Where (not) to Cross the Street

Leah van Oorschot[1], Dimitris Michailidis[1], Niek IJzerman[2] Shayla Jansen[2], and Diederik Roijers[2,3]

[1] Universiteit van Amsterdam, leahkarlijn@hotmail.com & d.michailidis@uva.nl
[2] City of Amsterdam, [shayla.jansen,n.ijzerman,d.roijers]@amsterdam.nl
[3] Vrije Universiteit Brussel, diederik.roijers@vub.be

Keywords: Crosswalk Detection · Point Clouds · Accessible Navigation.

Accessible crossings are crucial for determining suitable accesible routes. While locations where people cross are often determined by context, crosswalks are generally preferred as they are visible and grant priority to pedestrians [1]. In this thesis, we contribute to more inclusive navigation solutions by developing a novel framework to automatically detect crosswalks.

Most crosswalk detection studies use aerial imaging or point cloud data, but both data sources have limitations when used independently. Satellite imagery can be obscured by environmental factors, while point clouds are computationally intensive. We integrate both data sources to enhance crosswalk detection accuracy and scalability.

We build off Tile2Net [2] which extracts sidewalk, crosswalk, and footpath polygons from aerial images. However, Tile2Net is trained on North American data, which we found reduces its applicability to diverse urban environments like Amsterdam. Therefore we make the following contributions: after testing Tile2Net in a new urban context and analysing its performance, we develop a framework to enhance its performance, tailored to Amsterdam's infrastructure. By doing so we demonstrate that automatic crosswalk detection methods can be adapted to various urban contexts.

The method exists of several consecutive steps. The initial step extracts crosswalk polygons from aerial images using Tile2Net. These initial polygons often contain a high number of false positives and inaccuracies, necessitating further refinement. The next step matches these polygons to LiDAR point cloud data. Next, points with a reflective index below a set threshold are removed, isolating those that belong to road markings. The DBSCAN algorithm is then used to cluster these points based on spatial proximity, separating crosswalk stripes when multiple are present in one polygon. Next, these clusters are expanded using the cyclic KDTree algorithm to ensure completeness. Each cluster is then analyzed for shape and size, retaining only those that match the rectangular shape and size criteria for Dutch crosswalk standards. Clusters meeting these criteria are identified as crosswalk stripes, while others are disregarded.

We use a third dataset, Project Sidewalk (PSW), to validate the framework. Project Sidewalk crowdsources pedestrian accessibility information, including crosswalks, using Google Street View. We use precision and recall to assess the framework's accuracy. Finally, the exact areas of the crosswalks polygons are validated against manually labeled crosswalk polygons using the Jaccard index.

Table 1 shows the results, with precision increasing to 100% across four distinct city areas. This indicates that the framework effectively filters out false positives, making predictions more reliable. However, there is a minor 3% decrease in recall due to crosswalk polygons identified by Tile2Net being falsely filtered out. This is due to polygons that are too small to process accurately and noise in the point cloud data. Despite this, the substantial increase in precision outweighs the recall decrease, ensuring accurate crosswalk identification.

Area	**Metric**	**Tile2Net**	**Processed**
Venserpolder	PSW	32	32
	True Positive	24	23
	False Positive	46	0
	False Negative	8	9
	Precision	0.343	1.0
	Recall	0.75	0.719
De Baarsjes	PSW	62	62
	True Positive	47	47
	False Positive	25	0
	False Negative	15	15
	Precision	0.713	1.0
	Recall	0.805	0.805
Osdorp	PSW	107	107
	True Positive	99	96
	False Positive	31	0
	False Negative	8	11
	Precision	0.758	1.0
	Recall	0.907	0.897
Grachtengordel	PSW	69	69
	True Positive	49	45
	False Positive	25	0
	False Negative	20	24
	Precision	0.662	1.0
	Recall	0.710	0.652

Table 1. Crosswalk detection in 4 Amsterdam city districts.

The proposed framework improves crosswalk detection using aerial imagery and LiDAR point clouds, achieving perfect precision and reliably identifying crosswalks. These results suggest that models trained on different cities can be valuable when refined appropriately, significantly contributing to urban mapping and pedestrian safety by providing precise crosswalk polygons for applications like accessible navigation.

In future work, we aim to address recall limitations as the current framework is not capable of identifying crosswalks if they are not located by Tile2Net. Additionally, the framework should be applied to diverse urban environments to validate and enhance its performance further.

Acknowledgements This research was in part supported by the European Union's Horizon Europe research and innovation programme under grant agreement No 101120406. DM is supported by the Innovation Center for AI (ICAI, The Netherlands).

References

1. Ahmetovic, D., Manduchi, R., Coughlan, J.M., Mascetti, S.: 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. pp. 251–258 (2015)
2. Hosseini, M., Sevtsuk, A., Miranda, F., Cesar Jr, R.M., Silva, C.T.: Mapping the walk: A scalable computer vision approach for generating sidewalk network datasets from aerial imagery. Computers, Environment and Urban Systems **101**, 101950 (2023)