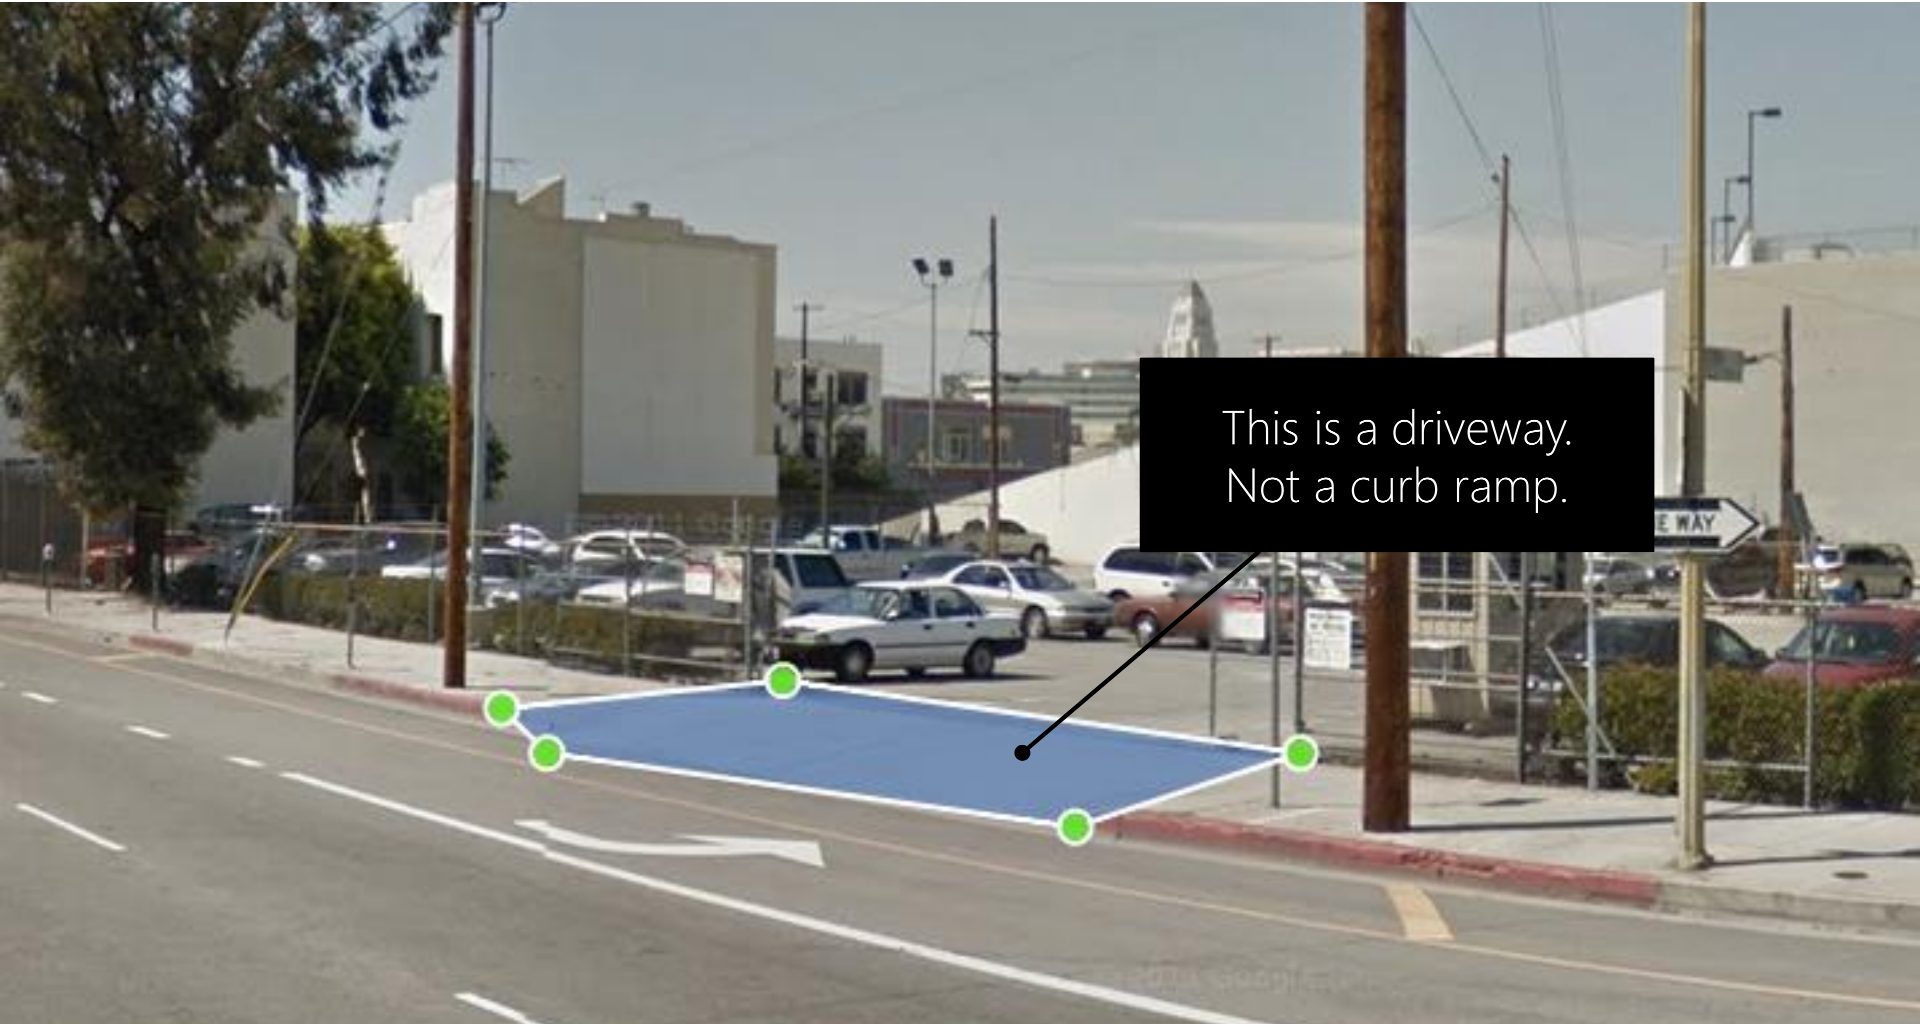

Evaluation | Example Labels from Manual Labeling

Evaluation | Example Labels from Manual Labeling

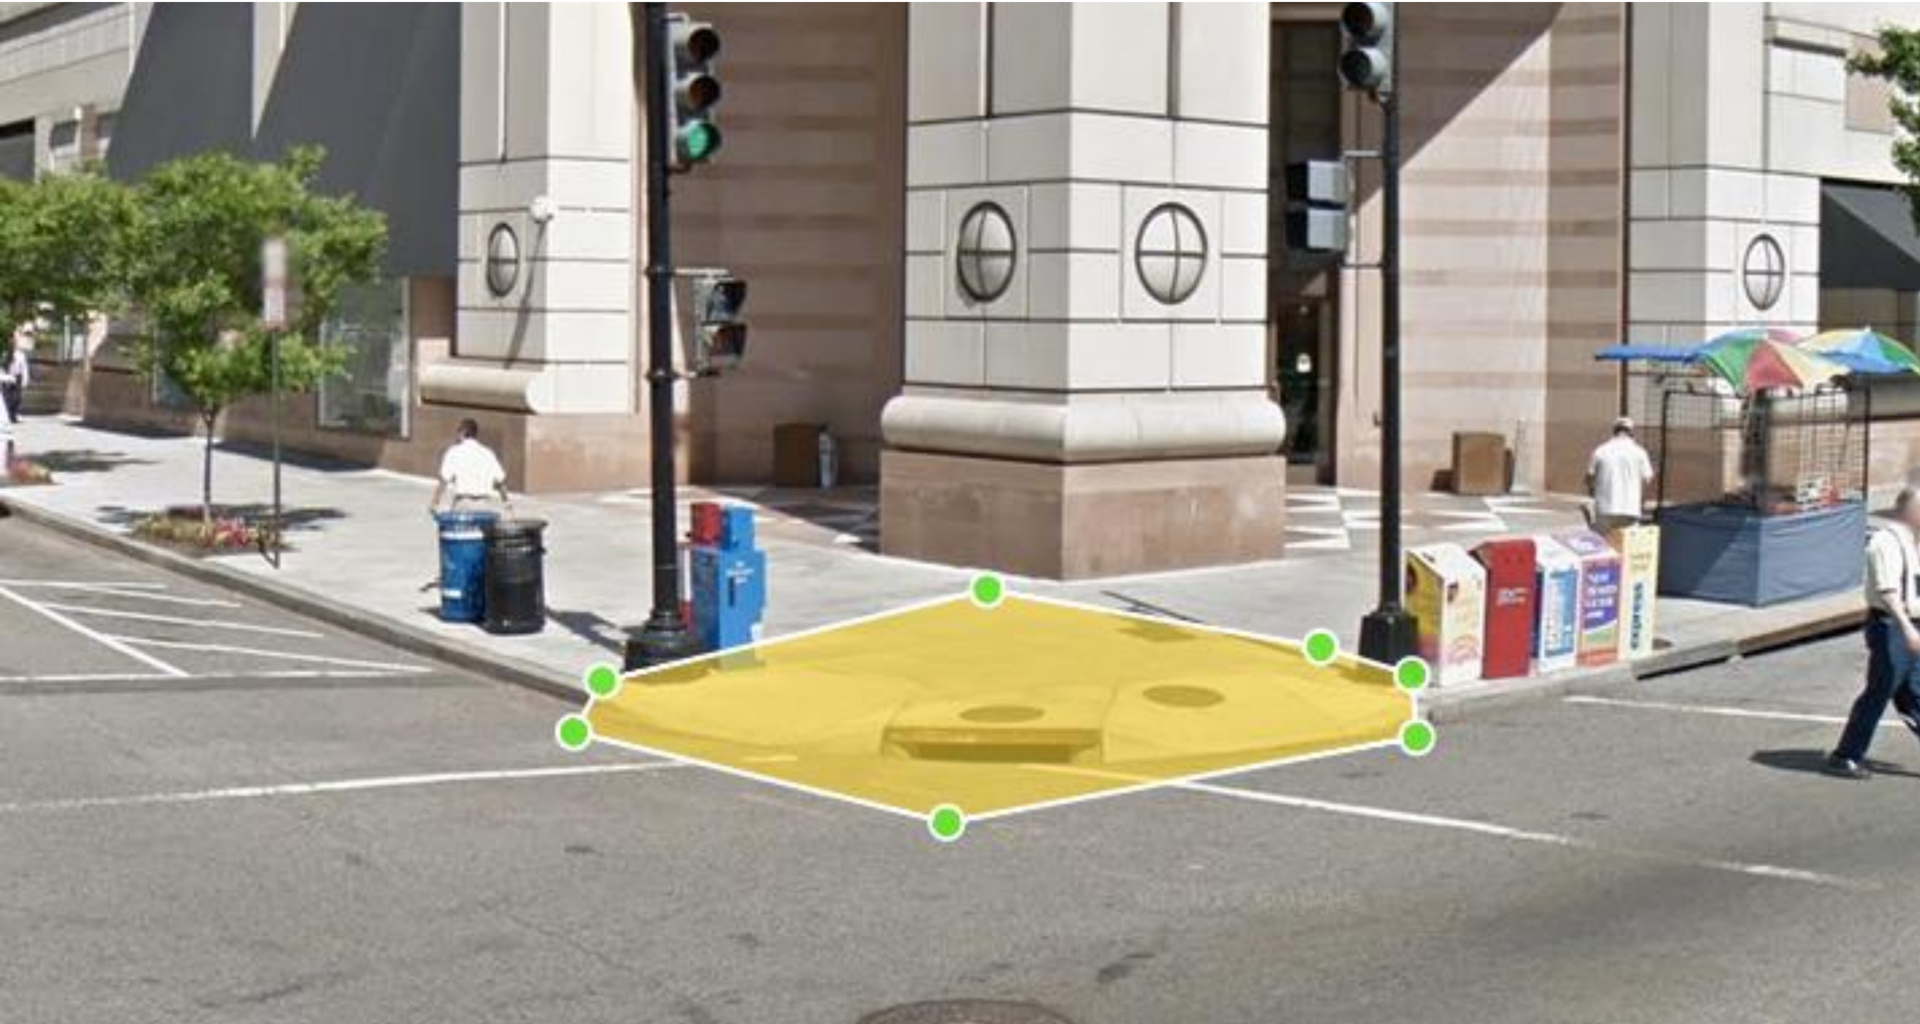

Examples Labels from CV + Verification

Raw Street View Image

Evaluation | Example Labels from CV + Verification

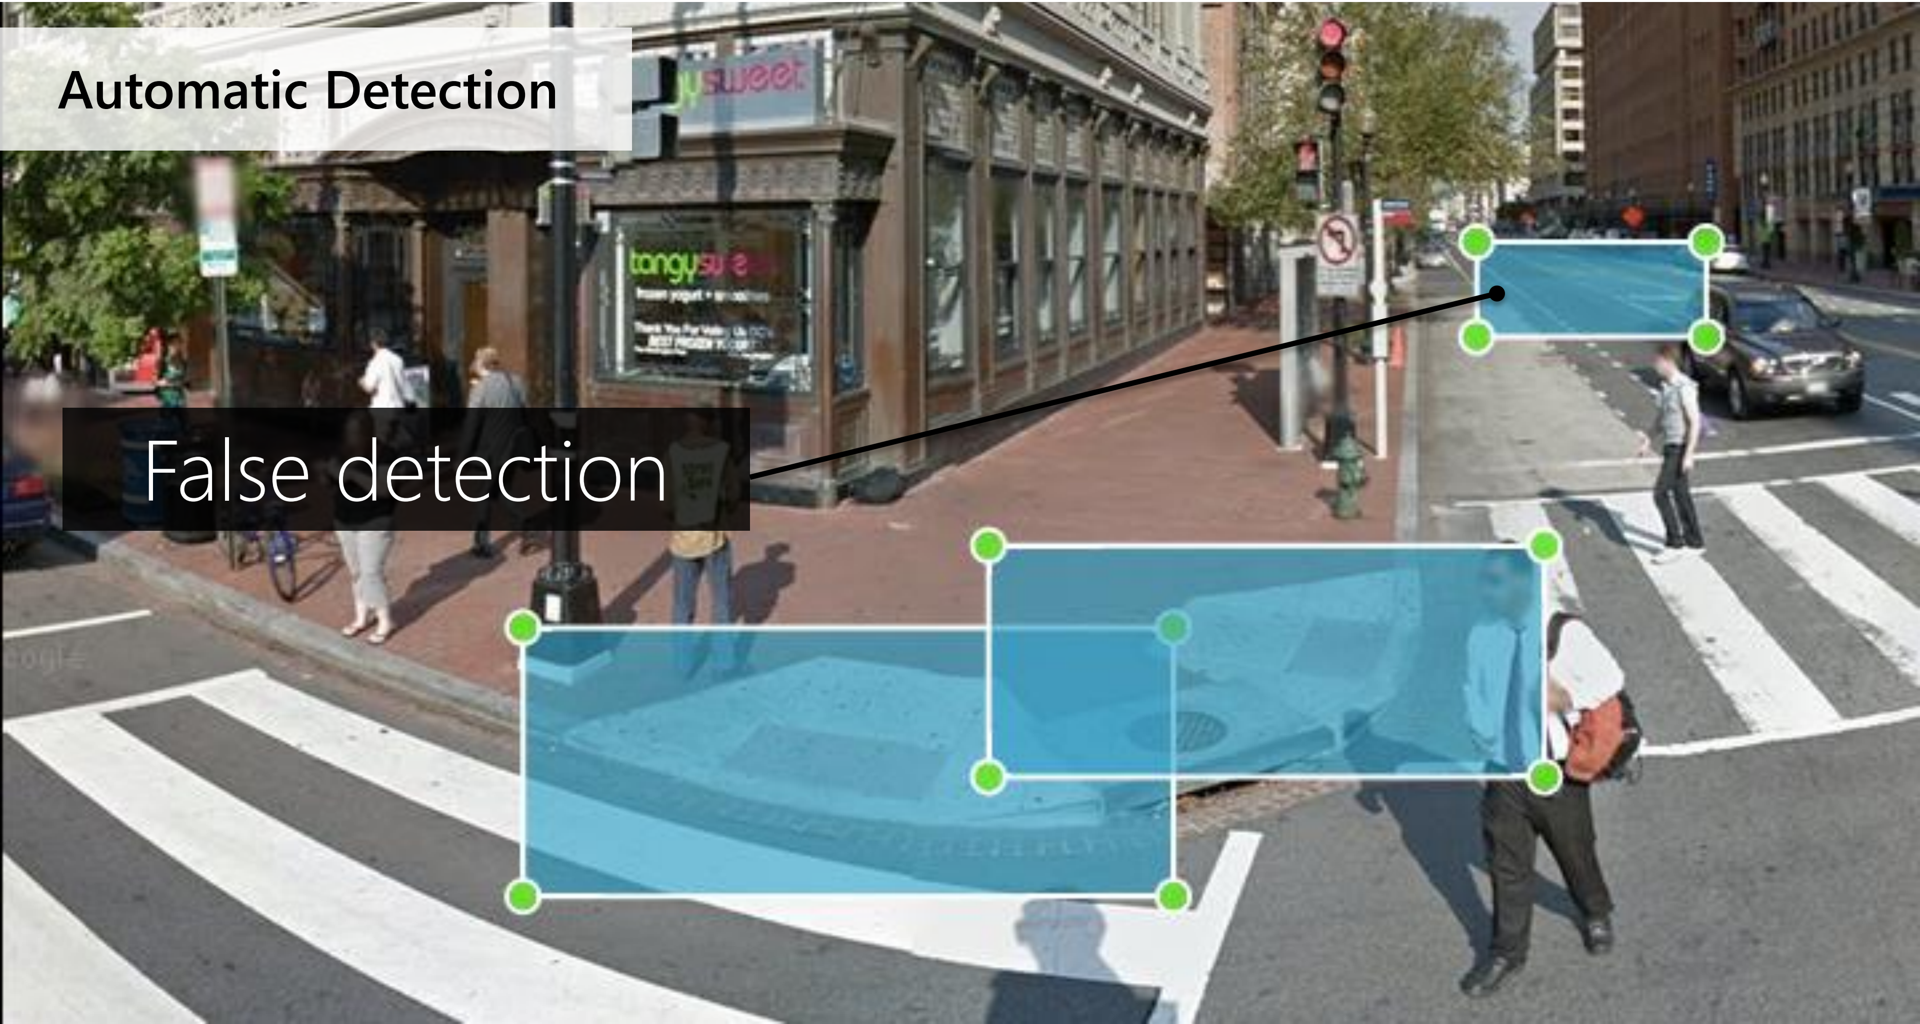

Automatic Detection + Human Verification

Summary

We developed a method that combines crowdsourcing and computation that **increased accessibility data collection efficiency without losing accuracy**

Back of Envelope Calculation

How long would it take to audit a city?

DC Total Street Distance: **1,238 mi**

Audit Speed: **7.9 mi / hour**

Estimated Total Time-cost: **157 hour ($1.1k)**

Estimated Total Time-cost with automation: **137 hour**

I think we can do better than this ☺

Future Work: Improving Detection Accuracy

Context integration & scene understanding

Using 3D-data for mensuration

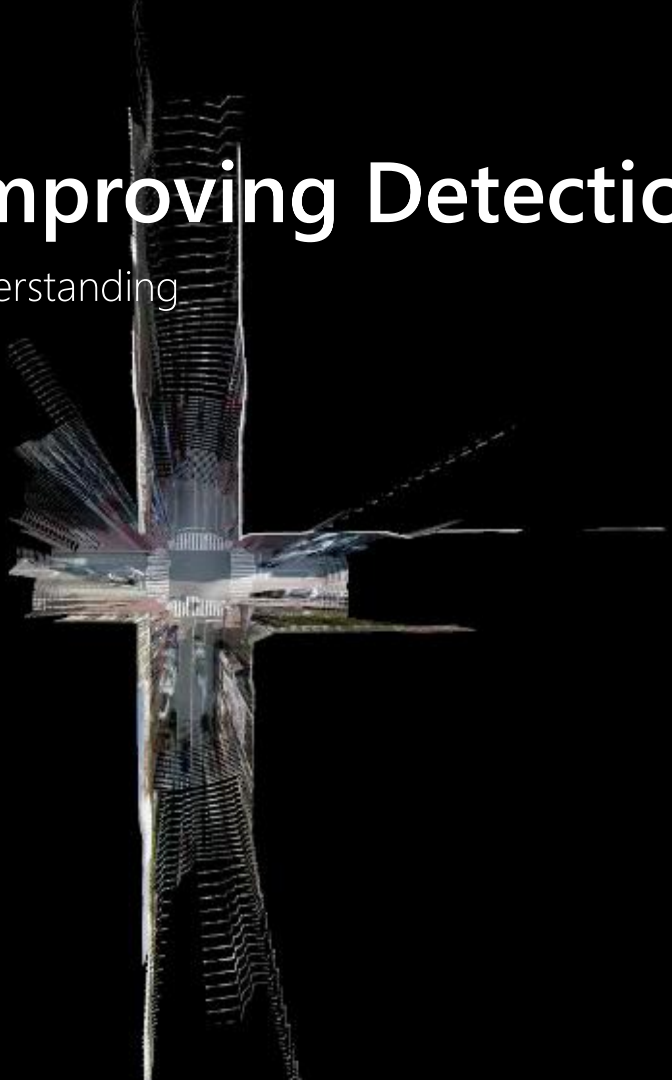

Future Work: Reacting to Changes

Using image dataset that is updated frequently, can we identify dynamic accessibility features like constructions?

Urthecast

Future Work: Research Beyond Sidewalk Accessibility

- Bus stop accessibility
- Urban vegetation
- Public health (*e.g.*, cleanliness)
- Urban planning (*e.g.*, bicycle roads)

Hara *et al.* [ASSETS 2013, TACCESS 2015]

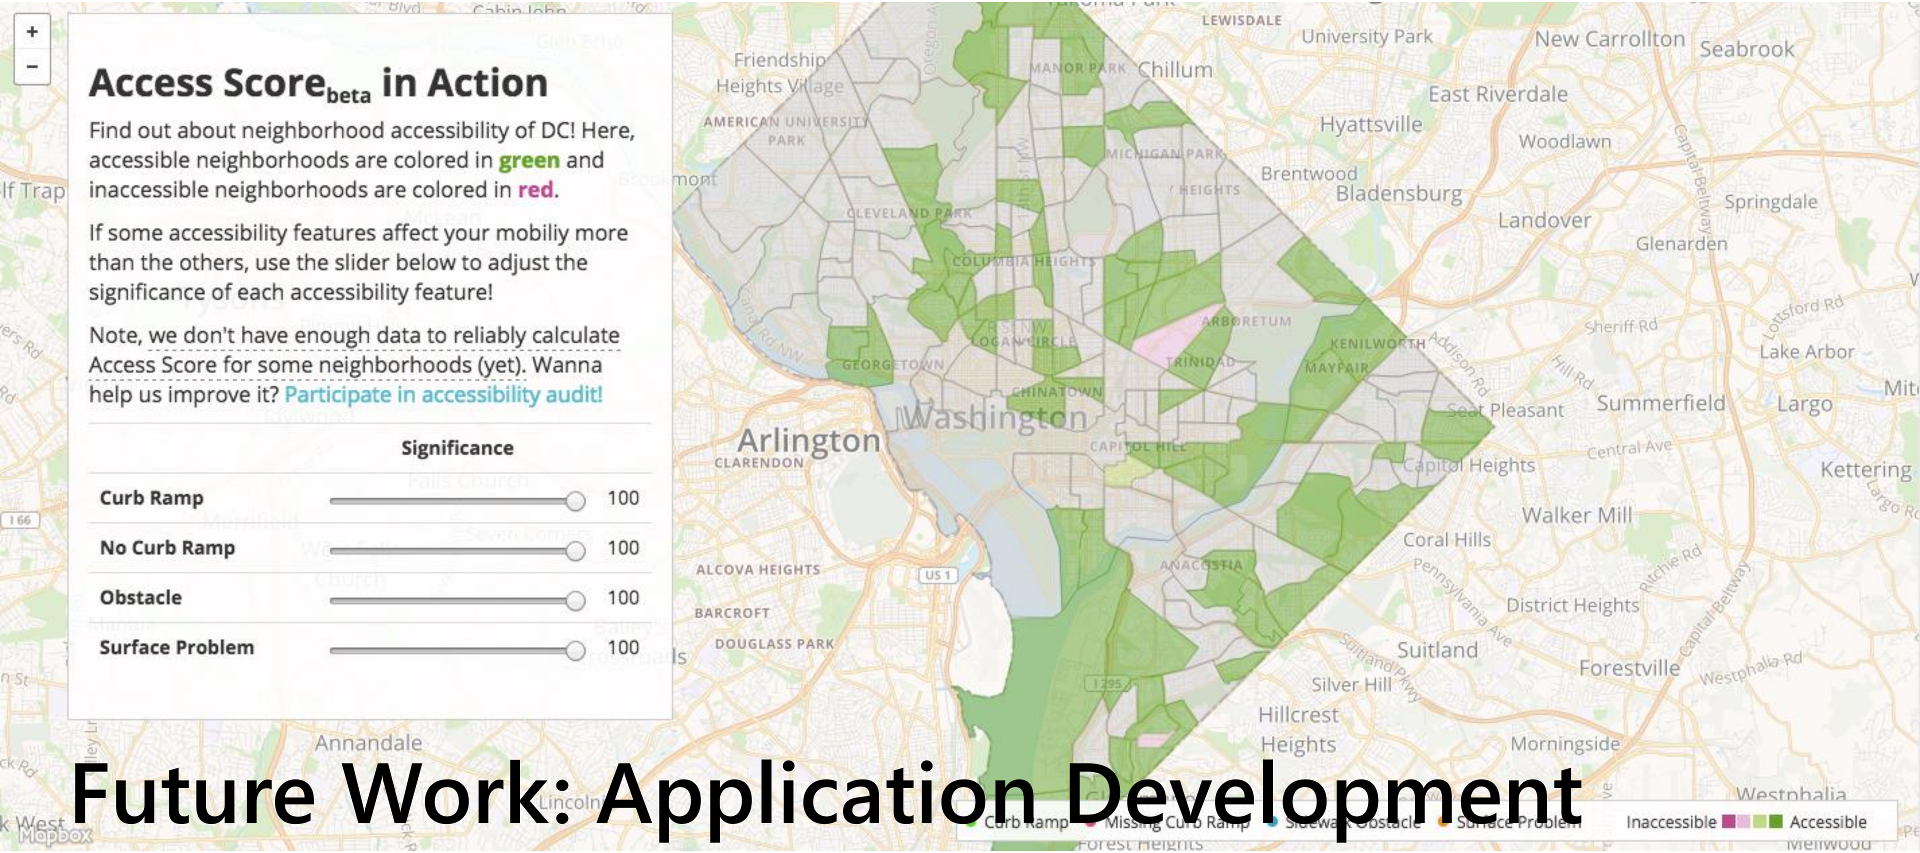

Future Work: Application Development

sidewalk.umiacs.umd.edu/demo

Collaborators

Advisor

Jon E. Froehlich

Professors | Researchers

Shiri Azencot and David Jacobs

Students

Cynthia L. Bennett, Megan Campbell, Christine Chan, Jonah Chazan, Vicki Le, Anthony Li, Kelly Minckler, Zachary Lawrence, Robert Moore, Rochelle H. Ng, Sean Pannella, Niles Rogoff, Manaswi Saha, Soheil Sehnezhad, Jin Sun, and Alex Zhang,

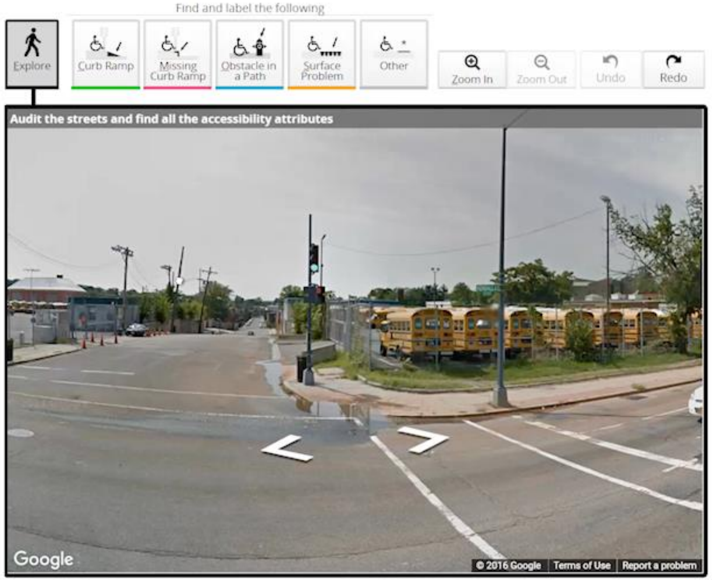

sidewalk.umiacs.umd.edu

Thank you!

Kotaro Hara | @kotarohara_en

Carnegie Mellon University

UNIVERSITY OF MARYLAND

COMPUTER SCIENCE UNIVERSITY OF MARYLAND

HCIL

makeability lab