

# APE: An Open and Shared Annotated Dataset for Learning Urban Pedestrian Path Networks

Yuxiang Zhang, Nicholas-J Bolten, Sachin Mehta, and Anat Caspi  
University of Washington  
Seattle, Washington, USA

## ABSTRACT

Inferring the full transportation network, including sidewalks and cycleways, is crucial for many automated systems, including autonomous driving, multi-modal navigation, trip planning, mobility simulations, and freight management. Many transportation decisions can be informed based on an accurate pedestrian network, its interactions, and connectivity with the road networks of other modes of travel. A connected pedestrian path network is vital to transportation activities, as sidewalks and crossings connect pedestrians to other modes of transportation. However, information about these paths' location and connectivity is often missing or inaccurate in city planning systems and wayfinding applications, causing severe information gaps and errors for planners and pedestrians. This work begins to address this problem at scale by introducing a novel dataset of aerial satellite imagery, street map imagery, and rasterized annotations of sidewalks, crossings, and corner bulbs in urban cities. The dataset spans  $2,700\text{km}^2$  land area, covering select regions from 6 different cities. It can be used for various learning tasks related to segmenting and understanding pedestrian environments. We also present an end-to-end process for inferring a connected pedestrian path network map using street network information and our proposed dataset. The process features the use of a multi-input segmentation network trained on our dataset to predict important classes in the pedestrian environment and then generate a connected pedestrian path network. Our results demonstrate that the dataset is sufficiently large to train common segmentation models yielding accurate, robust pedestrian path networks.

## KEYWORDS

dataset, accessibility, computer vision, pedestrian, urban network

## 1 INTRODUCTION

A connected pedestrian path network detailing the location and connectivity of sidewalks, crossings, and curbs is essential to any transportation network. Sidewalks and crossings connect all other transportation modes and are the key to building an accessible city. Creating a fully connected pedestrian path network is essential for many city planning tasks and wayfinding applications [7].

While automobile roads have been extensively mapped [17, 18], mapping information for the paths that serve pedestrians is incomplete or missing. This fundamental lack of information creates bottlenecks in the transportation network and artificially disconnects neighborhoods. More importantly, this information gap impacts pedestrians with disabilities' travel in the cities. A connected pedestrian path network map consisting of the location and connectivity

of sidewalks and crossings would close this information gap and improve accessibility for all.

A pedestrian map requires different information than a map for car navigation. It must show common paths (sidewalks, crossings), their connectivity, transitions, and additional attributes of those paths. Altogether this is difficult to gather manually due to the large amount of data and potential for error [6, 7]. In addition, mappings that are done by human mappers often contain errors due to many factors, including low image quality and ambiguous mapping protocols. We validated and evaluated several areas mapped through a human annotation campaign carried out by the OpenSidewalks Project [4]. We found that although mappers were trained before doing the mapping and were provided with instructions and mapping protocol, different types of errors were still made in the mapping process. Example errors are visualized in Figure 1, and Table 1 summarizes our analysis on human-generated maps, where we analyzed the three elemental geometry features (sidewalks, crossings, and sidewalk links that connect sidewalks to crossings)<sup>1</sup>. These geometry-based errors can lead to inaccurate depictions of pedestrian paths and internal path connectivity, hindering downstream applications and particularly impacting routing algorithms and network analysis. Machine learning methods that globally generate connected pedestrian path networks efficiently would improve mapping accuracy, reduce manual efforts, and improve transportation planning.



**Figure 1: Example of annotations from human mappers. Red text shows different types of features in a pedestrian path network map. Yellow boxes show examples of errors made by human mapper: (a) crossing lines are not connected to the sidewalk lines with link lines (b) the curb nodes are falsely mapped to the middle of the road.**

<sup>1</sup>The nodes that are analyzed in Table 1b include the curb nodes (the intermediate nodes shared by a crossing line and a link line that connects a crossing to a sidewalk), and the crossing nodes (the nodes at the intersection of a crossing line and a road line). Sidewalks are mapped as lines in common GIS data thus there is no designated tag for sidewalk nodes in the mapped annotations.

\*Preprint. Under review.

†Corresponding author: Yuxiang Zhang (email: yz325@uw.edu)

**Table 1: Analysis on human annotations**

(a) Statistics of validated annotations			
	Annotation count	Annotation length	
links	1526	9423.9 m	
sidewalks	3537	37527.1 m	
crossings	1594	8543.9 m	

(b) Number of errors on Nodes			
Node type	Location	Classification	Tag
curb	301	108	1
crossing	51	35	226

(c) Number of errors on LineStrings			
Line type	Location	Classification	Tag
links	337	N/A	59
sidewalks	396	N/A	59
crossing	375	15	0

In this work, we present a novel dataset, the Annotations for Pedestrian Environment (APE) dataset, created through a rasterization process of GIS mapping annotations. The APE dataset includes aerial satellite images, street map tiles, and rasterized annotations. The contributions of this work are two folds (1) we present a method for rasterizing Geographic Information System (GIS) mapping annotations to generate the APE dataset (Figure 2). (2) We also develop an end-to-end process (Figure 3) to infer a connected pedestrian path network, using a multi-input segmentation network trained on the APE dataset to predict pedestrian path locations and integrate them with street network data to complete a pedestrian path network. The APE dataset and code will be released upon publication.

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 introduces the APE dataset. Section 4 details the implementation of our process. Experiment results are shown in Section 5. Discussions and conclusions are made in Section 6.

## 2 RELATED WORK

We discuss related work in three domains: Section 2.1: studies on street/road network mapping, Section 2.2: methods to map pedestrian environments, Section 2.3: available aerial imagery datasets. The combined work highlights the need for an imagery-based dataset that targets the pedestrian environment.

### 2.1 Street and Road Network Map Generation

Studies have investigated using aerial imagery along with auxiliary data for street mapping. Wu et al. [28] used OpenStreetMap (OSM) centerlines as labeled data and extracted roads from very high-resolution (VHR) satellite images. Sun et al. [24] added crowd-sourced global positioning system (GPS) data to satellite images to extract roads with CNN-based semantic segmentation. Zhou et al. [30] fused remote sensing images and GPS for road detection

and extraction. Additional recent learning-based studies included Lu et al. [16] proposing a multi-scale residual network for road detection, Pan et al. [21] proposing a fully convolutional neural network using VHR remote sensing. Mattyus et al. [17] estimated road topology from aerial images, Mi et al. [18] generated road lane graphs from LiDAR data with a hierarchical graph generation model. Importantly, work in this domain solely focuses on automobile road detection and extraction, and does not address the generalization or extension of the proposed methods to the pedestrian environment. Methods for mapping the environments that serve pedestrians’ travel have not been widely studied.

### 2.2 Pedestrian Environment Mapping

Few studies have focused on mapping pedestrian environments compared to automobile roads. Karimi et al. [13] explored pedestrian map generation approaches in a small-scale area, demonstrating preliminary mapping results that heavily depended on the availability and quality of input data. Recent advancements in remote sensing have led to more imagery-based approaches, such as Ahmetovic et al. [1] detection of zebra crossings using satellite imagery and validation with street-level images. Likewise, Ghilardi et al. [9] classified and located crosswalks using an SVM classifier over data extracted from road maps, and Ning et al. [19] extracted sidewalks from aerial images with a neural network and restored occluded segments from street view images. These studies improved pedestrian environment mapping, but they do not generate a comprehensive, connected pedestrian path network needed for city planning and navigation. Other studies use on-the-ground data, such as Zhang et al. [29] automated collection of street-view images with auxiliary data to map sidewalk connectivity and sidewalk infrastructure., Hou et al. [10] extraction of sidewalks using LiDAR data and point cloud segmentation. However, these methods often require physical systems to cover a large area to generate a pedestrian path network. Other studies that map pedestrian path networks at scale are often based solely on existing street (road) data. For example, Li et al. [14]’s semi-automated method generated a sidewalk network using parcel-level data and roadway centerline, but it required human editing for quality control.

### 2.3 Aerial Imagery Datasets

There are several datasets containing aerial imagery for mapping. The TorontoCity dataset [26] contains aerial satellite images for road curb extraction and road centerline estimation. PRRS [2] presents a dataset for building extraction and Digital Surface Model (DSM) estimation using satellite data. In addition, the DeepGlobe dataset [8] and the ISPRS dataset [11] contain imagery and annotations for tasks including road extraction, building detection, and land cover classification. These datasets enable researchers to study different tasks involving the use of aerial imagery data, but they neither directly nor specifically target the pedestrian path networks, nor provide annotations needed for semantically understanding the pedestrian environment. In sum, currently, there are no extensive datasets focused on pedestrian environments, nor methods to generate pedestrian path networks for large areas.

### 3 THE ANNOTATIONS FOR PEDESTRIAN ENVIRONMENT DATASET

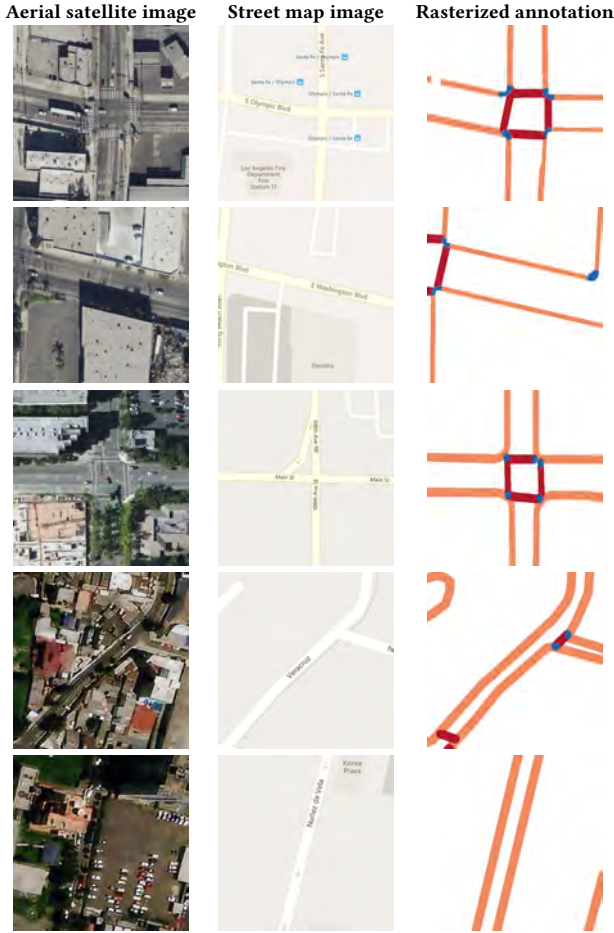


Figure 2: Samples from the APE dataset. The first three rows show samples from North American cities. The last two rows show samples from South American cities, where the available aerial satellite images have lower resolutions.

To address the lack of large-scale datasets that specifically target the pedestrian environment, we develop a method to rasterize more commonly found GIS annotations and present the Annotations for Pedestrian Environment (APE) dataset, which contains labels for essential classes that compose a connected pedestrian path network graph. The GIS annotations for pedestrian path environments, either created by city agencies or crowdsourcing by local mappers (e.g. OpenStreetMap (OSM)), are primarily created for use in vector GIS data systems. The commonly used GIS data formats (e.g. Shapefile map or GeoJSON) cannot be directly used in computer vision tasks that required input data to be imagery-like. In creating the APE dataset, we rasterize GIS data from multiple sources in multiple geographic areas to create the labeled data needed for understanding pedestrian environments in computer vision tasks.

#### 3.1 Coverage

The APE dataset covers select regions in the following areas: (1) Los Angeles, CA, United States (2) Bellevue, WA, United States (3) Quito, Ecuador (4) Sao Paulo, Brazil (5) Santiago, Chile (6) Gran Valparaiso, Chile. Consequently, the APE dataset encompasses diverse built environments from both North and South America. Cities in South America are more densely populated with higher street and intersection density, but simpler in shape compared to North American cities[22]. The OpenSidewalks Project [4] provides crowdsourced GIS annotations for select regions in each of these 6 cities in a consistent schema [5], with annotations created by trained and veteran mappers.<sup>2</sup> These annotations are available in standard GIS formats in OSM. In addition, the City of Los Angeles provides a GIS dataset with additional information for the Los Angeles area, including annotations for sidewalks, crossings, corner bulbs, and driveways [20]. The APE dataset consists of a total of 14,800 samples and spans approximately 2,700km<sup>2</sup> land area.

#### 3.2 Collecting Images and Annotations

Each dataset sample includes (1) the aerial satellite image, (2) the street map imagery tile, and (3) the rasterized annotations for pedestrian paths. The aerial satellite images and the street map imagery tiles are acquired from Bing Maps along every major road for each area described in Section 3.1. The GIS annotations mainly come in three geometry types: *Point*, *LineString*, and *Polygon*. The nodes representing crossings, curbs, and link endpoints are usually annotated as the *Point* geometry. For each of these annotated objects with the *Point* geometry, we convert the point to a circle by adding a buffer zone with a fixed radius and then rasterize the circle in the image. Sidewalks and crossings usually have two types of representations: the centerline representation and the polygon boundary representation. (1) If they are represented by their centerline as the *LineString* geometry, for each *LineString*, we convert it into a *Polygon* by adding a buffer zone to each side of the *LineString*, then rasterize the *Polygon* as a filled polygon shape in the image. (2) If they are already annotated by their boundary as the *Polygon* geometry, we directly rasterize them as filled polygon shapes in the image.<sup>3</sup> Each set of an aerial satellite image, a street map image, and a rasterized annotation are aligned with their precise geographic bounding boxes.

#### 3.3 Classes

Our dataset provides annotations for three distinct classes needed to semantically segment the pedestrian environment. Shown in Figure 2, these classes include (1) **Corner bulb** (2) **Sidewalk** (3) **Crossing**. Corner bulbs are commonly used when describing a transportation network since they serve as a transition zone connecting a sidewalk to curb ramps, crossings, or another sidewalk. The nodes representing sidewalk endpoints, link endpoints, and

<sup>2</sup>A mapper is considered trained if the mapper is given clear written instructions and had completed at least one mapping training session. A trained mapper is considered to be a veteran mapper if the mapper has completed several mapping tasks and the mapping outcomes are validated in a peer-review process.

<sup>3</sup>For the data currently in the APE dataset, the polygon boundary representation is only applied to the sidewalk annotations from the City of Los Angeles. The sidewalk annotations from the OpenSidewalks project are all in the centerline representation per the OpenSidewalks standard.

curbs are usually located within the corner bulbs. Sidewalks and crossings are essential elements in an urban pedestrian path network graph, as the lines representing sidewalks and crossings are essentially the edges a pedestrian will traverse. The focus of our work is on mapping pedestrian paths, thus, all other annotated classes (including roads, buildings, and trees) collectively comprise *background class* in the following experiments. However, these classes can also be rasterized with our method (Section 3.2) and added to the APE dataset as additional class labels.

### 3.4 Challenges

As previously noted in autonomous-driving datasets, imagery dataset bias can be introduced by non-representative geographic locations of images [12, 25, 27]. This bias generally refers to a systematic error that results from the training dataset only representing a limited geographic region or a limited type of environment. To expand APE’s ability to generalize to other regions and environments, we created a more balanced and representative dataset that includes images from a diverse range of regions and environments, including imagery from both North and South American urban contexts. However, this introduces a different challenge having to do with geopolitical biases in satellite image collection: only lower-resolution aerial satellite images were openly available for South American cities. This presented challenges in model learning. Figure 2 displays samples from the APE dataset, with the first three rows representing North American cities and the last two rows representing lower-resolution images in South American cities. Differences in model performance when predicting pedestrian networks in different cities are discussed in Section 5.

Another challenge for learning from aerial satellite images is that the important classes such as sidewalks and crossings are occluded in some satellite images [19]. For example, in the second and third rows shown in Figure 2, part of the pixels that are labeled as *sidewalk* are occluded by buildings, the shade from buildings, and vegetation in the aerial satellite images. In these cases, learning from aerial satellite images alone is very challenging and street map imagery tiles can provide crucial auxiliary information. Section 5 shows a quantitative analysis of the models trained with (1) only aerial satellite images (2) only street map imagery tiles, and (3) both aerial satellite images and street map imagery tiles.

## 4 INFERRING A CONNECTED PEDESTRIAN PATH NETWORK GRAPH WITH APE

The APE dataset enables us to automate the process of generating pedestrian path network data at scale. In this section, we introduce an end-to-end process to infer a connected pedestrian path network. As shown in Figure 3, the process consists of three main steps. First, Section 4.1 details the creation of a hypothesized graph with existing street network data using a tool we developed called *Pedestrianfer*. Second, Section 4.2 describes the multi-input segmentation network trained on the APE dataset for generating pixel-wise prediction masks for the important classes in the pedestrian environment. Lastly, Section 4.3 describes the process of using the predictions from the segmentation network to optimize the hypothesized graph into an accurate connected pedestrian path network graph.

### 4.1 Pedestrianfer

To infer a hypothesized pedestrian path network from incomplete information, we developed *Pedestrianfer* (Figure 4). *Pedestrianfer* performs three discrete inference tasks: (1) inferring an (optimistic) sidewalk network from the existing street networks, (2) inferring street crossing locations from a street network and a sidewalk network, and (3) creating a preliminary conjecture of curb interface locations around the street crossings. *Pedestrianfer* is similar in outcome to the method [14] discussed in Section 2 but with the following key differences: (1) the only required input to our method is a street network, and sidewalk locations are estimated from that street network, (2) full intersection-to-intersection sidewalk paths are estimated rather than being broken into 50-meter segments, and (3) street crossing pathways are generated via a cost function that weighs several properties and does not require manual intervention. Additionally, *Pedestrianfer* generates pathways according to the OpenSidewalks Schema specification [5].

*Pedestrianfer* first infers sidewalk networks from a street network under two alternative regimes, each representing a different hypothesis on the built environment. The alternatives depend on whether the data source is a vector layer of street network data with metadata on the presence (and optionally, offset distance) of sidewalks. If metadata is present, sidewalks are placed only where indicated. If not, *Pedestrianfer* retrieves a street network and hypothesizes full sidewalks, i.e. the scenario in which all streets have a connected sidewalk on either side. In the latter case, *Pedestrianfer* retrieves an existing street network, such as a street network from OSM. In either case, the *Pedestrianfer* process involves 4 steps: (1) creating a directed graph representation of the street network, (2) generating all right-hand-turn paths that start and stop at the same node (closed paths), (3) drawing sidewalks via line offset algorithm, and (4) trimming or joining them based on the path context, e.g. trimming overlapping sidewalk lines when they are neighbors along the path.

Next, *Pedestrianfer* uses a street network and a sidewalk network (generated by *Pedestrianfer* in the previous step if it does not already exist) to infer crossings. In this step, *Pedestrianfer* iterates over each street intersection node in the street network and generates (1) all street lines associated with the intersection, directed outwards from the intersection and up to half of the distance to the next intersection, and (2) a set of candidate sidewalks to connect with a crossing on each side of each street. *Pedestrianfer* then generates candidate crossings that are lines drawn from a sidewalk on the left side of a street to a sidewalk on the right side of a street. Multiple candidates are generated for each street of an intersection by selecting a series of points along that street (every 1 meter) and generating a crossing that connects the closest corresponding left and right sidewalks. Metrics of the candidates are generated, including the distance of the street point to the intersection, the crossing line length, and the angle between the crossing line and the street it crosses. A cost function is then used to heuristically select the *best* crossing: one that minimizes a linear combination of the distance to the street intersection, crossing line length, and non-orthogonal crossing angles. Therefore, *Pedestrianfer* estimates street crossing locations that are near intersections, which align with common (albeit not universal) policies and pedestrian safety measures. In the case that



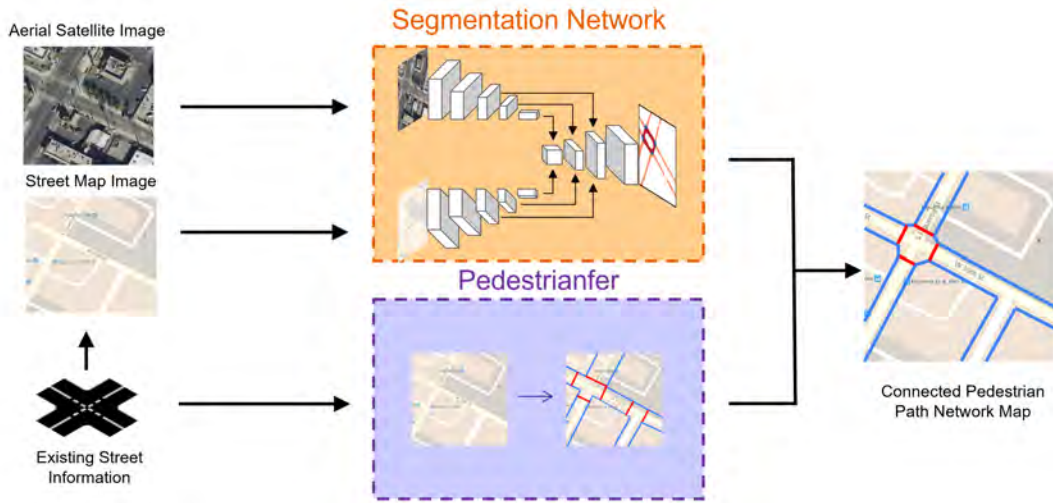


Figure 3: Overview of the end-to-end process for inferring pedestrian path network graph: Section 4.1 outlines *Pedestrianfer*, generating a pedestrian path network from incomplete information. Section 4.2 details the segmentation network using aerial and road map images to predict class labels of pixel locations in the pedestrian map. Section 4.3 explains how *Pedestrianfer* and segmentation network information are combined for accurate graph inference.



Figure 4: Hypothesized graph generated with *Pedestrianfer*. **Yellow** dots represent the hypothesized nodes for the curb nodes and the nodes on the sidewalk. **Blue** lines represent the hypothesized sidewalks. **Red** lines represent the hypothesized crossings. Links are part of the crossings between the curb nodes and the nodes on the sidewalk

crossing locations with ground markings are already known and present in an associated dataset (as *Point* data), *Pedestrianfer* will generate a crossing near that point by projecting to the nearest known sidewalk candidates on each side.

Lastly, *Pedestrianfer* splits crossing lines into three segments: (1) the originating sidewalk surface, (2) the street surface, and (3) the destination sidewalk surface. These correspond to typical surfaces a pedestrian would travel in an urban area. Sidewalk-street transitions, which often have vertical displacements or are where curb-cut ramps meet the street, can be interpreted as potential curb interfaces. *Pedestrianfer* does not use any metadata to determine where to split each crossing into these 3 segments, but rather

provides a low-information hypothesis on sidewalk-curb-street locations, which can be improved with manual mapping (GIS software, crowdsourcing in OSM) or learning-based method such as those described in Section 4.2 and Section 4.3.

## 4.2 Segmentation Network

In order to verify, correct, and refine the *Pedestrianfer* hypothesized path network graph, we use inference from a CNN-based segmentation network. The segmentation network has a siamese-like structure, allowing it to fuse and utilize the information from both the aerial satellite image and the street map tiles. As shown in Figure 3, the segmentation network has two identical branches, the aerial satellite images are used as the input in one branch and the street map tiles are used in the other. Each branch has its encoder-decoder structure. In addition, to fuse the information from both branches, we concatenate feature maps from both branches at different layers hierarchically during the up-sample process. Experiments described in Section 5 used FCN-8s [15] as the backbone model, but it can be replaced by any other segmentation model with a similar structure.

## 4.3 Optimization for Full Pedestrian Path Network graph

**4.3.1 Optimizing node geolocation.** *Pedestrianfer* generates a hypothesized path network graph that outlines the potential location and connectivity of sidewalks and crossings. To obtain a more accurate pedestrian path network, we use information from the segmentation network to refine the hypothesized path network graph. Our strategy is to first find the optimized geolocation of the nodes in the graph, then connect the nodes with edges to complete the graph.

At each intersection, there are nodes representing sidewalks endpoints and curbs in a hypothesized graph generated by *Pedestrianfer*.

Ideally in a correct pedestrian path network, these nodes should all locate in a connected region being segmented as the *corner bulb* class (an example is given in Figure 5 (c) and Figure 5 (d)). To infer the correct geolocation of these nodes, we aim to find a parameterized affine transformation to warp each set of hypothesized nodes in a corner, represented by their pixel coordinates, to a new set of coordinates, so they better align with the *corner bulb* class in the predicted segmentation mask. We start by connecting the nodes at each corner (shown in Figure 5a) to form a closed polygon (shown in Figure 5b), then we find a parameterized affine transformation so that the sum of the probability of each pixel under each closed polygon being the *corner bulb* class is the greatest.

Mathematically, this optimization process can be defined as follows. For the total of  $n$  points  $[(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)]$  in a corner in a given Image  $I$ , the affine transformation that warps them to a set of new points  $[(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_n, y'_n)]$  can be described as:

$$\begin{bmatrix} x'_1 & \dots & x'_n \\ y'_1 & \dots & y'_n \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} * \begin{bmatrix} x_1 & \dots & x_n \\ y_1 & \dots & y_n \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = AX + t \quad (1)$$

In words,  $X$  is a set of points representing the coordinates of the nodes in a hypothesized graph at a given street corner. The transformation parameters  $A$  and  $t$ , and the new set of coordinates  $[(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_n, y'_n)]$  are to be found using the information obtained from the segmentation network.

Assuming there are a total of  $m$  pixels that fall into the polygon enclosed by the  $n$  corner points  $X$ , and the probability (as predicted by the segmentation network) that each of these  $m$  image pixels to be the *corner bulb* class is denoted as  $p_i$ , then we define the function that finds these pixels and their probability as  $f$ :

$$f : X \mapsto [p_1, p_2, \dots, p_m] \quad (2)$$

and define a function  $g$  that sums  $[p_1, p_2, \dots, p_m]$  as:

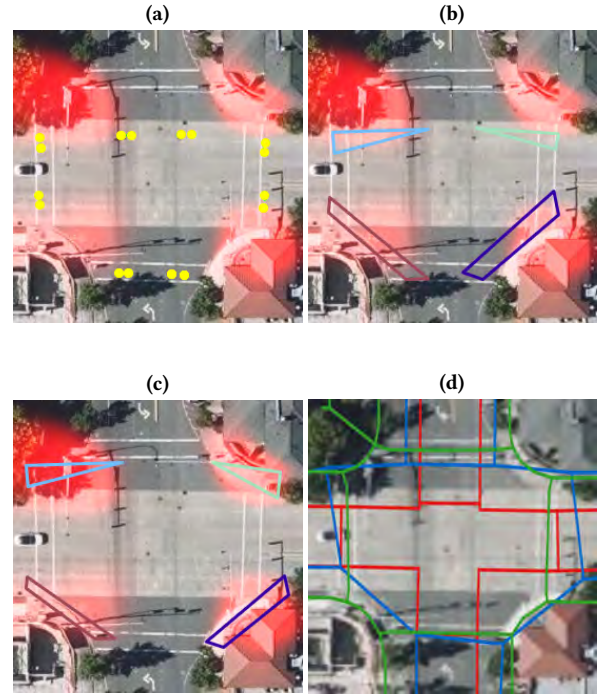
$$g(f(AX + t)) = g([p_1, p_2, \dots, p_m]) = \sum_{i=1}^m p_i \quad (3)$$

The goal is to maximize  $g$  so that the pixels under the new polygons have the greatest sum of probability (shown in Figure 5c). Thus, the optimization problem can be expressed as:

$$\begin{aligned} & \underset{A, t}{\text{minimize}} && -g(f(AX + t)) \\ & \text{subject to} && \forall i \in [1, n] \quad 0 < x'_i < I_{\text{width}}, 0 < y'_i < I_{\text{height}} \end{aligned} \quad (4)$$

There is no closed-form solution to Equation 4, and the gradient of  $f$  cannot be explicitly found. Hence we use the simultaneous perturbation stochastic approximation (SPSA) [23] method to find the optimal parameters  $A$  and  $t$  of the objective function  $g$ .

After optimizing sidewalk and curb node geolocation, we connect them with the information from the hypothesized graph to generate new sidewalks and crossing edges. The optimization process is shown in Figure 5. Figure 5a and 5b show the original node locations in the *Pedestrianfer* hypothesized graph. In this example, *Pedestrianfer* approximates the nodes erroneously in the middle of the road. Figure 5c shows the locations of the nodes post optimization, where the nodes and the enclosed polygon formed by the nodes are moved to the middle of a region predicted as the *corner*



**Figure 5: Illustration of node location optimization, the probability of each pixel being the *corner bulb* is shown as a heat map in red in (a) - (c). (a) *Pedestrianfer* hypothesized nodes (b) Polygons formed by the hypothesized nodes in each corner (c) New set of nodes optimized with information from the segmentation network (d) Compare to Human annotation: **Green** is human annotation graph **Red** is *Pedestrianfer* Hypothesized graph. **Blue** is the optimized graph.**

*bulb* class. The improved graph (post-optimization) in Figure 5d, shows graph elements considerably closer to the human-generated (ground truth) graph. A more detailed quantitative evaluation is made in Section 5.4.

**4.3.2 Graph refinement with class probability.** The optimized graph is used with probability masks of each class from the segmentation network to further improve the predicted path network map accuracy.

For a set of hypothesized nodes at a given corner, if the closed polygon they form does not overlap with enough high-probability pixels representing the *corner bulb* class, we consider these nodes to be falsely hypothesized. Mathematically, we define  $\mu_p$  as:

$$\mu_p = g(p)/m \quad (5)$$

If  $\mu_p$  is less than a set threshold, these nodes are considered to be false-hypothesized and therefore directly removed from the graph without optimizing with equation 4. The threshold is chosen to balance the edges' precision and recall in this post-processing step. Higher thresholds lead to higher precision.

For downstream applications (e.g. wayfinding) that use the predicted graph, confidence values are added to edges to inform the optimization for high-confidence routes, i.e., the application may choose to avoid low-confidence edges and instead optimize for higher-confidence paths. For each edge created by connecting two nodes in the graph, we assign a confidence value as an attribute of the edge as follows: each *LineString* representing the *crossing* class, is converted into a polygon by adding a buffer to each side of the *LineString*, then we compute the mean probability of the pixels in the polygon in the *crossing* class, similar to Equation 5. The mean probability is used as the confidence value and stored as an attribute of the edge in the graph. Similarly, the confidence values of the *sidewalk* edges are computed and stored as an attribute of the *sidewalk* edges. The confidence value improves the graph for various transportation network analyses, sidewalk network scoring, personalized routing [3], and other downstream applications.

## 5 EXPERIMENTS

### 5.1 Training the Segmentation Network

Three segmentation networks were trained with different setups: aerial satellite image branch only, street map image tile branch only, and both aerial and street map image tile branches. All three models used the same dataset split and data augmentation techniques. Performance comparisons are shown in Figure 6 and Table 2.

**5.1.1 Dataset.** We train and validate with an 80%/20% split of the APE imagery and annotation samples in the Los Angeles area. Specifically, we chose to use the rasterized data from the City of Los Angeles annotations because of the amount of data available. Bellevue and Quito annotations from OpenSidewalks are used as unseen test sets.

**5.1.2 Data Augmentation.** To improve the robustness of the segmentation network, data augmentations were applied to the training data during the training stage. First, random rotating, cropping, and resizing were applied. Again, we note the resolution differential in open aerial satellite images in North America (Los Angeles and Bellevue) as compared to those from South America (Examples are shown in Figure 2). We enhanced the segmentation network’s performance on low-resolution images by using Gaussian kernels with random sizes to artificially generate low-resolution images in training.

### 5.2 Qualitative Analysis

Figure 6 visualizes the segmentation results on the validation set. The models’ shape segmentation and identified locations for *sidewalk*, *crossing*, and *corner bulb* align well with the ground-truth segmentation. These qualitative examples show the difficulty in predicting pedestrian path network classes with single-source input and the improvement gained from adding other input sources. The first and second rows of Figure 6 show that different kinds of errors are produced when the model is trained with only one branch of input data. In the example shown in the first row, the prediction made with the model trained with only street images generated false-positive *sidewalk* predictions. Adding aerial satellite images in model training helped remove these spurious predictions. Predictions made with only aerial satellite images incurred other

limitations, as shown in the second-row example, when sidewalks were occluded in the satellite image by vegetation, the model trained with only those images generated false-negative *sidewalk* predictions. Adding street image tiles in training helped predict many of the occluded sidewalks. To further demonstrate the effectiveness of using both the aerial satellite images and the street image tiles as inputs to the model, the examples in row 3 and row 4 of Figure 6 show different cases in which using both branches of data generated better result than using one branch of input data alone. In the third row example, the model trained with only one branch of data generated different false-positive *sidewalk* predictions, and using both branches of data to train the model helped remove these false-positive predictions. Similarly, in the fourth-row example, models trained with only one branch of data generated false-positive predictions on *sidewalk* and *crossing*, while using both branches of data helped remove these false predictions. A quantitative analysis of the performance of our method when testing in different cities is provided in Section 5.4.

### 5.3 Quantitative Analysis

Table 2 shows the quantitative experiment results for the models trained with the 3 different setups: (1) trained with the aerial satellite image branch only, (2) trained with the street map image tile branch only, and (3) trained with both the aerial satellite images and street map image tiles. The metrics we included are (1) overall mIoU, (2) the mIoU for each of the 4 classes (*background* (BG), *sidewalk* (SW), *corner bulb*, and, *crossing*) in the dataset, (3) pixel accuracy, and (4) instance-wise precision and recall for the important *corner bulb* class.

From Table 2 we observe that the model trained with both the aerial satellite images and street map image tiles has a higher mIoU on each of the classes than the models trained with only one branch of data, demonstrating its better pixel-wise performance. In addition, higher precision and recall on the important *corner bulb* class demonstrates its better instance-wise performance. These metrics combined shows the effectiveness of the model for segmenting the pedestrian environment with both the aerial satellite images and the street map image tiles, confirming the discussions we made in Section 5.2 based on the observations from Figure 6. A graph-level analysis is made in Section 5.4.

### 5.4 Quantitative Evaluation of the Inferred Path Network Graph

Although the models perform well with the pixel-wise and instance-wise measures, these metrics do not fully reflect the accuracy of the connected graph prediction. mIoU (or other pixel-wise measures) cannot measure how close a predicted graph is to the ground truth graph. In order to measure the similarity between the predicted graphs to the ground truth graph, we compare (1) the pedestrian path network graph generated by *Pedestrianfer*, and (2) the graph optimized using the segmentation network, to the ground truth graph generated and validated by human mappers from the OpenSidewalks project. *Pedestrianfer* is used as the baseline method as it is similar to the method proposed by [14] and it represents the methods that derive a pedestrian path network purely from existing street network information. The metrics we used for evaluating the



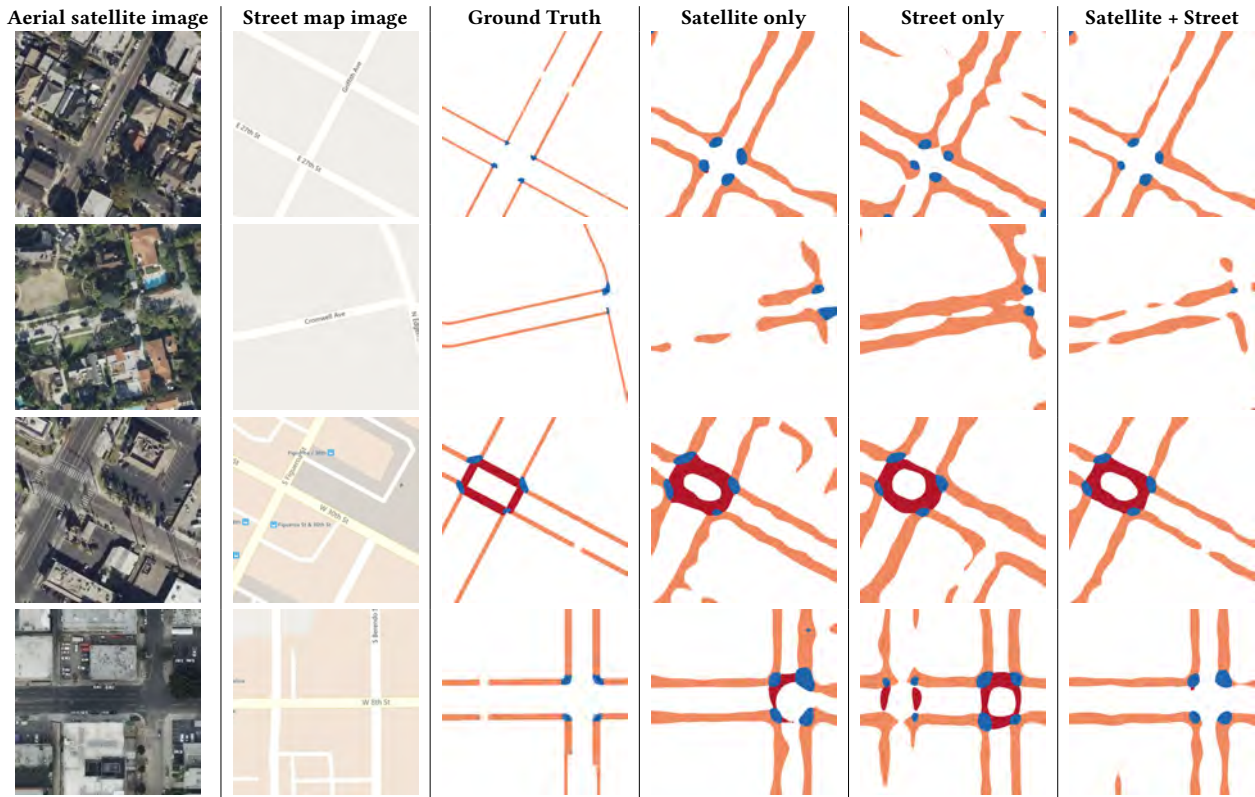


Figure 6: Qualitative results on the validation set. The segmentation results of 3 different models are shown in columns 4-6. (1) Trained with the aerial satellite image branch only (2) Trained with the street map tile branch only (3) Trained with both the aerial satellite image branch and the street map image tile branch. The model that uses both aerial satellite images and street map images generates better predictions than the models that use only one input. A detailed discussion of these samples is made in Section 5.2. Quantitative analysis of models’ performance is made in Section 5.3 and Table 2.

Table 2: Quantitative segmentation results: the model that uses both aerial satellite images and street map tiles outperforms models that use only one branch of data

Method	Pixel-wise						Instance-wise	
	BG IoU	SW IoU	Corner IoU	Crossing IoU	mIoU	Accuracy	Corner precision	Corner recall
Satellite Only	0.89	0.42	0.47	0.42	0.55	0.89	0.59	0.68
Street Only	0.85	0.33	0.48	0.37	0.51	0.86	0.60	0.63
Satellite + Street	0.91	0.47	0.57	0.49	<b>0.61</b>	0.92	0.73	0.69

graphs are similar to the ones proposed by [17] for road topology measurements, namely the precision, recall, and F1 score based on the assignments of predicted sidewalk edges (or crossing edges) to the corresponding edges in the ground truth graph. The results are summarized in Table 3. The first row of each sub-table shows the performance of the baseline model, and the second row of each sub-table shows the performance of our method where the information from the segmentation networks that use both aerial satellite images and street images is used to optimize the predicted graphs. Table 3a shows evaluation results on sidewalks and Table 3b shows evaluation results on crossings. The graph generated with the segmentation network outperforms the baseline model, especially in

the precision category because the baseline model that uses only street network information tends to be overly optimistic about the existence of pedestrian paths. High precision, recall, and F1 score demonstrate the effectiveness of our method in inferring a pedestrian path network with both the existing street network and the information learned with the segmentation network. In addition, our method maintains high precision and recall when tested on unseen datasets (Bellevue and Quito) and outperforms the baseline methods, despite the test areas having significantly different built environments and aerial satellite image resolution (with notably low-resolution imagery in Quito). Most importantly, since we start from a connected graph hypothesized by *Pedestrianfer* and utilize



**Table 3: Quantitative evaluation: graph level analysis**

(a) Sidewalks									
	Precision			Recall			F1		
	LA	Bellevue	Quito	LA	Bellevue	Quito	LA	Bellevue	Quito
<b>Pedestrianfer (Baseline)</b>	68.1	69.2	67.1	73.4	73.2	71.2	70.7	71.1	69.1
<b>Pedestrianfer + Segmentation</b>	83.5	82.4	75.6	77.8	77.0	73.6	80.5	79.6	74.6

(b) Crossings									
	Precision			Recall			F1		
	LA	Bellevue	Quito	LA	Bellevue	Quito	LA	Bellevue	Quito
<b>Pedestrianfer (Baseline)</b>	73.5	69.2	71.9	82.1	71.8	75.8	77.6	70.5	73.8
<b>Pedestrianfer + Segmentation</b>	76.7	72.2	74.4	87.7	75.1	78.7	81.8	73.6	76.5

the segmentation network alone to correct node locations and remove false-hypothesized geometries, the outcome graph stays as a connected graph that can be directly used as a routable network graph.

## 6 CONCLUSION

In this work, we introduce the APE dataset to address the dearth of data and methods to automate mapping for pedestrian path networks. APE includes aerial images, street map tiles, and annotated pedestrian environments. It provides the research community with a new and challenging dataset to address pedestrian network path predictions through machine learning methods. The dataset covers diverse urban areas and we demonstrated it can be used for computer vision tasks. Additionally, we developed a method to infer pedestrian path networks using segmentation and street data, and validated the method’s accuracy through comparison to human-annotated data. Through our experiments, we were able to both metricize human performance in pedestrian networks annotation tasks like drawing geometries and labeling (metrics achieved through double annotations) and also evaluated the baseline performance of a commonly used segmentation model. The APE dataset and our process provide valuable initial contributions to pedestrian environment research, specifically in contributing to the many wayfinding and planning tasks that rely on detailed pedestrian transportation networks. We hope this work will inspire further pedestrian network data collections at scale.

## ACKNOWLEDGMENTS

This work was funded in part by the Taskar Center for Accessible Technology, USDOT ITS4US NOFO No: 693JJ322NF00001, and Microsoft’s AI4Accessibility award. Thanks to the G3ict organization for its support of the Los Angeles, Quito, Sao Paulo, Santiago, and Gran Valparaiso, on-the-ground mapping efforts.

## REFERENCES

- [1] Dragan Ahmetovic, Roberto Manduchi, James M Coughlan, and Sergio Mascetti. 2015. Zebra crossing spotter: Automatic population of spatial databases for increased safety of blind travelers. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. 251–258.
- [2] Selim Aksoy, Bahadır Ozdemir, Sandra Eckert, Francois Kayitakire, Martino Pesarasi, Orsan Aytakin, Christoph C Borel, Jan Cech, Emmanuel Christophe, Sebnem Duzgun, et al. 2008. Performance evaluation of building detection and

- digital surface model extraction algorithms: Outcomes of the PRRS 2008 algorithm performance contest. In *2008 IAPR Workshop on Pattern Recognition in Remote Sensing (PRRS 2008)*. IEEE, 1–12.
- [3] Nicholas Bolten and Anat Caspi. 2019. AccessMap Website Demonstration: Individualized, Accessible Pedestrian Trip Planning at Scale. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. 676–678.
- [4] Nicholas Bolten and Anat Caspi. 2021. OpenSidewalks. <https://www.opensidewalks.com/>
- [5] Nicholas Bolten and Anat Caspi. 2021. The OpenSidewalks Schema. <https://github.com/opensidewalks/OpenSidewalks-Schema>
- [6] Nicholas Bolten and Anat Caspi. 2021. Towards routine, city-scale accessibility metrics: Graph theoretic interpretations of pedestrian access using personalized pedestrian network analysis. *PLoS one* 16, 3 (2021), e0248399.
- [7] Nicholas Bolten and Anat Caspi. 2022. Towards operationalizing the communal production and management of public (open) data: a pedestrian network case study: A pedestrian network case study in operationalizing communal open data. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS)*. 232–247.
- [8] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raskar. 2018. DeepGlobe 2018: A Challenge to Parse the Earth Through Satellite Images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [9] Marcelo Cabral Ghilardi, Julio CS Jacques, and Isabel Manssour. 2016. Crosswalk localization from low resolution satellite images to assist visually impaired people. *IEEE computer graphics and applications* 38, 1 (2016), 30–46.
- [10] Qing Hou and Chengbo Ai. 2020. A network-level sidewalk inventory method using mobile LiDAR and deep learning. *Transportation research part C: emerging technologies* 119 (2020), 102772.
- [11] ISPRS. 2022. ISPRS Test Project on Urban Classification, 3D Building Reconstruction and Semantic Labeling. <https://www.isprs.org/education/benchmarks/UrbanSemLab/default.aspx>
- [12] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 306–316.
- [13] Hassan A Karimi and Piyawan Kasemsupakorn. 2013. Pedestrian network map generation approaches and recommendation. *International Journal of Geographical Information Science* 27, 5 (2013), 947–962.
- [14] Hanyan Li, Jack Cebe, Sara Khoeni, Yanzhi Xu, Chelsea Dyess, and Randall Guensler. 2018. A semi-automated method to generate GIS-based sidewalk networks for asset management and pedestrian accessibility assessment. *Transportation research record* 2672, 44 (2018), 1–9.
- [15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.
- [16] Xiaoyan Lu, Yanfei Zhong, Zhuo Zheng, and Liangpei Zhang. 2021. GAMSNet: Globally aware road detection network with multi-scale residual learning. *ISPRS Journal of Photogrammetry and Remote Sensing* 175 (2021), 340–352.
- [17] Gellért Mátyus, Wenjie Luo, and Raquel Urtasun. 2017. Deeproadmapper: Extracting road topology from aerial images. In *Proceedings of the IEEE international conference on computer vision*. 3438–3446.
- [18] Lu Mi, Hang Zhao, Charlie Nash, Xiaohan Jin, Jiyang Gao, Chen Sun, Cordelia Schmid, Nir Shavit, Yuning Chai, and Dragomir Anguelov. 2021. HDMapGen: A hierarchical graph generative model of high definition maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4227–4236.

- [19] Huan Ning, Xinyue Ye, Zhihui Chen, Tao Liu, and Tianzhi Cao. 2022. Sidewalk extraction using aerial and street view images. *Environment and Planning B: Urban Analytics and City Science* 49, 1 (2022), 7–22.
- [20] City of Los Angeles. 2022. Los Angeles GeoHub Sidewalks (Mapped Areas). <https://geohub.lacity.org/maps/lahub::sidewalks-mapped-areas/about>
- [21] Deng Pan, Meng Zhang, and Bo Zhang. 2021. A generic FCN-based approach for the road-network extraction from VHR remote sensing images—using openstreetmap as benchmarks. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 14 (2021), 2662–2673.
- [22] Olga L Sarmiento, Andrés F Useche, Daniel A Rodriguez, Iryna Dronova, Oscar Guaje, Felipe Montes, Ivana Stankov, Maria Alejandra Wilches, Usama Bilal, Xize Wang, et al. 2021. Built environment profiles for Latin American urban settings: The SALURBAL study. *PLoS One* 16, 10 (2021), e0257528.
- [23] James C Spall. 1998. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on aerospace and electronic systems* 34, 3 (1998), 817–823.
- [24] Tao Sun, Zonglin Di, Pengyu Che, Chun Liu, and Yin Wang. 2019. Leveraging crowdsourced GPS data for road extraction from aerial imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7509–7518.
- [25] Antonio Torralba and Alexei A Efros. 2011. Unbiased look at dataset bias. In *CVPR 2011*. IEEE, 1521–1528.
- [26] Shenlong Wang, Min Bai, Gellert Mattyus, Hang Chu, Wenjie Luo, Bin Yang, Justin Liang, Joel Cheverie, Sanja Fidler, and Raquel Urtasun. 2016. Torontocity: Seeing the world with a million eyes. *arXiv preprint arXiv:1612.00423* (2016).
- [27] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. 2019. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097* (2019).
- [28] Songbing Wu, Chun Du, Hao Chen, Yingxiao Xu, Ning Guo, and Ning Jing. 2019. Road extraction from very high resolution images using weakly labeled OpenStreetMap centerline. *ISPRS International Journal of Geo-Information* 8, 11 (2019), 478.
- [29] Yuxiang Zhang, Sachin Mehta, and Anat Caspi. 2021. Collecting Sidewalk Network Data at Scale for Accessible Pedestrian Travel. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–4.
- [30] Kai Zhou, Yan Xie, Zhan Gao, Fang Miao, and Lei Zhang. 2021. FuNet: A novel road extraction network with fusion of location data and remote sensing imagery. *ISPRS International Journal of Geo-Information* 10, 1 (2021), 39.