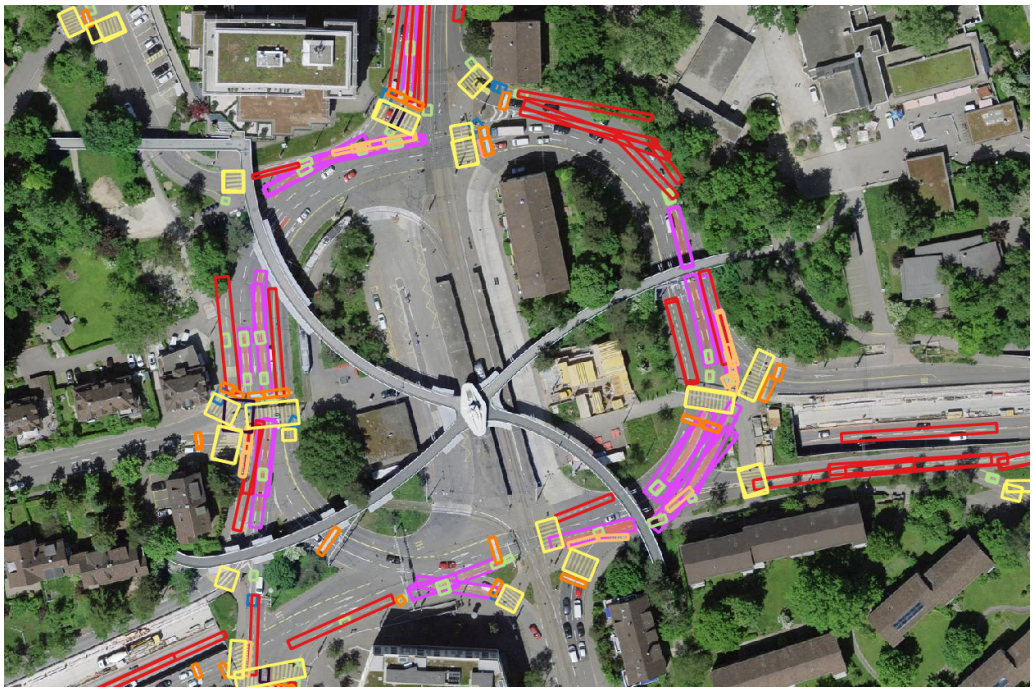

Mapping Cycling-Specific Infrastructure Using Object Detection on Remotely Sensed Images

Sebastiano P. Papini

David Zani

STRC Conference Paper 2025 **April 28, 2025**

STRC | **25th Swiss Transport Research Conference**
Monte Verità / Ascona, May 14-16, 2025

Mapping Cycling-Specific Infrastructure Using Object Detection on Remotely Sensed Images

Sebastiano P. Papini
D-MTEC
ETH Zurich
spapini@ethz.ch

David Zani
D-BAUG
ETH Zurich
zani@ibi.baug.ethz.ch

April 28, 2025

Abstract

The new Cycle Route Act (2020) aims to significantly expand Swiss cycling infrastructure within a decade. However, in stark contrast to this policy agenda, from a research perspective the lack of comprehensive data on cycling infrastructure is a significant barrier, particularly for studying the impact of infrastructure modifications on a junction-specific granularity (e.g. on induced cycling demand or crash risks). Existing data suffers from at least one of the following deficiencies: (1) fragmentation along administrative borders, (2) purely link-oriented data without information on junctions, (3) missing information on historical infrastructure changes, and (4) inadequate categorization from a road design perspective. The contribution of this paper is therefore threefold. First, a object detection method is utilized on aerial imagery to generate a dataset that addresses all four previously mentioned issues. For this purpose, we train a YOLOv8 (You Only Look Once, version 8) model, a common deep learning architecture for object detection, to detect ten different cycling-specific infrastructure features. Second, it is demonstrated that the method is valid by comparing a subset of the resulting data set to a external communal level data set. Third, the overall historical development of cycling-specific infrastructure in the ten largest Swiss agglomerations is discussed.

Keywords

Cycling Infrastructure; Switzerland; Object Detection; Remote Sensing; Transport Infrastructure

Contents

List of Tables

List of Figures

1 Introduction

In 2018, Swiss voters approved a constitutional amendment on cycling by a large majority. The resulting law, the new Cycle Route Act, came into force in 2020. It aims to significantly expand cycling infrastructure. Consequently, the Office for Spatial Development (ARE) projects the modal share of bikes to double until the year 2050. Considering the prevailing political agenda, the lack of data on bicycle infrastructure is striking, particularly from a research perspective. Researchers seeking to use such data for the study of induced cycling demand (e.g. Song *et al.*, 2017; Wysling and Purves, 2022; Pritchard *et al.*, 2019; Hwang and Guhathakurta, 2023) or crash risks (e.g. Tait *et al.*, 2022; Vandenbulcke *et al.*, 2014; Smith and Welsh, 1988; Reynolds *et al.*, 2009; Van Petegem *et al.*, 2021; DiGioia *et al.*, 2017), which are often presumed to be associated with enhanced infrastructure, face significant data constraints.

Although certain communes or cantons have made notable efforts to map their streets according to bicycle accessibility, a federally consistent data set remains unavailable. If data is available at all, it suffers from at least one of three shortcomings:

1. It is predominantly link-oriented, such as OpenStreetMap (OSM). A considerable number of infrastructure modifications focus on junctions and adjacent street segments. Consequently, these alterations are frequently not documented. However, these nodes are the locations where the impacts may be the most substantial, as evidenced by the fact that most bike crashes occur near junctions.
2. A longitude (time) component is almost always absent, which complicates the formulation of a robust identification strategy for research designs related to bicycle infrastructure (Vanparijs *et al.*, 2015).
3. It lacks detailed records of concrete infrastructure measures that comprise the possible alternatives given the physical, economical, and legal constraints in the road design process. This limitation hinders the development of research that could help planners make informed decisions based on scientific evidence (Sohail *et al.*, 2023).

The objective of this paper is threefold. First, a deep learning-enabled object detection method is utilized on aerial imagery to generate a data set that satisfies all three criteria. Second, it is shown that the result is reliable. Third, the overall development of cycling-specific infrastructure in the largest Swiss agglomerations is analyzed. For this purpose, we train YOLOv8 (You Only Look Once, version 8), a predefined deep learning architecture for object detection, to detect ten different cycling-specific infrastructure features.

1.1 Related Work

Cycling infrastructure databases are an important element for cycling-related research. Past efforts to catalog cycling infrastructure have used OSM and publicly available datasets, including aerial imagery from Google (Ferster *et al.*, 2023). Machine learning methods have been widely used to identify road infrastructure and provide the necessary data for subsequent analyses, including intersection designs and their safety (Wijnands *et al.*, 2021), road marking extraction (Jin *et al.*, 2012), YOLO applications to identify lane markings for improved road safety and management (Antwi *et al.*, 2024a) and to locate school zones from satellite imagery (Antwi *et al.*, 2024b), and using Google street view images to map and manage street signs (Campbell *et al.*, 2019; Habibi Aghdam *et al.*, 2016). The work presented in this paper builds on these methods by incorporating the component of time, analyzing how cycling infrastructure has evolved.

2 Data

Our study area focuses on the major agglomerations in Switzerland, as dense urban areas and their suburban surroundings exhibit the highest share of cycling modal splits today. Their high share of short trips by other modes offers great potential for further shifts toward cycling (Bundesamt für Statistik and Bundesamt für Raumentwicklung, 2023). In order to capture these areas well, we take the ten largest administrative communes by population and draw a bounding box around their agglomeration. Figure 1 displays the chosen bounding boxes that represent our study area in orange. Although some of these bounding boxes extend into neighboring countries, we focus only on the Swiss territory within these bounding boxes.

We utilize the natural color orthophoto mosaic SWISSIMAGE provided by the Federal Office of Topography (swisstopo). The timing of recording aerial imagery is characterized by an alternating pattern, with different Swiss regions undergoing imaging in successive years. This results in a coverage cycle that is approximately three-yearly since 1998. The imagery has a ground resolution of 0.1 meters in images taken since 2018, 0.25 meters since 2005, and 0.5 meters for earlier images, as indicated in Table 3.

To generate image-label pairs for model training, over 40 square kilometers of aerial imagery from each of the ten agglomerations and all available years from 1998 to 2023 and

Figure 1: Map of the study area

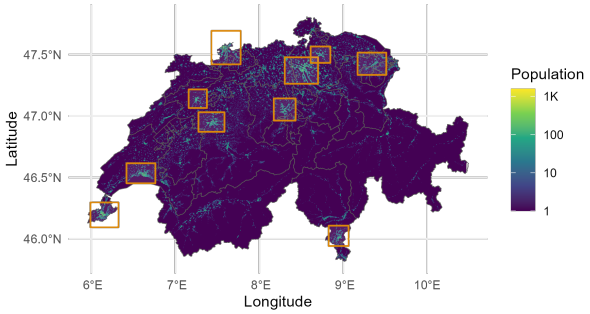

therefore in each of the three available resolutions, were manually labeled using QGIS. In each image, all relevant cycling infrastructure was captured with a manually drawn polygon. The ten relevant cycling infrastructure objects are crosswalks, shark teeth (indicating yielding relations), circles (roundabouts), bike stopping spaces ("Veloaufstellflächen", where cyclists can stop ahead of cars at traffic lights), retaining bars (indicating stopping lines at intersections), pedestrian islands, bike pictograms, red colored bike lanes, advanced retaining bars ("vorgezogene Haltebalken", where bike lanes advance closer to the intersection than car lanes), and (yellow) bike lanes. Each label was validated by a second annotator in line with the four-eyes principle. This resulted in over 15,175 labeled instances (see Tables 4 and 5 in Appendix B for an overview).

The labels (i.e., polygons) were transformed into minimum rotated rectangles. The resulting directed bounding boxes formed the basic image-label pairs necessary for training the YOLO model. Portions of the image-pair labels were reserved for model validation.

3 Method

A pretrained YOLOv8 model was used for the analysis. It was trained on the manually labeled instances of cycling infrastructure. One model was trained for all of the possible object classes (infrastructure types) and resolutions (10, 25, 50 centimeters), as this was found to perform better than one specific model for each class or resolution. There were 4782 (∼80%) images for training, 543 for validation (∼10%), and 520 for testing (∼10%).

One input tile consist of a 102×102 meters (or 1024×1024 pixels) cutout of the aerial images. Orthophotos with 25cm and 50cm resolution were up-sampled to ensure each tile had the same number of pixels. A sliding window was implemented to cover all agglomeration bounding boxes, with an 62 meter overlap to minimize instances missed due to placement at the edge of a tile.

After obtaining the model predictions, some post-processing steps were implemented. The predictions from the model provided instances of cycling infrastructure in all of the ten agglomerations over all of the available aerial imagery years (1,127,885 detected instances in total over 23 years). All detected instances with all confidence levels were kept and used in the post-processing steps. The intersection over union (IoU) was computed for all instances of an object class, over all years. The polygons of any objects of the same class with an IoU > 0.5 were merged into a single object.

Preliminary inspection of the results showed some significant "gaps" in detected objects, i.e., objects that were detected in one year but not the following year. Manual checks proved most of these cases to be a result of some type of occlusion. For example, an object was detected in one year but then not visible in the next year because it was covered by trees, shadows, cars, etc. This issue of occlusion is common when working with aerial imagery. To avoid this problem, the missing detections at the agglomeration level were "infilled": if an object was detected in one year, not detected in the next year, and then again detected in the following year, the missing case was overridden (based on a threshold IoU value).

Lastly, the results were aggregated to the intersection level, and only for "major intersections". To define major intersections, intersections of OSM lines were first buffered by 35 meters. All line crossings within a proximity of 25 meters were clustered, and the clustered buffers then merged into a single geometry. For each resulting geometry, the traffic volumes from the national traffic model passing through the geometry were summed. The intersections with the top 25 % highest traffic volumes were defined as "major". These are therefore intersections where cycling infrastructure is the most needed, as opposed to smaller intersections. While this step reduces some of the available detail of the results, it also improves the accuracy when correctly classifying infrastructure. For example, the model may not be able to correctly predict all instances of cycling infrastructure at every arm of an intersection, but it very likely is able to predict whether an intersection overall has a case of a specific cycling infrastructure. For the purposes of this work, this level of detail is sufficient and worth the additional prediction accuracy.

3.1 Performance Metrics

The performance metrics used to estimate the quality of the model are precision (the ratio of true positives to the sum of all positives), recall (the ratio of true positives to the sum of true positives and false negatives), and F1 (the geometric mean of the former measures). Performance metrics were computed for the overall model, for each of the three resolutions, and for each of the ten agglomerations.

Model performance was also analyzed using external validation datasets: OSM and Zurich city's publicly available data on its bike lanes. By comparing the roads where the trained model predicted different types of bike lanes and where OSM and the city actually has those types of bike lanes, the model's accuracy was validated against data that were not used in the model training process at all (either during image labeling, model training, or model validation).

4 Results

An example of the model's output is shown in Figures 3, 4, 5, and 6 in Appendix C. Table 1 shows an overview of the model performance. The values are F1 scores, in parentheses are precision and recall. Each column is a subset of the test set. Performance for each of the ten objects varies, from poor (0.42 for bike lanes) to good (0.83 for crosswalks). The performance also varies by image resolution. Interestingly, images with better resolution (10 centimeters) do not always produce better predictions than the lower resolutions (25 centimeters) (e.g., for bike stopping spaces). There is also variation in the model performance by agglomeration. Only object classes with at least four instances were used to evaluate model performance, which explains why some resolutions and agglomerations have missing values. Overall, the model does not perform well for a range of objects. Some aggregation is necessary to improve the results and their usefulness for further analysis.

In part due to the relatively poor performance of the model at the individual object level, and since that level of detail was not necessary for the goals of this paper, the predictions were aggregated to the intersection level. For each major intersections, the presence of the individual object classes was measured. For example, if a bike pictogram was predicted anywhere within the "major intersection" polygon, then that intersection was labeled as having a bike pictogram. In this way, even if individual bike pictograms were not

Table 1: Model performance

	Full Test Set	10	25	50	Basel	Bern	Biel	Genf	Lausanne	Lugano	Luzern	St. Gallen	Winterthur	Zurich
crosswalk	0.83 (0.80, 0.86)	0.84 (0.81, 0.88)	0.85 (0.81, 0.89)	0.80 (0.78, 0.82)	0.80 (0.77, 0.82)	0.86 (0.80, 0.93)	0.81 (0.79, 0.83)	0.84 (0.82, 0.87)	0.84 (0.83, 0.85)	0.89 (0.88, 0.90)	0.85 (0.82, 0.88)	0.81 (0.86, 0.77)	0.94 (0.89, 1.00)	0.81 (0.74, 0.88)
shark_teeth	0.67 (0.72, 0.63)	0.78 (0.80, 0.76)	0.74 (0.71, 0.76)	0.10 (0.20, 0.07)	0.50 (0.71, 0.38)	0.75 (0.81, 0.70)	0.67 (0.53, 0.91)	-	0.78 (0.78, 0.78)	0.83 (0.75, 0.92)	0.55 (0.50, 0.62)	0.67 (0.82, 0.56)	0.92 (0.92, 0.92)	0.65 (0.67, 0.63)
circle	0.66 (0.71, 0.62)	0.91 (1.00, 0.83)	0.71 (0.62, 0.83)	-	-	0.71 (0.62, 0.83)	-	-	-	-	-	-	-	-
bike_stopping_space	0.66 (0.70, 0.63)	0.63 (0.77, 0.53)	0.72 (0.64, 0.82)	-	-	0.67 (0.80, 0.57)	-	-	0.50 (0.57, 0.44)	-	-	0.80 (0.67, 1.00)	-	-
retaining_bar	0.63 (0.69, 0.58)	0.65 (0.74, 0.57)	0.64 (0.64, 0.63)	0.60 (0.68, 0.53)	0.48 (0.62, 0.39)	0.60 (0.57, 0.62)	0.75 (0.74, 0.76)	0.64 (0.70, 0.58)	0.68 (0.81, 0.59)	0.60 (0.81, 0.48)	0.49 (0.52, 0.46)	0.77 (0.77, 0.77)	0.82 (1.00, 0.70)	0.70 (0.68, 0.72)
pedestrian_island	0.54 (0.51, 0.56)	0.55 (0.45, 0.71)	0.59 (0.50, 0.71)	0.48 (0.60, 0.39)	0.41 (0.50, 0.35)	0.67 (0.57, 0.80)	0.37 (0.29, 0.50)	0.40 (0.36, 0.44)	0.57 (0.55, 0.59)	0.75 (0.60, 1.00)	0.75 (0.82, 0.69)	0.78 (0.75, 0.82)	0.44 (0.33, 0.67)	0.50 (0.41, 0.65)
bike_pictogram	0.51 (0.61, 0.44)	0.57 (0.61, 0.54)	0.48 (0.59, 0.41)	-	0.40 (0.66, 0.28)	0.43 (0.65, 0.32)	0.32 (0.25, 0.45)	0.61 (0.73, 0.52)	0.72 (0.84, 0.63)	0.40 (0.30, 0.60)	0.46 (0.44, 0.48)	0.50 (0.60, 0.43)	1.00 (1.00, 1.00)	0.48 (0.55, 0.42)
red_colored_bike_lane	0.45 (0.60, 0.36)	0.46 (0.58, 0.38)	0.58 (0.61, 0.55)	-	-	0.60 (0.50, 0.75)	-	0.52 (0.62, 0.45)	0.33 (0.50, 0.25)	-	0.55 (0.62, 0.50)	-	-	-
advanced_retaining_bar	0.43 (0.44, 0.43)	0.37 (0.36, 0.38)	0.55 (0.52, 0.58)	-	0.15 (0.25, 0.11)	0.50 (0.43, 0.60)	0.62 (0.62, 0.62)	-	-	-	0.64 (0.62, 0.67)	0.29 (0.25, 0.33)	-	0.33 (0.33, 0.33)
bike_lane	0.42 (0.44, 0.41)	0.43 (0.44, 0.41)	0.49 (0.48, 0.50)	0.29 (0.34, 0.25)	0.18 (0.24, 0.15)	0.55 (0.56, 0.54)	0.43 (0.61, 0.33)	0.35 (0.44, 0.29)	0.47 (0.43, 0.53)	0.45 (0.32, 0.73)	0.51 (0.56, 0.47)	0.41 (0.46, 0.37)	0.56 (0.44, 0.76)	0.41 (0.41, 0.40)

identified by the model, one true positive would suffice to correctly identify the presence of pictograms at that intersection. This is feasible since 1) it is likely that the presence of one instance of cycling infrastructure at an intersection indicates the presence of more instances, and 2) this study aims to explore the evolution of cycling infrastructure at intersections in general, and not for individual arms of an intersection.

The model predictions were compared to two external data sets not used in training: OSM and Zurich city's official cycling network data. These networks were first buffered, and predicted objects above a certain IoU threshold with those networks were considered to be true positives. Table 2 therefore shows the precision and recall of the model's predicted bike lanes relative to the OSM and Zurich city data.

Table 2: Comparison of model performance relative to external Zurich bike lanes data sets. OSM = OpenStreetMap

Dataset	**Recall (%)**	**Precision (%)**
Zurich City Data	73.8	76.6
Zurich OSM Data	71.5	68.5

5 Findings and Discussion

In general, all agglomerations show an increasing quantity of cycling infrastructure over time (see Figure 7 in the appendix). However, some of the increases are likely due to improvements in the image resolution, specifically in the years when 25-centimeter resolutions first become available (2005 and after). The infrastructure types most directly related to cycling show the strongest increases, especially since 2015 (bike lanes, bike

Figure 2: Share of major intersections without cycling infrastructure over time. Note the different scales on the y-axis

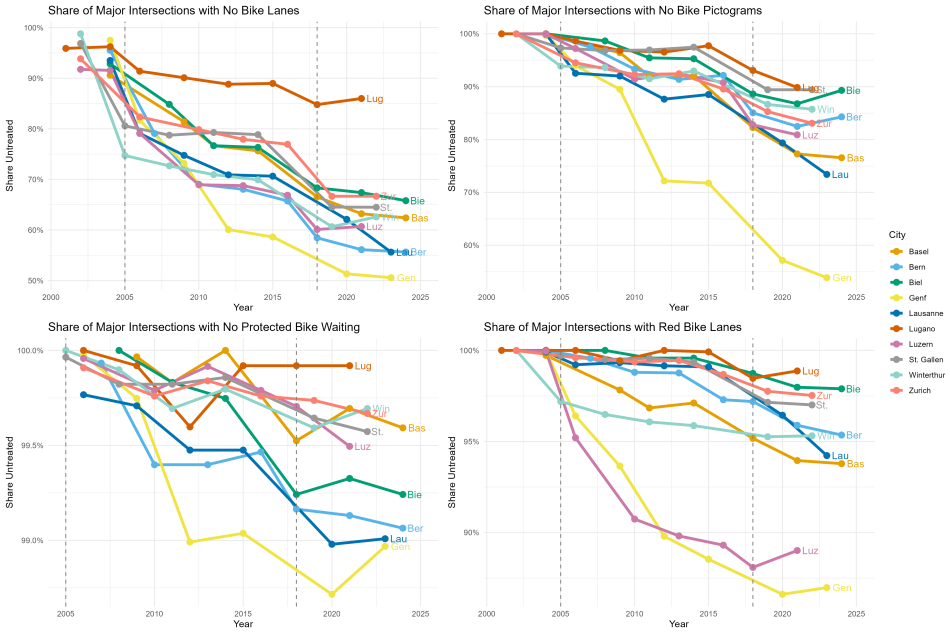

pictograms, and red colored bike lanes). More recently popularized bicycle infrastructure, such as the advanced retaining bar and bike stopping space, see little or only more recent improvement. These quantities are expected to increase in the near future as aerial images are updated.

Figure 2 shows the share of major intersections in each agglomeration that do not have at least one instance of a specific type of cycling infrastructure. In other words, this figure shows the share of intersections that do not have cycling-friendly infrastructure. For all agglomerations, over time, the share of such intersections is decreasing. Notable jumps are present when the resolution improves, due to the improved image quality and not necessarily because more infrastructure was implemented in those years. The plots for the different infrastructure types show that some infrastructure is more widely implemented (bike lanes and pictograms), while other types are relatively more rare (protected waiting areas and red bike lanes).

Geneva shows the most significant improvement and has the lowest share of intersections without cycling infrastructure, outperforming all other agglomerations. Luzern shows a relatively strong reduction of intersections without red bike lanes, but otherwise performs similarly to other agglomerations. Lugano has the highest share of intersections without cycling infrastructure (85 % or higher). Figures 8 and 9 (appendix) show the infrastructure

developments over time for each agglomeration by settlement type and road ownership.

These results show a significant expansion of cycling infrastructure in all agglomerations over time, regardless of settlement type and road ownership. It is interesting to note the relative performance of the agglomerations, where Geneva can be seen to have relatively more infrastructure than the other agglomerations, and national roads are also relatively more advanced in their cycling infrastructure. This makes sense, as federal roads are (formally) governed by the same rules for all agglomerations, while cantonal and communal practices may differ widely. However, national roads intersections are usually larger and more complex, so the presence of cycling infrastructure may be an effect of intersection size and not directly caused by the owners' decisions. The wide difference between agglomerations regarding cantonal and communal road cycling infrastructure indicates more efforts are needed to ensure cycling infrastructure is being implemented at all levels of government to achieve national cycling modal share goals.

This sort of historical data on cycling infrastructure is also useful for further analyses, such as crash risk analysis. Historical crashes cannot always be matched to the infrastructure at the location of the crash, since that infrastructure is likely to change over time. With the historical results provided by this work, a matching of crashes to infrastructure is more accurate, beyond what is currently possible using, e.g., OSM repositories. OSM data quality changes over time with changing tags and completeness, while the aerial imagery is not affected by these shortcomings. Combining different datasets for more complete infrastructure overview may be useful. Zurich city's own bike lane data is incomplete and required additional manual labeling by the authors to make a comparison with the model predictions possible. Even though model performance was poor for bike lanes (precision of about 40 %), the precision relative to the Zurich city data was significantly better (about 77 %) (see Table 2).

6 Limitations, Outlook, and Conclusion

While the results in this work provide useful information for planners and safety-related studies, there are some shortcomings. First, the machine learning model is not able to detect separate cycling infrastructure, i.e., cycling paths. This limitation does not strongly affect analyses in Swiss agglomerations, since cycle paths are rare in these locations. Other agglomerations, such as in the Netherlands, would be more strongly affected. The

historical overview is somewhat limited to the years after the 25-centimeter aerial images were introduced (2005 and later), as the 50-centimeter resolution did not allow labeling or identification of most cycling infrastructure classes. The resolution changes further affected the prediction results by introducing artificial jumps in the amount of predictions. Object occlusion through trees, shadows, cars, etc. also negatively affected the predictions, though the process of infilling gaps in the predictions mitigated this. The model performance was overall relatively poor considering the prediction of individual infrastructure classes, especially linear infrastructures (bike lanes) that could span over several image tiles. A different modeling approach may be necesary to better predict linear elements. Lastly, though the model is able to predict presence of cycling infrastructure at the intersection level, mere presence does not explain the possible impacts on infrastructure safety or usage. Further data would be necessary to explore the actual contribution made (either to attractiveness or safety) by expanding cycling infrastructure.

Future work could improve the model for better predictions, specifically of linear elements such as bike lanes. Subsequent years of images could be grouped together to reduce the error from object occlusion or false negatives. Additional validation with other external datasets would be useful to improve the model reliability. Future work should also begin exploring how these data can be used in analyses of cycling mode shift and subjective and objective safety.

This work showed how machine learning on aerial imagery can provide a useful overview of cycling infrastructure development in urban areas. The results are useful to planners and policymakers who are interested in achieving goals pertaining to increased cycling modal share. The results also provide useful groundwork for further analyses of subjective and objective cycling safety, going beyond what is currently available through open source data such as OSM and official government data.

7 Acknowledgements

The authors would like to thank Marius Jauer for supporting the image labeling and Fengjiao Ji for supporting the model training and prediction process. The Department of Civil, Environmental and Geomatics Engineering (D-BAUG) at ETH Zürich provided funding via the E-Bike City lighthouse project (https://ebikecity.ch).

8 References

Antwi, R. B., M. Kimollo, S. Y. Takyi, E. E. Ozguven, T. Sando, R. Moses and M. A. Dulebenets (2024a) Turning Features Detection from Aerial Images: Model Development and Application on Florida's Public Roadways, *Smart Cities*, **7** (3) 1414–1440, ISSN 2624-6511. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Antwi, R. B., S. Takyi, A. Karaer, E. E. Ozguven, R. Moses, M. A. Dulebenets and T. Sando (2024b) Detecting school zones on florida's public roadways using aerial images and artificial intelligence (ai2), *Transportation Research Record*, **2678** (4) 622–636.

Bundesamt für Statistik and Bundesamt für Raumentwicklung (2023) *Mobilitätsverhalten der Bevölkerung: Ergebnisse des Mikrozensus Mobilität und Verkehr 2021*, Bundesamt für Statistik and Bundesamt für Raumentwicklung, Neuchâtel und Bern.

Campbell, A., A. Both and Q. C. Sun (2019) Detecting and mapping traffic signs from google street view images using deep learning and gis, *Computers, Environment and Urban Systems*, **77**, 101350, ISSN 0198-9715.

DiGioia, J., K. E. Watkins, Y. Xu, M. Rodgers and R. Guensler (2017) Safety impacts of bicycle infrastructure: A critical review, *Journal of Safety Research*, **61**, 105–119, ISSN 1879-1247.

Ferster, C., T. Nelson, K. Manaugh, J. Beairsto, K. Laberee and M. Winters (2023) Developing a national dataset of bicycle infrastructure for Canada using open data sources, *Environment and Planning B: Urban Analytics and City Science*, 239980832311599, ISSN 2399-8083, 2399-8091.

Habibi Aghdam, H., E. Jahani Heravi and D. Puig (2016) A practical approach for detection and classification of traffic signs using convolutional neural networks, *Robotics and Autonomous Systems*, **84**, 97–112, ISSN 0921-8890.

Hwang, U. and S. Guhathakurta (2023) Exploring the Impact of Bike Lanes on Transportation Mode Choice: A simulation-based, route-level impact analysis, *Sustainable Cities and Society*, **89**, 104318, ISSN 22106707.

Jin, H., Y. Feng and M. Li (2012) Towards an automatic system for road lane marking extraction in large-scale aerial images acquired over rural areas by hierarchical image

analysis and Gabor filter, *International Journal of Remote Sensing*, **33** (9) 2747–2769, ISSN 0143-1161, 1366-5901.

Pritchard, R., D. Bucher and Y. Frøyen (2019) Does new bicycle infrastructure result in new or rerouted bicyclists? A longitudinal GPS study in Oslo, *Journal of Transport Geography*, **77**, 113–125, ISSN 09666923.

Reynolds, C. C., M. A. Harris, K. Teschke, P. A. Cripton and M. Winters (2009) The impact of transportation infrastructure on bicycling injuries and crashes: a review of the literature, *Environmental Health*, **8** (1) 47, ISSN 1476-069X.

Smith, R. L. J. and T. Welsh (1988) Safety impacts of bicycle lanes, *TRANSPORTATION RESEARCH RECORD*, **1168**, 70–77.

Sohail, A., M. A. Cheema, M. E. Ali, A. N. Toosi and H. A. Rakha (2023) Data-driven approaches for road safety: A comprehensive systematic literature review, *Safety Science*, **158**, 105949, ISSN 09257535.

Song, Y., J. Preston and D. Ogilvie (2017) New walking and cycling infrastructure and modal shift in the uk: A quasi-experimental panel study, *Transportation Research Part A: Policy and Practice*, **95**, 320–333, ISSN 0965-8564.

Tait, C., R. Beecham, R. Lovelace and S. Barber (2022) Is cycling infrastructure in london safe and equitable? evidence from the cycling infrastructure database, *Journal of Transport Health*, **26**, 101369, ISSN 2214-1405.

Van Petegem, J. H., P. Schepers and G. J. Wijlhuizen (2021) The safety of physically separated cycle tracks compared to marked cycle lanes and mixed traffic conditions in Amsterdam, *European Journal of Transport and Infrastructure Research*, 19–37 Pages. Artwork Size: 19-37 Pages Publisher: European Journal of Transport and Infrastructure Research.

Vandenbulcke, G., I. Thomas and L. Int Panis (2014) Predicting cycling accident risk in Brussels: A spatial case–control approach, *Accident Analysis & Prevention*, **62**, 341–357, ISSN 00014575.

Vanparijs, J., L. Int Panis, R. Meeusen and B. de Geus (2015) Exposure measurement in bicycle safety analysis: A review of the literature, *Accident Analysis & Prevention*, **84**, 9–19, ISSN 00014575.

Wijnands, J. S., H. Zhao, K. A. Nice, J. Thompson, K. Scully, J. Guo and M. Stevenson (2021) Identifying safe intersection design through unsupervised feature extraction from satellite imagery, *Computer-Aided Civil and Infrastructure Engineering*, **36** (3) 346–361, ISSN 1093-9687, 1467-8667.

Wysling, L. and R. S. Purves (2022) Where to improve cycling infrastructure? assessing bicycle suitability and bikeability with open data in the city of paris, *Transportation Research Interdisciplinary Perspectives*, **15**, 100648, ISSN 2590-1982.

9 Glossary

ARE	Office for Spatial Development
IoU	Intersection over Union
OSM	OpenStreetMap
YOLOv8	You Only Look Once, version 8

A Input Data Overview

Table 3: Overview of agglomeration orthophoto availability and resolution (in centimeters)

	2024	2023	2022	2021	2020	2019	2018	2017	2016	2015	2014	2013	2012	2011	2010	2009	2008	2007	2006	2005	2004	2003	2002	2001	2000	1999	1998
Zurich			10			10			25			25			25				25				50				
Genf		10			10					25			25			25			25		50						50
Basel	10			10			10				25			25		25					50						50
Lausanne		10			10					25			25			25			25		50						50
Bern	10			10			10		25			25			25			25			50						50
Winterthur			10			10					25			25			25			25		50					
Luzern				10			10		25			25			25				25		50		50				
St. Gallen			10			10					25			25			25			25			50				
Lugano				10			10			25			25			25			25		50			50			
Biel	10			10			10				25			25			25				50						50

B Labels Instances

Table 4: Label counts by class and year

	1998	2002	2004	2013	2014	2015	2016	2018	2019	2020	2021	2022	2023
advanced_retaining_bar	0	0	28	5	24	21	54	39	21	13	31	14	7
bike_lane	192	163	201	118	278	225	326	174	177	168	325	186	109
bike_pictogram	25	1	84	64	94	252	220	192	123	241	279	54	126
bike_stopping_space	0	0	13	3	4	22	21	7	7	32	24	5	13
circle	0	8	11	14	2	12	20	6	3	6	6	2	5
crosswalk	1358	168	810	273	305	316	475	313	257	243	338	120	87
pedestrian_island	271	51	173	55	62	68	85	49	43	33	77	38	12
red_colored_bike_lane	9	0	14	2	8	27	99	22	13	18	74	11	22
retaining_bar	560	80	343	74	105	180	275	144	133	182	240	67	75
shark_teeth	17	62	302	90	120	97	172	126	90	38	135	63	48

Table 5: Label counts by class and agglomeration

	Basel	Bern	Biel	Genf	Lausanne	Lugano	Luzern	St. Gallen	Winterthur	Zurich
advanced_retaining_bar	25	56	37	22	17	8	30	15	13	34
bike_lane	298	379	208	254	349	88	217	242	201	406
bike_pictogram	308	242	91	376	261	47	95	49	60	226
bike_stopping_space	6	42	6	12	54	2	11	6	6	6
circle	5	27	7	5	13	5	2	3	2	26
crosswalk	774	838	566	535	999	172	145	179	131	724
pedestrian_island	137	163	96	112	203	22	54	51	38	141
red_colored_bike_lane	36	35	1	52	22	5	130	10	14	14
retaining_bar	322	408	213	383	465	102	112	81	64	308
shark_teeth	267	350	94	26	114	91	61	68	75	214

C Prediction Examples

Figure 3: Raw prediction output: Intersection Lagerstrasse and Langstrasse, Zürich

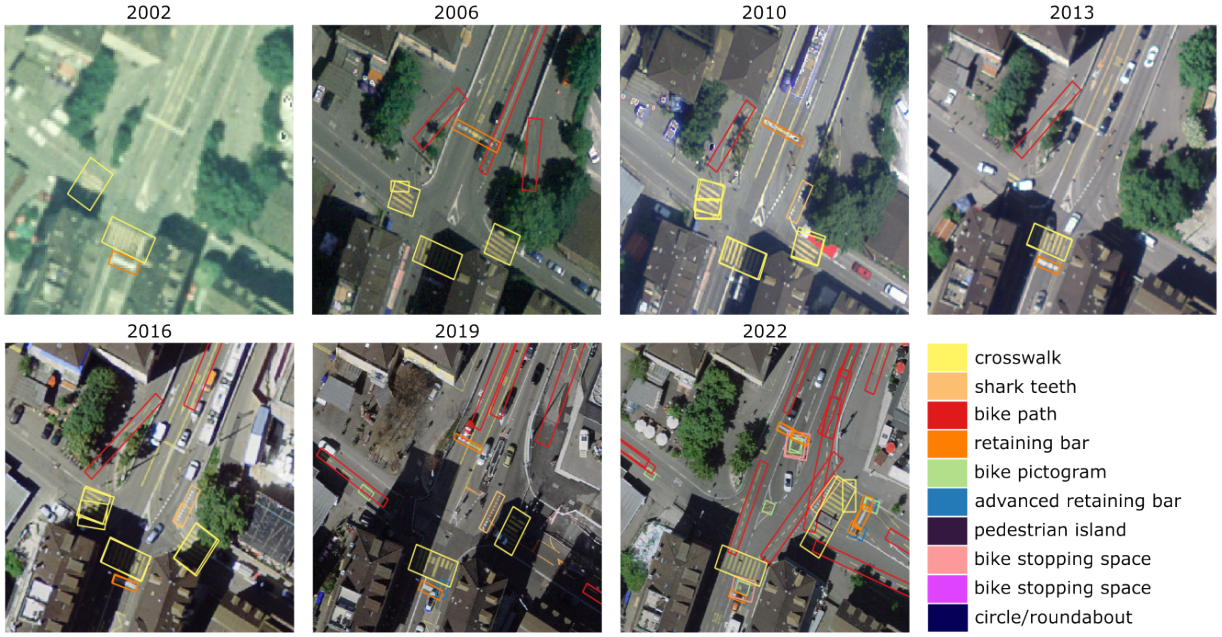

Figure 4: Raw prediction output: Intersection Museumstrasse, Bahnhofquai and Walchebrücke, Zürich

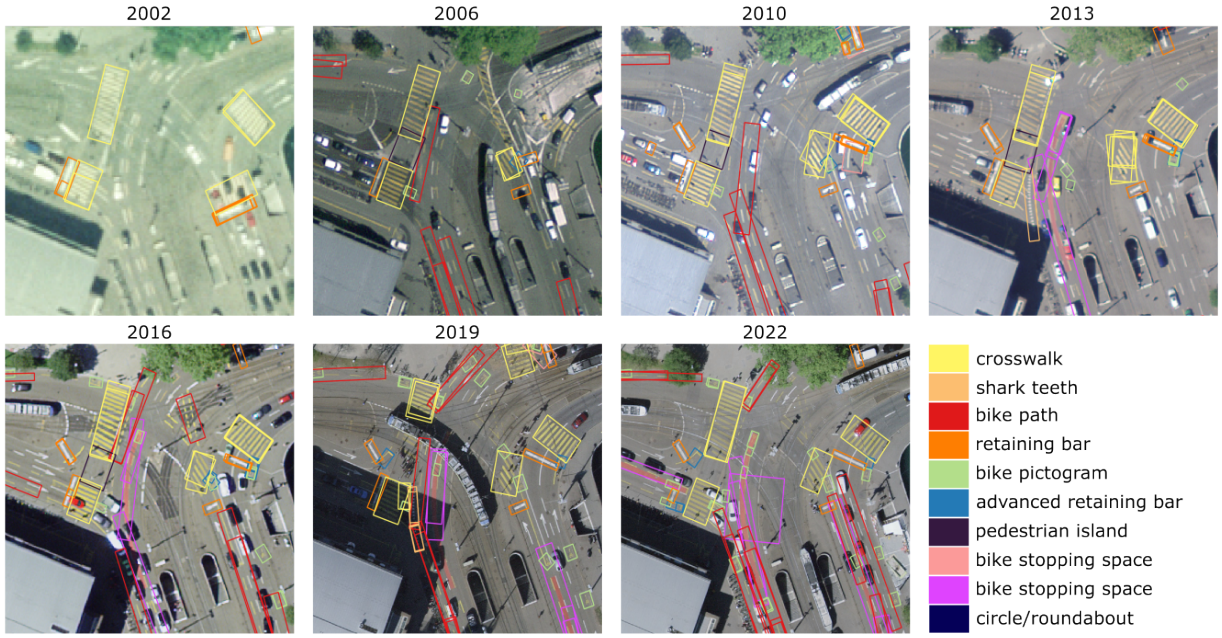

Figure 5: Raw prediction output: Bucheggplatz, Zürich

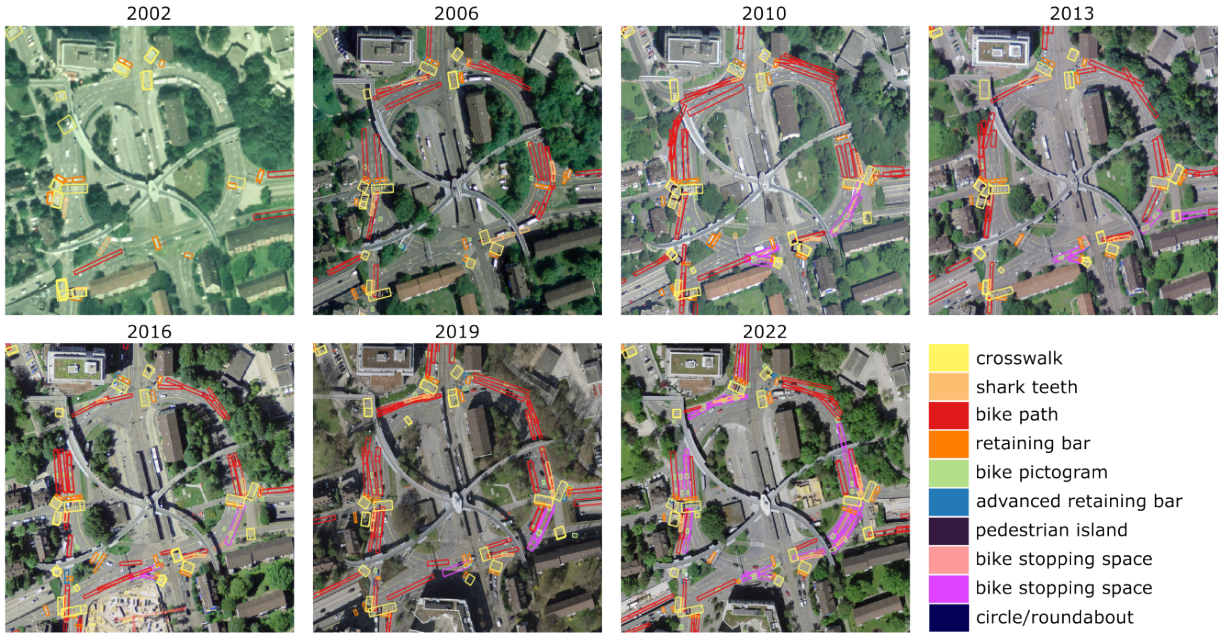

Figure 6: Raw prediction output: Intersection Hofwiesenstrasse and Wehntalerstrasse, Zürich

D Raw Prediction Outputs

Figure 7 shows the predicted instances of each type of infrastructure for each agglomeration over time.

Figure 7: Predicted infrastructure instances per agglomeration, normalized by road length. Dashed vertical lines show the year in which the aerial image resolution was upgraded (from 50 to 25 to 10 centimeters)

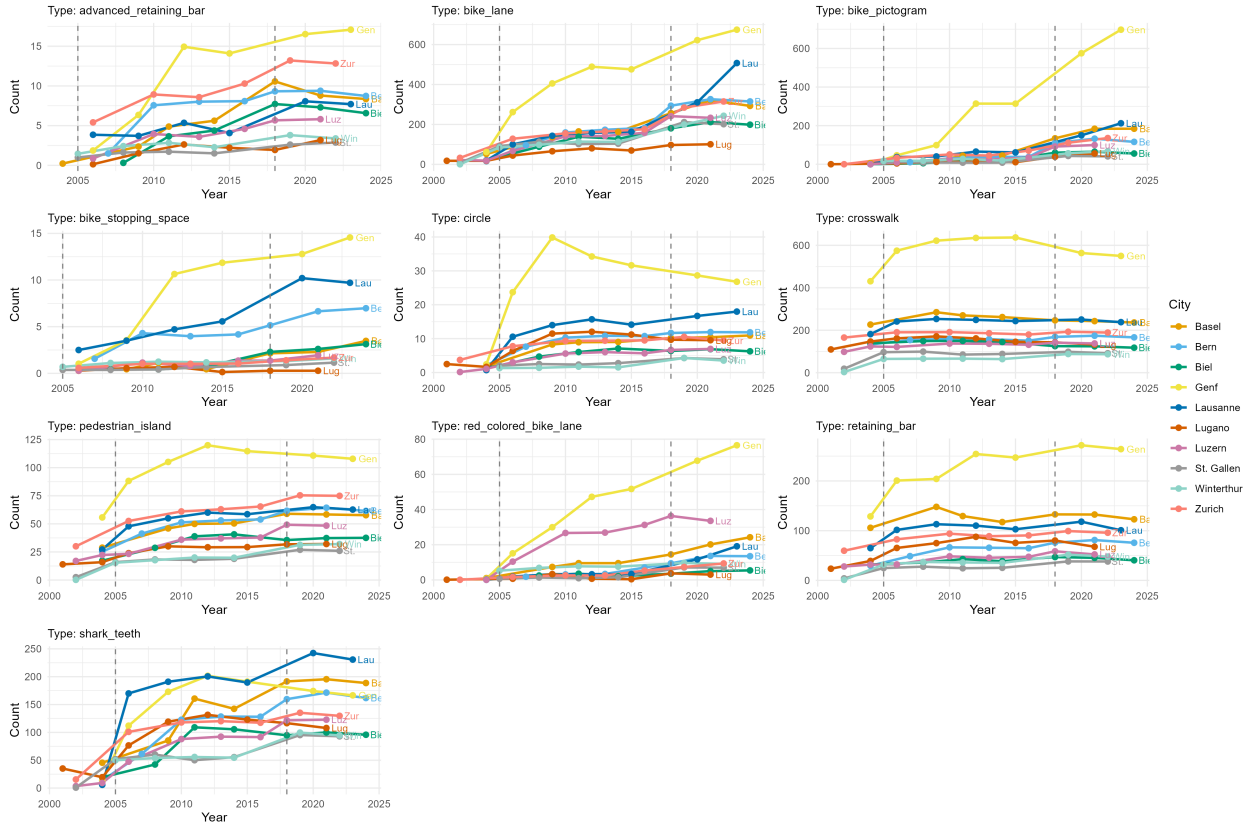

E Further Analysis at the Intersection Level

Figure 8 shows the share of major intersections without bike pictograms for each agglomeration by settlement type (core urban area, urban area, periurban area). Bike pictograms was chosen as the proxy for cycling infrastructure presence since pictograms are relatively ubiquitous in all agglomerations. In general, more urban areas have fewer intersections with no cycling infrastructure, with a decreasing trend over time for all agglomerations. Geneva has the fewest intersections without pictograms in all settlement areas. Luzern performs relatively better in periurban areas than the other agglomerations.

Figure 8: Share of major intersections with no bike pictograms by settlement type

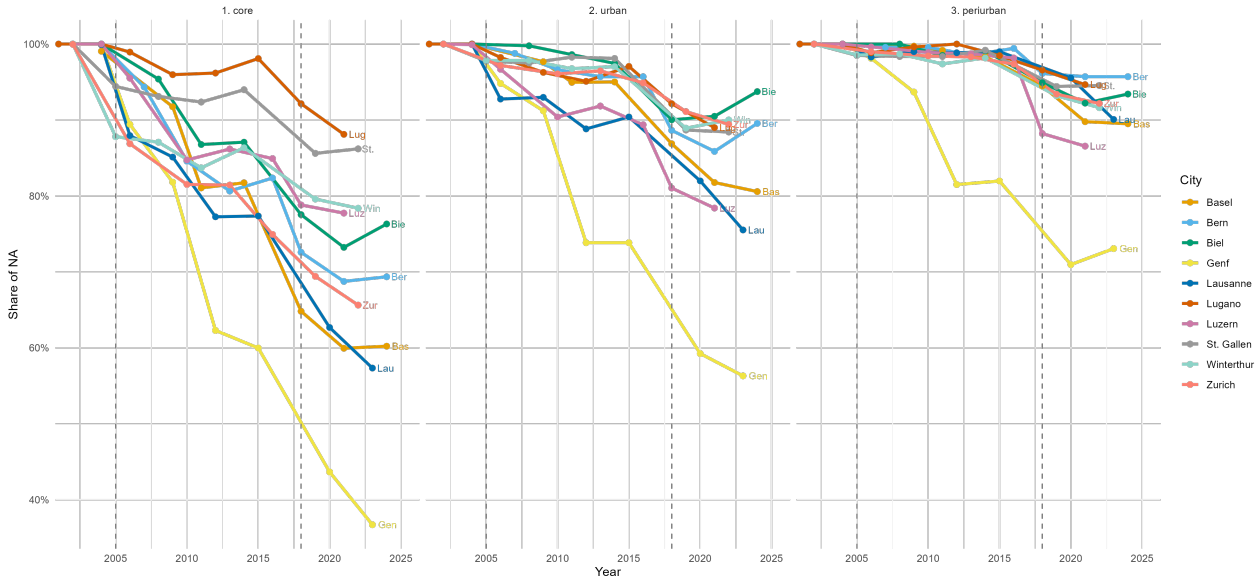

Figure 9 shows the share of major intersections without bike pictograms for each agglomeration by road ownership (communal, cantonal, federal). In general, federal roads have fewer intersections with no cycling infrastructure, with a decreasing trend over time for all roads. As before, Geneva has the fewest intersections without pictograms. The range between all the agglomerations is similar for all roads (about 40 %), but Geneva's lead is less pronounced for federal roads.

Figure 9: Share of major intersections with no bike pictograms by road ownership