

Scaling pedestrian crossing analysis to 100 U.S. cities via AI-based segmentation of satellite imagery

Marcel E. Moran, San José State University, marcel.moran@sjsu.edu
Arunav Gupta, New York University, arunav.gupta@nyu.edu
Jiali Qian, New York University, tony.qian@nyu.edu
Debra Laefer, New York University, debra.laefer@nyu.edu

Abstract: Accurately measuring street dimensions is essential to evaluating how their design influences both travel behavior and safety. However, gathering street-level information at city-scale with precision is difficult given the quantity and complexity of urban intersections. To address this challenge in the context of pedestrian crossings — a crucial component of walkability — we introduce a scalable and accurate method for automatically measuring crossing distance at both marked and unmarked crosswalks, applied to America’s 100 largest cities. First, OpenStreetMap coordinates were used to retrieve satellite imagery of intersections throughout each city — totaling roughly three million images. Next, Meta’s Segment Anything Model was trained on a manually labelled subset of these images to differentiate drivable from non-drivable surfaces (i.e., roads vs. sidewalks). Third, all available crossing edges from OpenStreetMap were extracted. Finally, crossing edges were overlaid on the segmented intersection images, and a grow-cut algorithm was applied to connect each edge to its adjacent non-drivable surface (e.g., sidewalk, private property, etc.), thus enabling the calculation of crossing distance. This achieved 93% accuracy in measuring crossing distance, with a median absolute error of 2 feet 3 inches (0.69 meters), when compared to manually verified data for an entire city. Across the 100 largest U.S. cities, median crossing distances ranged from 32 feet to 78 feet (9.8 – 23.8m), with detectable regional patterns. Median crossing distance also displayed a positive relationship with the cities’ year of incorporation, illustrating in a novel way how American city planning increasingly emphasizes wider (and more car-centric) streets. These findings identified opportunities to improve pedestrian safety and increase walkability at multiple scales, from the individual block to the entire city.

Keywords: satellite imagery, transportation, computer vision, pedestrians

1 Introduction

Making informed planning, design, and construction decisions regarding streets requires accurate information that is both granular at the block level and comprehensive to the municipal scale. Presently, datasets on intersection layouts are often missing, outdated, incomplete, and/or not machine readable. A fundamental component of intersections are pedestrian crossing distances, which influence both pedestrian safety and walkability (Abley et al., 2011; Martin, 2006). Thus, being able to generate accurate, citywide datasets of crossing distances can empower urban planners and traffic engineers to visualize patterns related to this feature of the built environment and strategically target streetscape investments. Optimizing the location of pedestrian refuge islands and sidewalk extensions are particularly important for ensuring that urban streets become

more walkable, especially for persons with mobility impairments, parents with strollers, and older adults, who may be most at risk within and deterred by long crossings (Wilmot & Purcell, 2022). That the U.S. is in the midst of a significant increase in pedestrian fatalities (following years of declines), makes such work particularly relevant (Cova, 2024; Schneider, 2020).

Scholarship in the last two decades has established an association between longer crossing distances and increased crashes. This flows intuitively from the fact that longer crossings increase the time pedestrians are exposed to oncoming automobiles. For example, an analysis of 10 years of pedestrian crashes at roughly 1,600 intersections in Utah found that for every 12 additional feet (3.7m) of crossing distance, pedestrian crashes increased by 5% (Islam et al., 2022). In Oregon, review of crash data between 2007-2014 at 191 crossings revealed that as the number of traffic lanes of the street increased, so did the risk-ratio for pedestrian crashes (Monsere et al., 2016). Crossing distances also impact more than just safety; in evaluating what aspects of streets in Mumbai, India were relevant to pedestrian flow, Kadali and Vedagiri (2015) determined that increasing crossing distance decreased intersections' perceived walkability. In the field of transportation planning, leading street-design guides (such as from NACTO, the Global Designing Cities Initiative, and the New York City Department of Transportation) advocate for crossing distance to be as short as possible, to make walking both safer and more likely (Bailey & Vega-Barachowitz, 2012; Duncan et al., 2016; Feuer et al., 2020).

For such reasons, urban scholars increasingly pursue means to identify and measure pedestrian crossings at larger scales. Recently, Verma and Ukkusuri (2023) mapped urban crosswalks with high accuracy in four American cities by training a deep-learning model to identify crosswalk markings from satellite imagery. That study extended an earlier computer-vision based attempt at identifying crosswalks by Ahmetovic et al. (2017) and complements the work of Li et al. (2023), who similarly trained a computer-vision model to recognize crosswalks from street-view imagery. Critically, those studies primarily rely on crossing markings (e.g., white or yellow "zebra" markings) as a primary means for identification, which raises the importance of methodologies that can also capture unmarked crossings, which can comprise a large percentage of total crossings in American cities, such as at 42% of intersections in San Francisco (Moran, 2022). In this vein, Moran and Laefer (2024) used a combination of data programmatically extracted from OpenStreetMap, along with manually labeled satellite imagery to measure crossings (both marked and unmarked) across three cities and found both intra- and inter-urban patterns that likely discourage walking in spatially distinct ways. Though highly accurate, that approach was labor intensive and time consuming, thus hampering its scalability.

More broadly, employing satellite imagery to study urban forms has grown significantly, as image resolution and availability have increased. In the realm of surface transportation, such efforts have included identification of roads, on-street parking, and traffic patterns (Haverkamp, 2002; Larsen et al., 2009; Moran, 2020). Computer-vision models' increasing ability to segment images has led to their increased use by urban researchers (Albert et al., 2017; Bagwari et al., 2023), including Meta's Segment Anything Model (SAM) (Osco et al., 2023; Ren et al., 2024). Building on these converging strands of research, the goal of this study is to leverage both street-network information embedded within OpenStreetMap, in combination with Meta's SAM, to automate accurate measurement of marked and unmarked pedestrian crossing distances at city-scale, thereby unlocking new insights on pedestrian networks nationally.

2 Methodology

The scope of this study is to automatically measure pedestrian crossing distances throughout the 100 largest U.S. cities by population (see Supplementary Materials S1 for list of cities). This requires maintaining accuracy across municipalities that vary in size, urban form (e.g., dense and walkable vs. sprawling and car-centric), topography, building height (and resulting shade), and tree cover. In addition, to ensure this method's reproducibility and potential for its adaptation to other urban questions, only freely available imagery and models are employed, and reasonable processing time (i.e., roughly one hour per city) was prioritized. Guided by these goals, the analysis is comprised of three primary steps (see Figure 1):

1. **Intersection Image Segmentation:** Train Meta's Segment Anything Model (SAM) to differentiate between drivable and non-drivable surfaces on satellite-imagery tiles, centered on street intersections.
2. **Crossing Edge Extraction:** Extract all available pedestrian crossing edges from OpenStreetMap, and generate additional edges by expanding crossing nodes into edges.
3. **Grow-Cut Algorithm:** Overlay crosswalk edges extracted from OpenStreetMap on segmented satellite imagery and deploy a grow-cut algorithm on each edge, resulting in accurately measured pedestrian crossings.

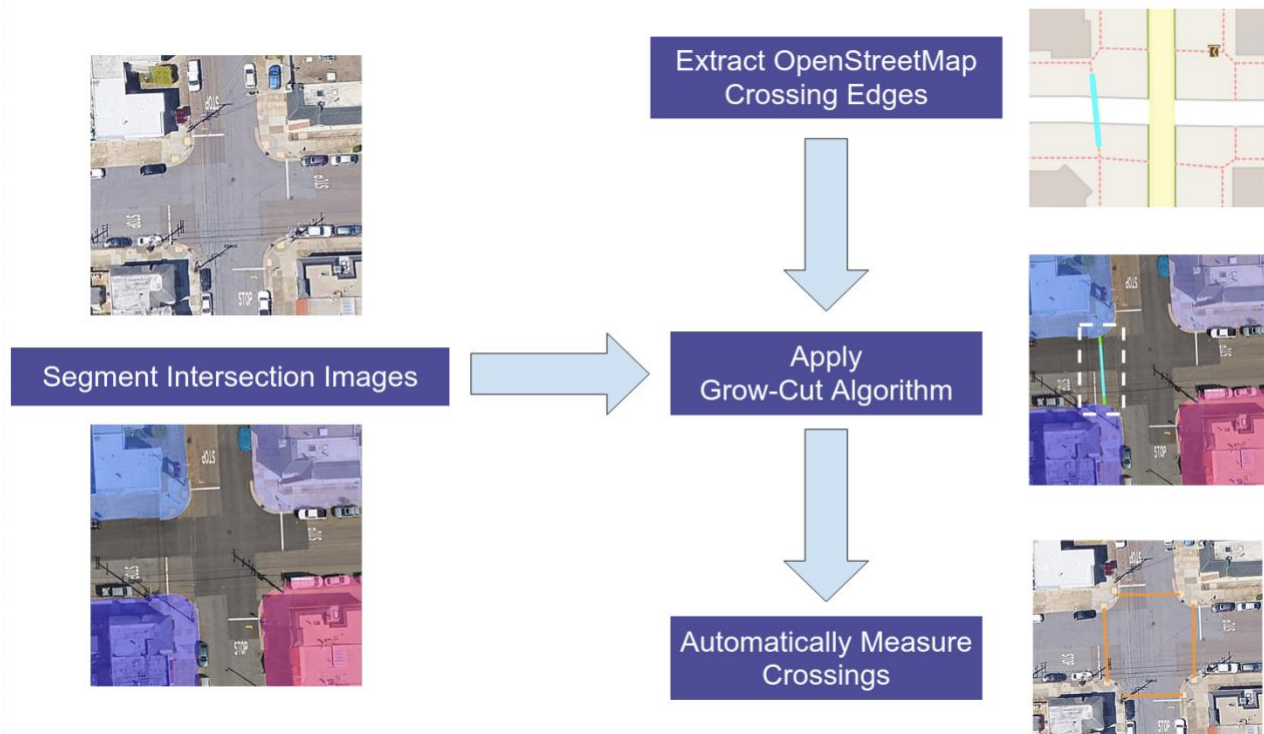


Figure 1. Diagram of analysis steps, including (1) segmentation of intersection satellite imagery, (2) extraction of crossing edges from OpenStreetMap, and (3) application of the grow-cut algorithm to the composite data set

2.1 Intersection image segmentation

First, the Python package OSMnx (Boeing, 2017) was used to generate the latitude-longitude coordinates of all street intersections and mid-block crossings for the 100 most populous cities in the United States. Those coordinates were then used to retrieve satellite-imagery tiles from Google Maps API, centered on each intersection. Each resulting image encompasses a 25-meter radius from its center point and is 1629 x 1629 pixels, with each pixel representing a 30x30cm patch on the ground. These images capture a wide range of intersection sizes and layouts and include contextual information from surrounding land uses (e.g., street trees, buildings, sidewalks, etc.). In total, approximately three million satellite imagery tiles were queried, totaling 8.3 terabytes of data. Tile retrieval was implemented with a parallelized approach that utilized multiple “server-less” workers simultaneously fetching tiles from different parts of a city, enabling a peak retrieval rate of 2,000 tiles per minute. Chicago required the largest number of tiles to be fetched, roughly 80,000, which took approximately 40 minutes.

After compiling these intersection tiles, the core computer-vision task was segmenting non-drivable surfaces (e.g., sidewalks, buildings, parks, and pedestrian-refuge islands). This was critical because accurately measuring crossing distance entails determining a crossing’s extent from one non-drivable segment to another (i.e., from sidewalk to sidewalk or from sidewalk to refuge island). As a “zero-shot” model, Meta’s SAM 2.1 (hereafter, SAM; Ravi et al., 2024), can segment objects without training data, but its performance can be improved via fine-tuning with manually labeled samples (Ma et al. 2024). To begin, satellite images of street intersections from 14 cities – covering a wide range of densities and urban form (Arlington, TX, Baltimore, MD, Boston, MA, Denver, CO, Durham, NC, Glendale, AZ, Huntsville, AL, Irvine, CA, Jacksonville, FL, Oklahoma City, OK, Omaha, NE, Reno, NV, San Francisco, CA, and Stockton, CA) – were manually annotated to include masks for all of the non-drivable surfaces, while leaving drivable surfaces (e.g., roads) unlabeled (see Figure 2).

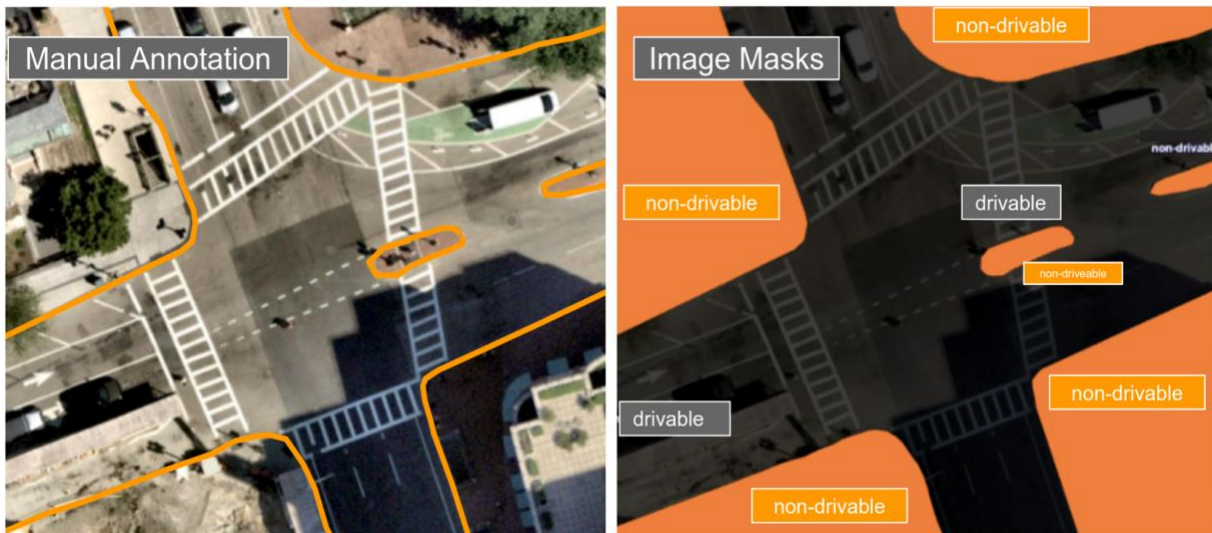


Figure 2. Manually annotated intersection image (left), with masks for non-drivable surfaces and the same intersection viewed as mask polygons (right)

Instead of manually labeling images at random, high-entropy intersection images (ranked internally for each city) were chosen. These high-entropy images contained greater visual complexity (see Figure 3) in contrast to low-entropy images, which typically depicted more homogeneous locations such as parking lots and industrial areas. Entropy was calculated using the standard log-probability sum (Shannon, 1948), where the probability distribution was defined using grayscale image pixel values rescaled to a 0-1 range. Tests using all three-color channels (red, green, blue) to define the pixel-value distribution resulted in no changes to the ranking of images by entropy for each city, so gray-scaling was used for computational efficiency.

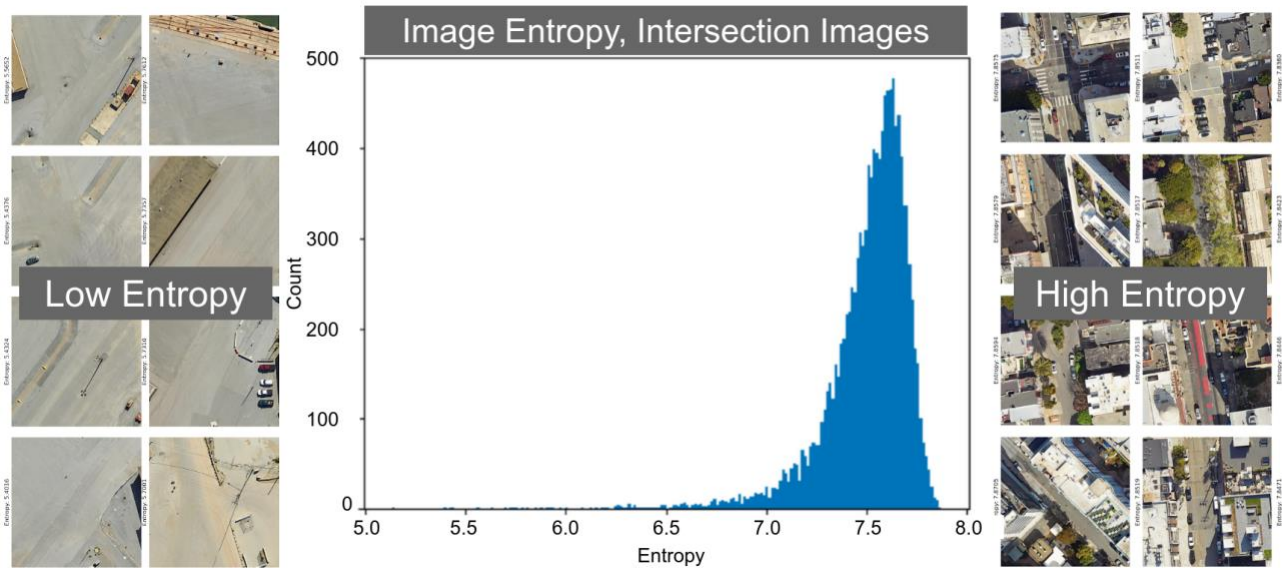


Figure 3. Entropy distribution for satellite imagery of street intersections in San Francisco (center), with examples of low-entropy images (left), and high-entropy images (right)

In total, 193 total high-entropy intersection images were selected for manual annotation. To increase the size of this training set, all manually annotated images were horizontally and vertically flipped, rotated (clockwise, counter-clockwise, and upside down), Gaussian blurred (up to 1.5px), and modified in terms of random noise (up to 1.05% of pixels). These modifications artificially expanded the training set to a final collection of 5,790 images.

Following this fine-tuning, SAM produced pixel-level masks of non-drivable surfaces within each of the three million intersection images drawn from all 100 cities (Dwyer et al., 2024). To enable this, each intersection image was first resized to 1024x1024 pixels to meet the model’s input requirements. Next, each image was converted into a 16x16 point grid, wherein each point serves as an “anchor” for SAM to view the image. SAM starts at each of these anchor points and gradually expands outward to form potential boundaries for the objects for which it is trained to segment. During fine-tuning, a denser point grid (24x24) in SAM was tested but produced more erroneous segmentation of objects in the roadway (i.e., false positives), for objects such as parked cars. As the point-grid density heavily influences the computational load of image segmentation, a 16x16 point grid was found to produce the best balance between accuracy and processing speed.

As part of the segmentation process, masks of non-drivable spaces were post-processed to extract contours that were then converted to polygonal geometries (see

Figure 4). To maintain spatial accuracy, each segmented image's file name was encoded with its center-point coordinates, which was later used to compute the corresponding bounding box. This allowed the resulting polygons to be mapped back to their real-world geographic locations, thereby ensuring alignment with original satellite imagery.



Figure 4. Segmentation of intersection images for non-drivable surfaces by SAM

Using the geographic bounding boxes derived from filenames, the mask contours were mapped to latitude and longitude coordinates and transformed into polygons. Next, for each image, the inverse selection of these polygons was used to identify all *drivable* areas. These inverted polygons were then exported as ESRI Shapefiles, thereby making the outputs compatible with a wide range of GIS tools. The decision was made to segment non-drivable areas as opposed to *drivable* areas, because refuge islands could not be carved out of drivable polygons during the manual-annotation phase.

To increase image-segmentation accuracy in differentiating between drivable and non-drivable surfaces, two additional post-processing steps were applied. The first aimed to reduce drivable areas being incorrectly segmented as non-drivable. This was achieved by applying a buffer to all street centerlines encoded within OpenStreetMap. Specifically, any mask polygon with at least 10m² or 5% of its area overlapping with the centerline-buffer was removed, successfully eliminating many false positives (see Figure 5). Buffers from street centerlines were dynamically set given differences in width; roads tagged as “primary” in OpenStreetMap had centerline buffers set at 2 meters, and untagged roads (primarily neighborhood streets) had centerline buffers set at 1.2 meters.

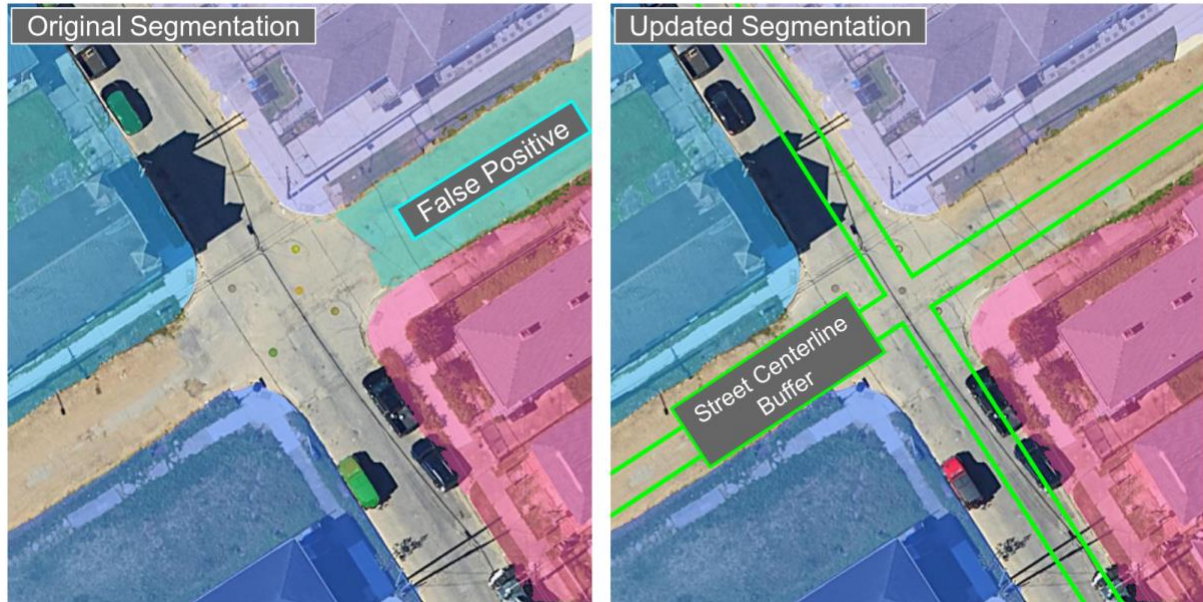


Figure 5. Intersection image with a false positive, non-drivable surface polygon (in turquoise, left). Segmentation improvement following the addition of a buffer (green lines, right) from the street centerline.

The second post-processing step aimed to include improperly omitted non-drivable areas (false negatives). To accomplish this, for each intersection image, all building footprints were extracted from OpenStreetMap as polygons and merged, creating a new non-drivable polygon that served as a substitute non-drivable mask, if one was not generated during image segmentation (see Figure 6).




Figure 6. Example of an intersection image with an omitted non-drivable surface (false negative, left). The addition of merged building footprints drawn from OpenStreetMap (center). Newly-segmented mask polygon based (right).

2.2 Crossing edge extraction

OpenStreetMap contributors generally encode pedestrian crossings as edges (connected nodes). Given this convention, the Python package OSMnx was used to extract all pedestrian crossing edges for each city. However, some pedestrian crossings in OpenStreetMap are only encoded as single nodes, in the center of the street between sidewalks. In those cases, additional crossing edges were generated by matching each of those nodes to the closest edge in OpenStreetMap with the tag “highway.” That allowed for construction of additional crossing edges; in some cities these newly generated crossing edges comprised over 50% of the total (see Supplementary Materials S2). To avoid the issue of duplicating crossings where both a crossing edge *and* a crossing node were present, all edges with endpoints that were within 5m of each other’s spatial location and within 5 degrees in terms of compass heading were merged into single edges. Though crossing edges in OpenStreetMap are not necessarily added by contributors in a way to precisely measure crossing distance (some do not reach sidewalks, whereas others end deep within sidewalks), they are of significant value as part of the grow-cut algorithm, described below.

2.3 Grow-cut algorithm

The final step of the analysis overlays the crossing edges extracted from OpenStreetMap atop the segmented intersection images. These edges provided a basis for applying a grow-cut algorithm, a common tool in image segmentation (Vezhnevets & Konouchine, 2005). For each crossing edge, the grow-cut algorithm expanded the line by 5% from both endpoints, while maintaining its original bearing/orientation (see Figure 7). Next, the location at which the newly-grown edge and the non-drivable polygon intersected was identified. Finally, the algorithm cut each edge where it intersected with the non-drivable polygons, resulting in a final layer of edges that represent pedestrian crossings.



Figure 7. Grow-cut algorithm sequence, in which the crosswalk edge from OpenStreetMap (left) is first grown by a magnification factor (center), and then cut based on the boundary of the non-drivable space segmentation (right)

Two additional post-processing steps were applied after completion of the grow-cut algorithm. The first removed any remaining duplicate crossings by constructing a buffer around all edges by 6.5ft (2m). Then all buffered edges that intersected and were within 10 degrees compass heading of each other were merged. Second, crossing edges smaller than 6 ft (1.83m) and larger than 160 ft (61m) were removed. These bounds were set

because visual inspection indicated there were effectively no pedestrian crossings shorter than six feet (the shortest being a crossing between a sidewalk and refuge island that is divided by a bike lane), and a very high proportion of crossings above 160 feet were determined to be erroneous as well. Figure 8 represents an expanded flow chart of the analysis, including data sources, data preparation, processing pipeline, and output products.

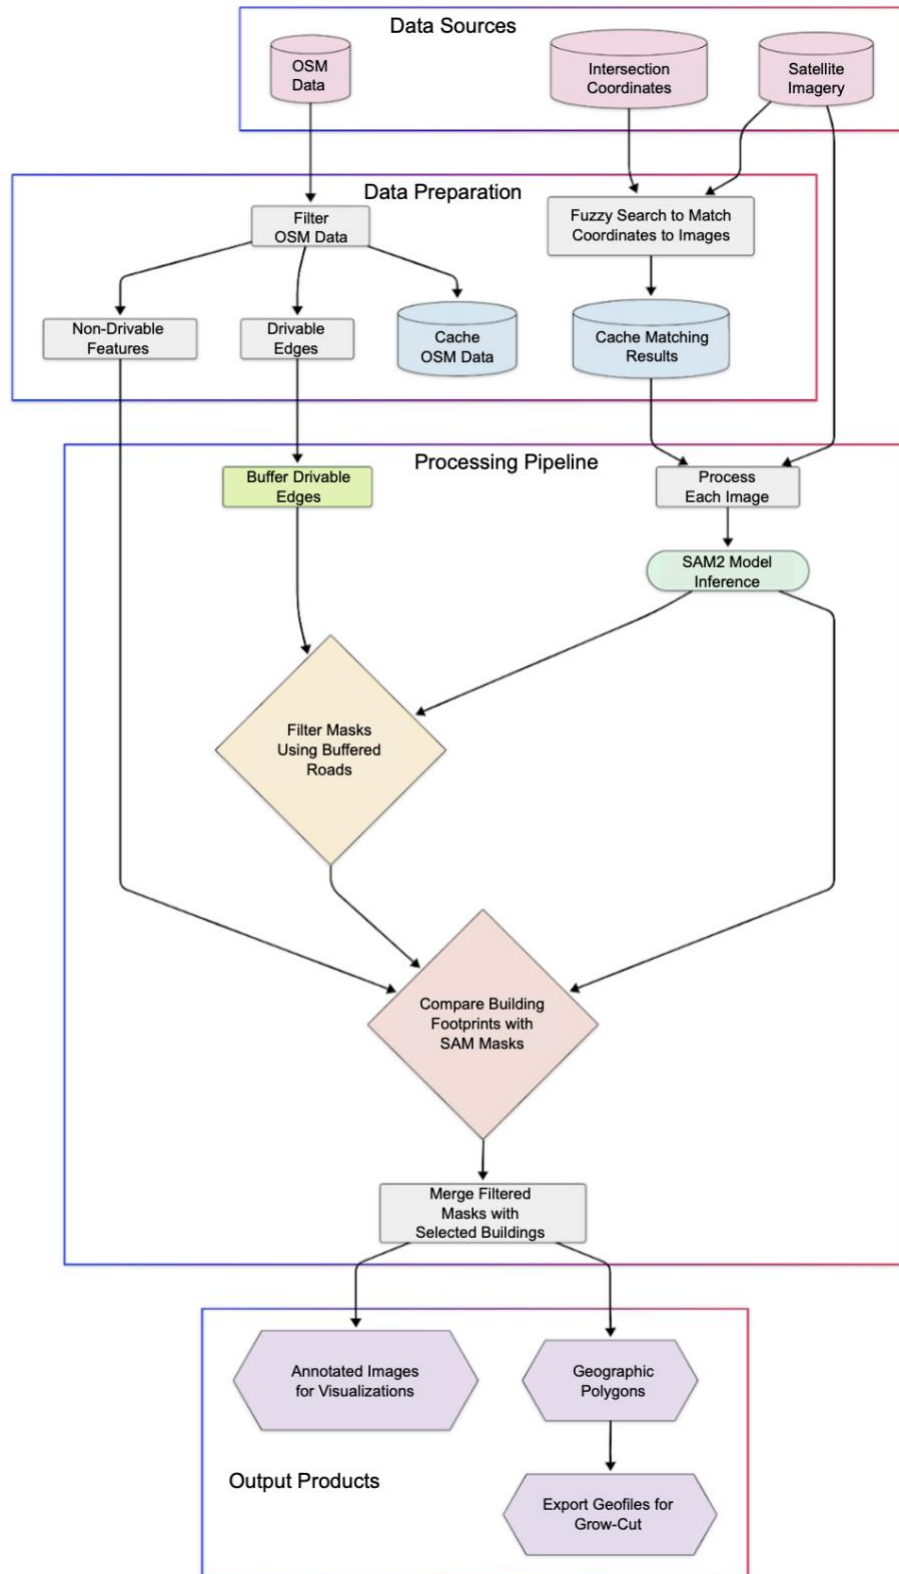


Figure 8. Flow chart of technical analysis steps, including original data sources (top), data preparation, processing pipeline (middle), and output products (bottom)

3 Results

The combination of intersection-image segmentation, crossing edge extraction from OpenStreetMap, and grow-cut algorithm application achieved automatic measurement of roughly 800,000 pedestrian crossings throughout the 100 largest cities in the U.S. This approach computed crossing distances for each city in roughly one hour, with the bulk of the processing time being dedicated to image fetching and segmentation by SAM. Visual inspection of automated crossings demonstrates successful placement throughout a wide range of urban densities, land-uses, and intersection layouts (see Figure 9).



Figure 9. Automatically generated crossing distances at intersections of varying complexity in San Diego (left), Philadelphia (center), and Baltimore (right), including both marked and unmarked crossings

3.1 Crossing distance accuracy

When compared to a citywide, manually verified crossing-distance dataset of San Francisco comprised of 12,948 crossings (drawn from Moran & Laefer, 2024), automated crossing distances achieved an accuracy of 93%, with a median absolute error of 2 feet, 3 inches (0.69m). Comparing the distributions of crossing distance for both the manually verified and automated datasets of San Francisco illustrates that accuracy remains high throughout the range (see Figure 10). In terms of coverage – the extent to which the automated process generated the *number and spatial extent* of crossings manually verified – the San Francisco comparison yielded 72% coverage, given availability of crossing edges in OpenStreetMap to draw on to seed the grow-cut algorithm.

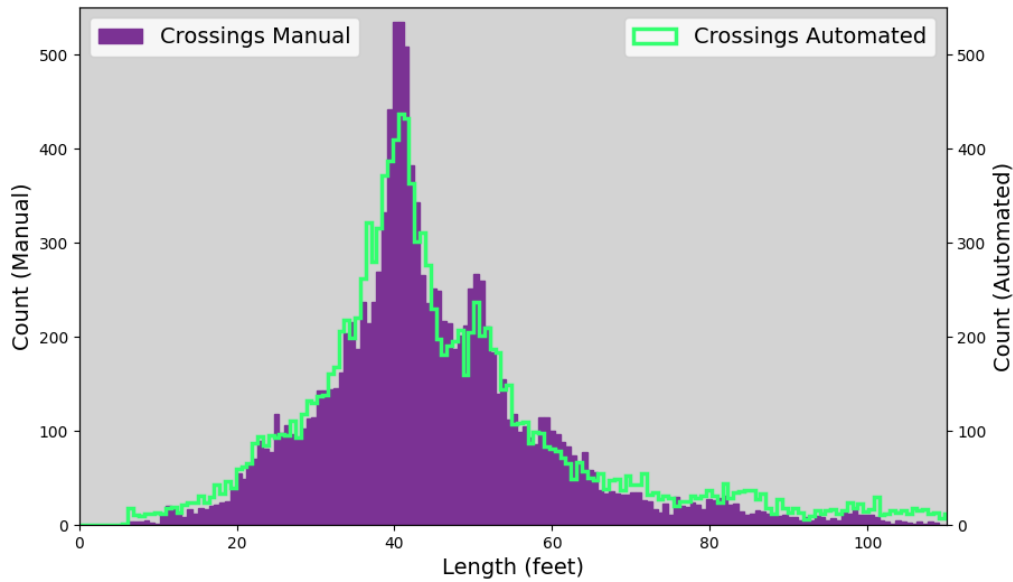


Figure 10. Distribution of pedestrian crossing distances in San Francisco, both via manual-verification (purple) and automated generation (green)

3.2 Nationwide analysis of crossing distance

This approach enables the first nationwide assessment of pedestrian crossing distance, with a number of emergent patterns. First, viewed at the scale of an entire city, crossing distance exhibits two spatial signatures: clusters, which represent neighborhoods where crossing distance is similar throughout, and corridor effects, relating to specific arterials where crossing distance is nearly uniform along a linear stretch of road (see Figure 11). Indeed, in the maps of both Chicago and New York, clusters of shorter crossings are evident throughout, often separated from one another by wider arterial corridors (exhibiting longer crossings). Further, many city centers (such as the Loop and River North neighborhoods in Chicago) include long crossings, due to blocks housing civic institutions (e.g., courthouses, museums, parks, etc.) that are separated by wide streets. Similar patterns appear in other cities as well; see Supplementary Materials for analogous maps of Milwaukee, WI, and Austin, TX.

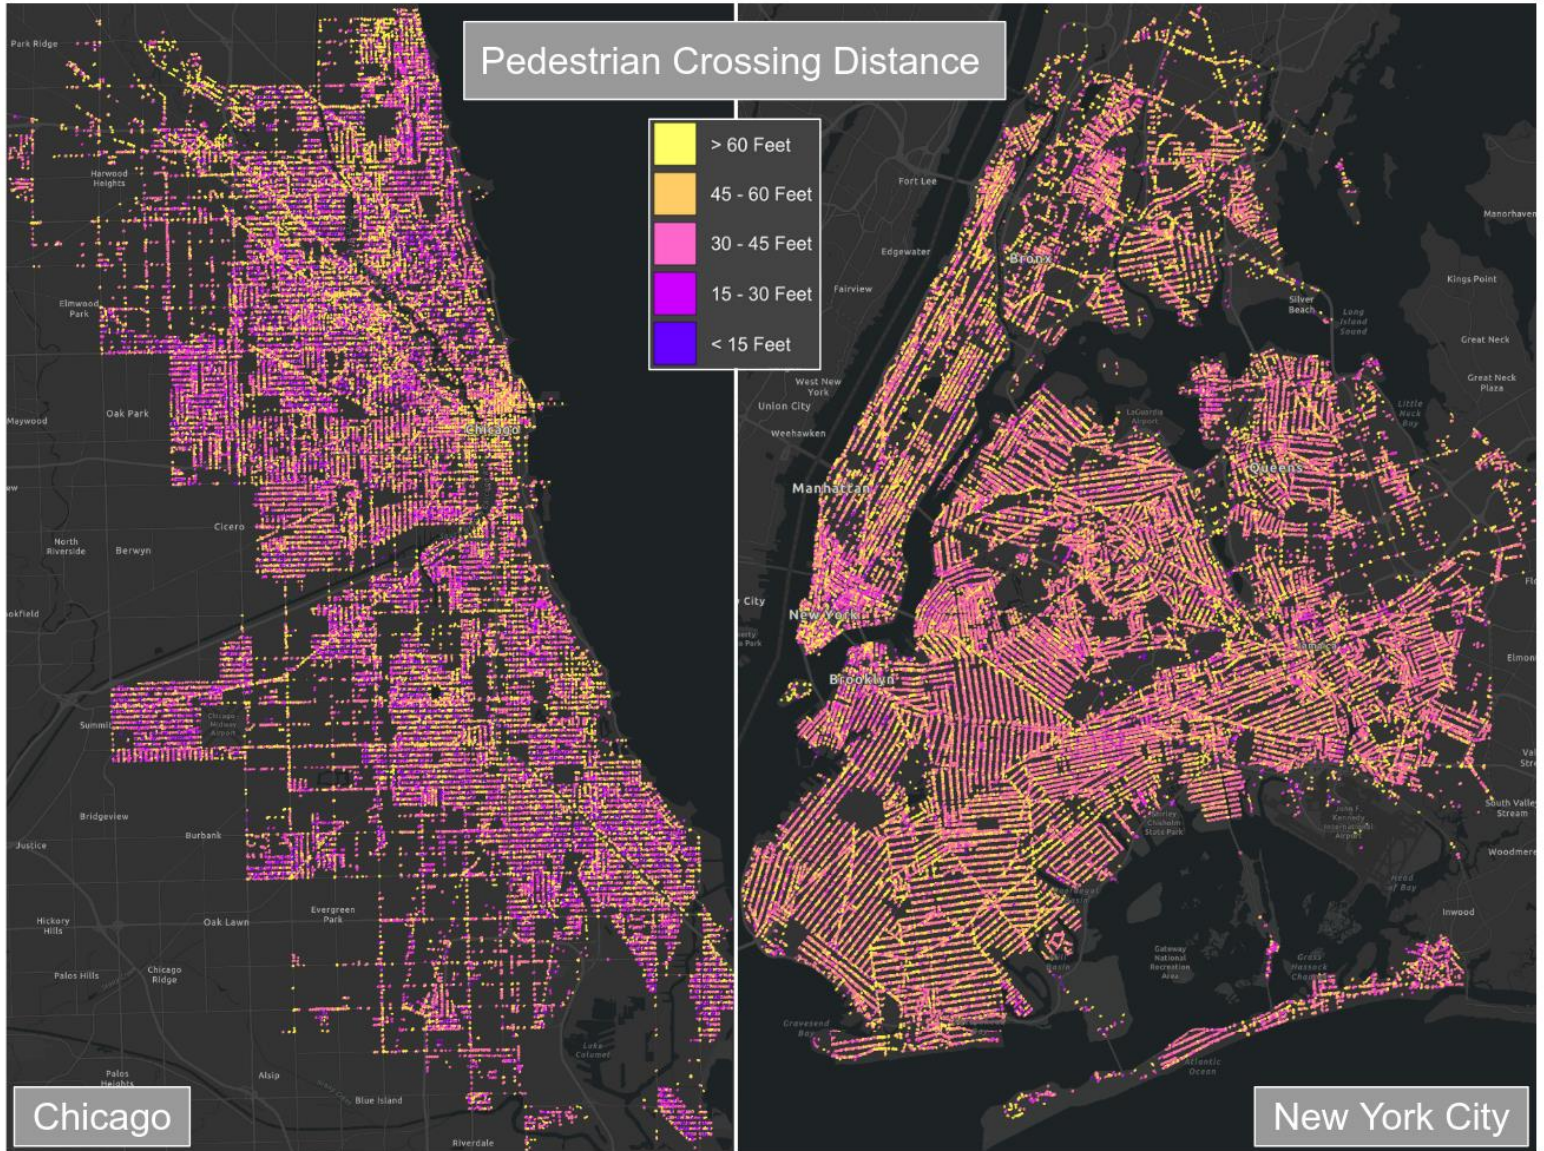


Figure 11. Maps of the spatial distribution of pedestrian crossing distances throughout Chicago (left), and New York City, excluding Staten Island (right) (For additional citywide maps, see Supplementary Materials)

Scaled to the 100 largest cities in the U.S., this also yields inter-city and regional comparisons; shorter median crossing distances clustered in the Northeast and Midwest (particularly in Michigan, Ohio, and Indiana), and longer median crossing distances are far more prevalent in the Southwest and West Coast (especially in the Dallas-Fort Worth area, southern Arizona, Southern Nevada, and Southern California; see Figure 12). To this point, Toledo, OH had the shortest median crossing distance (31 ft, 9.5 m) and North Las Vegas the longest (78 ft, 23.8 m).

median crossing distances. A correlation analysis (Spearman rank correlation) determines that there is a statistically- significant positive relationship between median crossing distance and year of incorporation (ρ rho = 0.510, $P < .0001$).

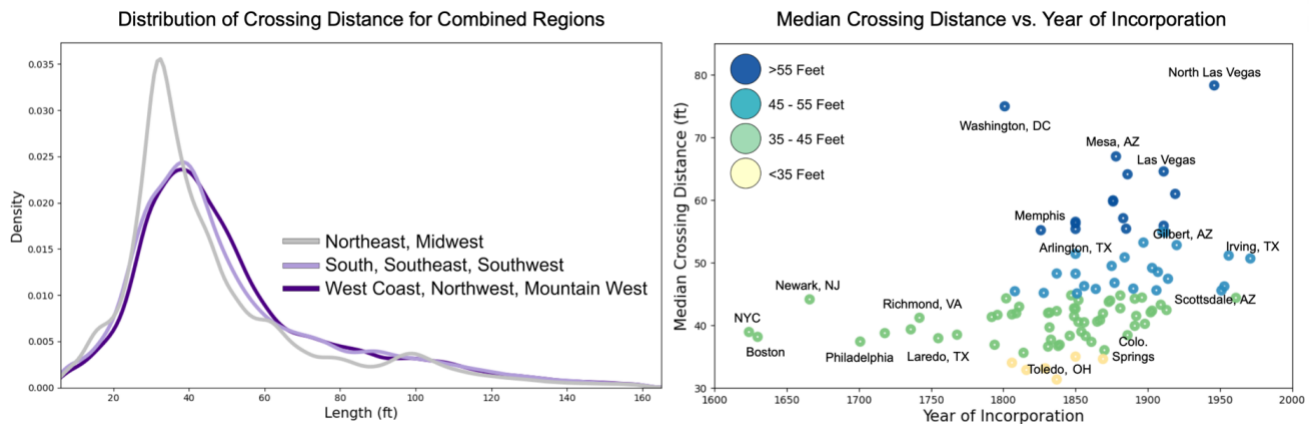


Figure 13. Kernel Density Estimation (KDE) Plot of the distribution of crossing distance, for combined regions (left, 13a). Scatterplot of median crossing distance vs. year of founding for the 100 largest U.S. cities (right, 13b).

4 Discussion and conclusion

Deploying an advanced computer-vision model to segment satellite imagery of street intersections, in combination with information drawn from OpenStreetMap, generated crossing distance measurements throughout the 100 largest American cities (totaling approximately 800,000 crossings), with high accuracy when compared to manually verified data. This establishes practitioner-relevant datasets using publicly-accessible data, and opens up analysis of this fundamental attribute of the pedestrian realm to new scales, both within and across cities. Unlike previous research which applied computer vision to pedestrian infrastructure, this approach captures both marked and *unmarked* crossings, a meaningful advancement given the latter comprises a large portion of total crossings in many American cities (in some cases at close to half of all intersections). In doing so, it demonstrates that crossing distance follows a universal pattern across U.S. cities: left-shifted normal distributions with a peak value count for a given distance that represents the average crossing along neighborhood streets, followed by declining values of increasing crossing distance. However, these datasets also exhibit differences associated with region and incorporation year. Specifically, median crossing distance is longer in cities in the South and West of the country, compared to those in the Northeast and Midwest. Chronologically, median crossing distance is also longer in cities incorporated after 1800, illustrating in a new way the increasing emphasis of car-centric streets and its effect on the pedestrian realm. In so doing, this study builds on scholarship exploring how street features maintain shared characteristics at multiple scales, and relate to the urban history of the United States (Barrington-Leigh & Millard-Ball, 2015; Boeing, 2021; Jiang, 2007).

The implications of this study are threefold. First, in terms of practice, these datasets can provide planners, policymakers, and pedestrian-safety advocates with insights as to patterns in crossing distance in their own cities, potentially pinpointing neighborhoods, corridors, and specific intersections for traffic-calming investments. Indeed, shortening

long crossings can reduce pedestrians' exposure to automobile traffic, and bring streets in line with current best practice in Urban Design. This also creates the opportunity to combine these datasets with other locally relevant attributes, such as speed limits, traffic signals, pedestrian volumes, historic crash data, and neighborhood demographics to further inform decision-making regarding street re-designs. In addition, there is now also the chance to determine if the associations between crossing distance and pedestrian safety, studied previously in smaller geographic contexts, hold at the national scale.

Second, there is likely strong interest in comprehensively measuring pedestrian crossing distance internationally; many nations have set ambitious goals of reducing pedestrian fatalities (Björnberg et al., 2022), and these methods can be adapted to cities globally, especially as high-resolution satellite imagery becomes available in more regions. Indeed, it would be fruitful to understand if such an approach could generate accurate measurements for cities in lower-income countries, including rapidly urbanizing areas where road safety is a significant public health issue (Barrington-Leigh & Millard-Ball, 2015; Boeing, 2021; Jiang, 2007).

Finally, this methodology represents the continued advancement of computer-vision models in the domain of urban research. These segmentation results portends a wide range of applications for identifying features within urban settings and at the sub-block scale. Future research could modify this methodology such that it captures bicycle infrastructure, transit-priority lanes, or other specific road markings. To that end, the full code underlying this analysis is available on GitHub for those who wish to proceed in these, or other directions.

There are several limitations to note. First, the use of crossing edges recorded in OpenStreetMap to seed the grow-cut algorithm inherently means that not all crossings were measured in each city. The comprehensiveness of pedestrian crossings in OpenStreetMap in American cities varies, a product of the volunteer-based nature of the map (Bennett, 2010; Demetriou, 2016). Though coverage was strong when compared to the manually verified dataset for San Francisco (72% of all crossings present in the automated dataset), that figure likely varies across these 100 cases and is difficult to measure. Second, when training SAM via the highest-ranked entropy intersection images, entropy values were not normalized across cities (entropy rankings were separate for each city), which raises the potential issue of comparability and consistency from city to city. To resolve this, future iterations of this analysis approach can consider ranking the entropy of training images collapsed across different cities.

In terms of accuracy, there remain two areas for which future model improvements should be pursued. First, SAM has difficulty accurately segmenting intersection images where tree-cover is heavier, given canopies often visually block underlying sidewalks. For this reason, crossings in leafier neighborhoods are likely less accurate relative to areas with fewer trees, which can result in measurement errors. One potential method of reducing this source of error would be to segment satellite imagery of intersections drawn from *winter* months, when deciduous trees do not have leaves. Second, it is difficult to remove all false-positive crossing edges from these automated datasets, particularly in the case of highly complex and idiosyncratic intersections, some of which include both crossing edges and separate crossing nodes in OpenStreetMap. Further modifying the grow-cut algorithm and post-processing steps will likely lead to improved outcomes on this front.

Lastly, Anchorage, Alaska, the 74th largest city in the United States by population was excluded from analysis because its very high latitude entails limited satellite imagery available from Google Maps. In its place, Frisco, Texas, the 101st largest city in the U.S. by population, was analyzed and incorporated into all subsequent analyses.

Acknowledgments

The authors acknowledge the NYU Center for Data Science, whose capstone project encouraged the development of this study.

Author contribution

The authors confirm their contribution to the paper as follows: conceptualization, writing, and project administration: M. Moran and D. Laefer; formal analysis and methodology: A. Gupta and J. Qian.

Appendix

Supplementary Materials available at <https://doi.org/10.5198/jtlu.2026.2771>.

References

- Abley, S., Turner, S. & Singh, R. (2011). Predicting walkability. *IPENZ Transportation Group Conference*, Auckland, NZ. https://www.researchgate.net/profile/Shane-Turner-5/publication/335339674_Predicting_Walkability/links/5d5fb5a6299bf1f70b05e08d/Predicting-Walkability.pdf
- Agbangba, C. E., Aide, E. S., Honfo, H., & Glèlè Kakai, R. (2024). On the use of post-hoc tests in environmental and biological sciences: A critical review. *Heliyon*, *10*(3), e24131. <https://doi.org/10.1016/j.heliyon.2024.e25131>
- Ahmetovic, D., Manduchi, R., Coughlan, J. M., & Mascetti, S. (2017). Mind your crossings: Mining GIS imagery for crosswalk localization. *ACM Transactions on Accessible Computing*, *9*(4), 1–25. <https://doi.org/10.1145/3046790>
- Albert, A., Kaur, J., & Gonzalez, M. C. (2017). Using convolutional networks and satellite imagery to identify patterns in urban environments at a large scale. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1357–1366. <https://doi.org/10.1145/3097983.3098070>
- Bagwari, N., Kumar, S., & Verma, V. S. (2023). A comprehensive review on segmentation techniques for satellite images. *Archives of Computational Methods in Engineering*, *30*(7), 4325–4358. <https://doi.org/10.1007/s11831-023-09939-4>
- Bailey, L., & Vega-Barachowitz, D. (2012). *Urban street design guide*. NACTO. <https://www.nyc.gov/html/dot/downloads/pdf/2012-nacto-urban-street-design-guide.pdf>
- Barrington-Leigh, C., & Millard-Ball, A. (2015). A century of sprawl in the United States. *Proceedings of the National Academy of Sciences*, *112*(27), 8244–8249. <https://doi.org/10.1073/pnas.1504033112>
- Barrington-Leigh, C., & Millard-Ball, A. (2020). Global trends toward urban street-network sprawl. *Proceedings of the National Academy of Sciences*, *117*(4), 1941–1950. <https://doi.org/10.1073/pnas.1905232116>
- Bennett, J. (2010). OpenStreetMap. Packt Publishing.
- Björnberg, K. E., Hansson, S. O., Belin, M. Å., & Tingvall, C. (2022). *The vision zero handbook: Theory, technology and management for a zero casualty policy*. Springer Nature.
- Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, *65*, 126–139. <https://doi.org/10.1016/j.compenvurbsys.2017.05.004>
- Boeing, G. (2021). Off the grid...and back again? The recent evolution of American street network planning and design. *Journal of the American Planning Association*, *87*(1), 123–137. <https://doi.org/10.1080/01944363.2020.1819382>
- Cova, E. (2024, April 3). *Pedestrian fatalities at historic high*. Smart Growth America. <https://smartgrowthamerica.org/pedestrian-fatalities-at-historic-high/>
- Demetriou, D. (2016). Uncertainty of OpenStreetMap data for the road network in Cyprus. In *Fourth International Conference on Remote Sensing and Geoinformation of the Environment (RSCy2016)*, 9688, 43–52. <https://doi.org/10.1117/12.2239612>
- Duncan, S., Chacra, A., Prakash, A., & Prati, F. (2016). *Global street design guide*. Global Designing Cities Initiative. <https://globaldesigningcities.org/publication/global-street-design-guide/>
- Dwyer, B., Nelson, J., & Hansen, T. (2024). *Roboflow (Version 1.0)*. Roboflow. <https://roboflow.com>

- Feuer, W., Smith, P. Aun, E., Browne, P., Cao, J., Colasacco, E., Gudjonsson, S., Mezey, C., Micka-Maloy, K., & Pettinati, N. (2020). *Street design manual* (3rd ed.). New York City Department of Transportation.
- Haverkamp, D. S. (2002). Extracting straight road structure in urban environments using IKONOS satellite imagery. *Optical Engineering*, 41(9), 2107–10. <https://doi.org/10.1117/1.1496785>
- Islam, A., Mekker, M., & Singleton, P. A. (2022). Examining pedestrian crash frequency, severity, and safety in numbers using pedestrian exposure from Utah traffic signal data. *Journal of Transportation Engineering, Part A: Systems*, 148(10), 04022084. <https://doi.org/10.1061/JTEPBS.0000737>
- Jiang, B. (2007). A topological pattern of urban street networks: Universality and peculiarity. *Physica A: Statistical Mechanics and its Applications*, 384(2), 647–655. <https://doi.org/10.1016/j.physa.2007.05.064>
- Kadali, B. R., & Vedagiri, P. (2015). Evaluation of pedestrian crosswalk level of service (LOS) in perspective of type of land-use. *Transportation Research Part A: Policy and Practice*, 73, 113–124. <https://doi.org/10.1016/j.tra.2015.01.009>
- Larsen, S.Ø., Koren, H., & Solberg, R. (2009). Traffic monitoring using very high resolution satellite imagery. *Photogrammetric Engineering & Remote Sensing*, 75(7), 859–869. <https://doi.org/10.14358/PERS.75.7.859>
- Li, M., Sheng, H., Irvin, J., Chung, H., Ying, A., Sun, T., Ng, A.Y., & Rodriguez, D.A. (2023). Marked crosswalks in US transit-oriented station areas, 2007–2020: A computer vision approach using street view imagery. *Environment and Planning B: Urban Analytics and City Science*, 50(2), 350–369. <https://doi.org/10.1177/23998083221112157>
- Ma, X., Zhang, X., Pun, M. O., & Huang, B. (2024). MANet: Fine-tuning segment anything model for multimodal remote sensing semantic segmentation. arXiv. <https://doi.org/10.48550/arXiv.2410.11160>
- Martin, A. (2006). *Factors influencing pedestrian safety: A literature review*. PPR-241. TRL Limited. <https://content.tfl.gov.uk/factors-influencing-pedestrian-safety-literature-review.pdf>
- McKight, P. E., & Najab, J. (2010). Kruskal-Wallis test. In *The Corsini Encyclopedia of Psychology*. John Wiley & Sons, Ltd. <https://doi.org/10.1002/9780470479216.corpsy0491>
- Monsere, C., Figliozzi, M., Kothuri, S. M., Razmpa, A., & Hazel, D. (2016). *Safety effectiveness of pedestrian crossing enhancements* (SPR 778). Oregon Department of Transportation, FHWA, USDOT. https://pdxscholar.library.pdx.edu/cgi/viewcontent.cgi?article=1353&context=cengin_fac
- Moran, M. E. (2020). What’s your angle? Analyzing angled parking to aid bike-network planning. *Environment and Planning B: Urban Analytics and City Science*, 48(7), 1912–1925. <https://doi.org/10.1177/2399808320954205>
- Moran, M. E. (2022). Where the crosswalk ends: Mapping crosswalk coverage via satellite imagery in San Francisco. *Environment and Planning B: Urban Analytics and City Science*, 49(8), 2250–2266. <https://doi.org/10.1177/23998083221081530>
- Moran, M. E., & Laefer, D. F. (2024). Multiscale analysis of pedestrian crossing distance. *Journal of the American Planning Association*, 1–15. <https://doi.org/10.1080/01944363.2024.2394610>
- Oscó, L. P., Wu, Q., de Lemos, E. L., Gonçalves, W. N., Ramos, A. P. M., Li, J., & Marcato, J. (2023). The segment anything model (SAM) for remote sensing applications: From zero to one shot. *International Journal of Applied Earth*

- Observation and Geoinformation*, 124, 103540.
<https://doi.org/10.1016/j.jag.2023.103540>
- Ravi, N., Gabeur, V., Hu, Y. T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Vasedev Alwala, K., Carion, N., Wu, C., Girshick, R., Dollar, P., & Feichtenhofer, C. (2024). SAM 2: Segment anything in images and videos. Meta. <https://ai.meta.com/research/publications/sam-2-segment-anything-in-images-and-videos/>
- Ren, S., Luzzi, F., Lahrachi, S., Kassaw, K., Collins, L. M., Bradbury, K., & Malof, J. M. (2024). Segment anything, from space? *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 8355–8365.
https://openaccess.thecvf.com/content/WACV2024/html/Ren_Segment_Anything_From_Space_WACV_2024_paper.html
- Schneider, R. J. (2020). United States pedestrian fatality trends, 1977 to 2016. *Transportation Research Record*, 2674(9), 1069–1083.
<https://doi.org/10.1177/0361198120933636>
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/doi:10.1002/j.1538-7305.1948.tb01338.x>
- Verma, R., & Ukkusuri, S. V. (2023). Crosswalk detection from satellite imagery for pedestrian network completion. *Transportation Research Record*, 2678(7), 845–856.
<https://doi.org/10.1177/03611981231210545>
- Vezhnevets, V., & Konouchine, V. (2005). Growcut: Interactive multi-label Nd image segmentation by cellular automata. In *Proceedings of Graphicon, 1*, 150–156.
<https://www.graphicon.ru/html/2005/proceedings/papers/VezhntvetsKonushin.pdf>
- Wilmot, K., & Purcell, C. (2022). Why are older adults more at risk as pedestrians? A systematic review. *Human Factors*, 64(8), 1269–1291.
<https://doi.org/10.1177/0018720821989511>