Comparative Assessment of GeoAI-Based Frameworks for Automatic Urban Tree Cover Mapping

Francesco Pirotti[1,2]*, Linda Avesani[3], Matteo Pianetti[4], Riccardo Greco[3], Catherine Dezio[2], Davide Quaglia[4]

[1] Interdepartmental Research Center of Geomatics (CIRGEO), University of Padova, Legnaro, Italy
[2] Department of Land, Environment and Agro-Forestry (TESAF), University of Padova, Legnaro, Italy
(francesco.pirotti,catherine.dezio)@unipd.it
[3] Department of Biotechnology, University of Verona, Verona, Italy – (riccardo.greco,linda.avesani)@univr.it
[4] Department of Informatics, University of Verona, Verona, Italy
matteo.pianetti@studenti.univr.it, davide.quaglia@univr.it

Keywords: urban trees, artificial intelligence, YOLO, Mask R-CNN, Deepforest, VHRTrees

Abstract

Accurate mapping of urban tree canopy is essential for quantifying ecosystem services and assessing the impact of green infrastructure on wellbeing and public health. This study evaluates and compares three Geospatial Artificial Intelligence (GeoAI) frameworks for the automated detection and segmentation of tree cover. The frameworks are YOLO, Detectree, and TreeEyed Utilizing high-resolution aerial imagery (0.2 m and 0.5 m ground sampling distance), the research tests different deep-learning paradigms, including object detection and semantic segmentation. The results indicate that while object-based models like YOLO align closely with statistical baselines (30.83% vs 30.11%), pixel-based models such as Detectree may underestimate fragmented urban vegetation. The study highlights the effectiveness of the TreeEyed QGIS plugin for urban applications and emphasizes the necessity of local LiDAR-derived data for model validation. Further studies would benefit from ad-hoc training with correct co-registration and consistent coordinate reference systems across layers.

1. Introduction

Urban tree cover plays a fundamental role in enhancing environmental quality, regulating urban microclimates, and supporting biodiversity within increasingly dense and anthropized landscapes. Trees contribute to the mitigation of urban heat islands, improve air quality by filtering pollutants, and provide ecosystem services that directly influence human health and well-being. In this context, the accurate mapping and quantification of urban vegetation are essential components for sustainable urban planning, climate adaptation strategies, and ecosystem service assessment.

Traditionally, estimates of urban tree cover have relied on field surveys or point-based sampling approaches, which, while robust, are often time-consuming, labor-intensive, and limited in spatial coverage. The rapid development of remote sensing technologies, combined with advances in machine learning and computer vision, has opened new possibilities for large-scale, high-resolution assessment of vegetation (Pirotti et al., 2025). Deep learning models applied to aerial and satellite imagery enable automated detection and classification of urban features with increasing accuracy, offering a promising alternative to conventional methods which require aerial or ground 3D data, such as LiDAR surveys (Pirotti et al., 2013).

Within this evolving methodological context, the present study focuses on the city of Verona, where urban green infrastructure is of growing importance for both environmental sustainability and citizens’ quality of life. The research is conducted in the framework of the èVRgreen project, funded by Fondazione Cariverona (EVRgreen, 2026), which aims to integrate environmental monitoring and nature-based solutions to improve urban resilience. Among the project’s objectives is the investigation of relationships between urban greenery and human well-being, including the hypothesis that the number of trees visible from residential environments is negatively correlated with perceived stress levels (De Lucia et al., 2025).

To support such analyses, reliable spatial information on the distribution and extent of tree canopy is required. In this study, urban tree cover was assessed using a combination of approaches, including the point-based estimation tool i-Tree Canopy and state-of-the-art deep learning models, specifically YOLO, Detectree and other models form the QGIS plugin TreeEyed. By comparing these methodologies, the study aims to evaluate their respective performance, identify potential biases, and explore their applicability for urban forestry and ecosystem service assessment. This comparative framework also provides insights into how emerging artificial intelligence techniques (Gazzea et al. 2023) can complement or enhance established tools in the analysis of complex urban environments. Furter detection of tree species is also a topic of interest (Michelini et al. 2022) for further enrichment of the products. A detection could foster improved species identification by delimiting only canopies for further classification.

2. Materials and methods

2.1 Study area

The study area is the city of Verona, located in north-eastern Italy along the Adige River. Verona is characterized by a heterogeneous urban landscape that includes a dense historical centre, residential and industrial districts, and peri-urban areas transitioning into agricultural land and hilly terrain on the north (see figure 1). The northern sector is marked by wooded hills and higher vegetation density, while the southern portion is predominantly flat and intensively cultivated. This variability in land cover and topography makes Verona a suitable case study for testing and comparing different approaches to urban tree cover estimation, as it encompasses a wide range of vegetation structures and urban forms within a relatively compact area.

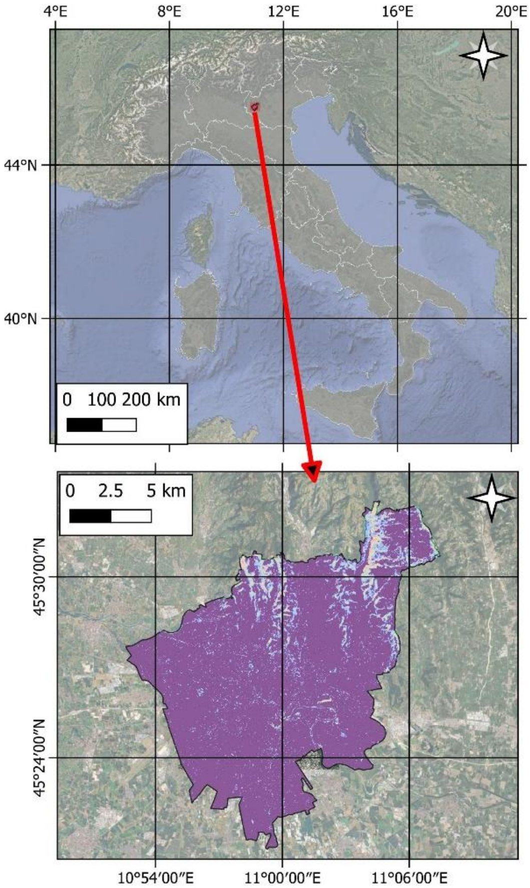

Figure 1: municipality of Verona, with canopy height model showing higher trees in the northeast part.

2.2 Deep learning

The methodological approach was based on the application of deep learning techniques for tree detection and canopy mapping. The YOLO version 8 and version 11 (You Only Look Once) (Redmon et al. 2016) object detection framework and the Detectree (Bosch 2024) framework were selected for an initial training strategy to represent complementary paradigms in image analysis: object-based detection and pixel-based segmentation. A third framework was used for the second strategy, where models are already provided and the QGIS interface used together with a high-resolution orthoimage.

The two different strategies are used to explore two different approaches, the first one, with YOLO and Detectree is more technical and requires significant manual training and interpretation, the second one is more "user friendly" and is integrated as QGIS plugin, thus integrating a tool that is already commonly used in spatial data analysis. This would make this last approach more easily accessible for new users which have a lower technical capacity.

2.2.1 YOLO: YOLO was employed to identify and delineate individual tree crowns from high-resolution aerial imagery through supervised training on annotated datasets, offering advantages in detecting discrete objects and estimating tree counts (Moussaoui et al., 2024). For YOLO, the training process involved configuring key parameters such as the number of epochs, image size, and batch size, as well as defining class labels through a structured dataset.

2.2.2 Detectree: Detectree was implemented as a semantic segmentation approach, where convolutional neural networks classify each pixel as either tree canopy or non-canopy, producing continuous canopy cover maps that are more directly comparable with area-based metrics. Detectree training required paired inputs of aerial images and corresponding binary masks representing tree cover. Both models were trained on datasets specifically prepared to capture the heterogeneity of the study area, including urban, agricultural, and forested contexts.

2.2.3 TreeEyed: A third framework was tested, the TreeEyed QGIS plugin (Ruiz-Hurtado 2025). TreeEyed is an open-source computational framework designed for the extraction and analysis of tree canopy information from high-resolution remote sensing imagery using deep learning techniques. The platform supports both object detection and semantic segmentation workflows, integrating convolutional neural network architectures for the identification of individual tree crowns and the delineation of continuous canopy cover. Its design enables scalable processing, standardized data handling, and reproducible model training, providing a flexible environment for urban vegetation mapping and quantitative analysis of tree-related ecosystem variables. Available models in this framework are listed in table 1 below.

Model	Source	Ideal resolution	Computer Vision Task
HighResCanopyHeight	https://github.com/facebookresearch/HighResCanopyHeight	1 m	Pixel-wise Regression
Mask R-CNN	Custom trained	5 m	Instance Segmentation
Deepforest	https://github.com/weecology/DeepForest	0.5 m	Object Detection
VHRTrees	https://github.com/RSandAI/VHRTrees	0.5 m	Object Detection

Table 1. List of models implemented in TreeEyed (afruizh, 2026).

The combined use of these methods allows not only for cross-validation of results but also for the evaluation of their respective strengths and limitations, providing a more robust framework for urban tree cover estimation and supporting informed methodological choices in urban forestry applications.

2.3 i-Tree

Prior to implementing this methodology, the city's tree cover was also estimated using i-Tree Canopy, an online tool developed by the USDA Forest Service that applies a random point sampling approach. The method consists of generating randomly distributed points within a user-defined area and manually classifying each point according to the observed land cover type from aerial imagery. This approach is fundamentally a statistical sampling procedure based on photo-interpretation, where the proportion of points assigned to each class is used to estimate overall land cover percentages and associated uncertainty. Its simplicity and low data requirements have led to widespread adoption in urban forestry studies, although the accuracy of the results depends on the number of sampled points, the interpreter's ability to correctly classify each observation and the stratification of the samples which has to catch the variability of the context (Nowak et al., 2018).

2.4 Training

The training phase followed the two different strategies. Strategy one used orthoimagery from Google base maps downloaded as tiles at approximately 0.2 m ground sampling distance (GSD). The other strategy used the 2021 Orthoimagery provided by the Veneto Region (Regione del Veneto, 2026) and clipped over Verona at 0.5 m GSD.

2.4.1 Training strategy 1 consisted in manual visual interpretation over Verona canopies which were converted to bit-masks. For YOLO, the training workflow involved initializing the pre-trained model and configuring key parameters, including the number of epochs, input image size, and batch size. A configuration file (.yaml) was used to define the dataset structure, specifying the paths to training and validation images as well as the associated class labels. In the case of Detectree, the training dataset consisted of aerial imagery paired with corresponding binary masks, where tree canopy was represented by white pixels and all other classes by black pixels. Model training was performed by providing the paths to the images and masks along with a structured train/test split defined in a dataframe. This setup enabled supervised learning for pixel-wise classification of canopy cover. Both training procedures were specifically tailored to account for the heterogeneity of the study area. To ensure adequate representation of different landscape conditions, three separate datasets were created: one for the dense urban fabric of the city centre, one for the southern sector dominated by agricultural land (primarily ploughed fields), and one for the northern area characterized by hilly terrain and woodland cover. This stratification aimed to improve model generalization across contrasting land-use and vegetation patterns.

2.4.2 Training strategy 2 used an existing high resolution canopy height model (CHM) map at 0.5 m resolution from the Veneto Region over an area from the Alpine space (Regione del Veneto - Agricoltura, foreste e biodiversità, 2026). This was obtained from an aerial laser scanning survey in 2019 and allows to determine a high number of potential training and testing canopies for the model. The four models from TreeEyed are already trained, and, as seen in table 1, each has its own ideal resolution. For this we did not have to train the model, we assessed the results by providing a orthoimage resampled at the ideal resolution for each model.

2.5 Accuracy metrics

Rigorous accuracy would need ground-truth data to test for recall and precision. Recall is how many correct trees detected (TP – true positives) divided the sum of TP and misses (trees not detected, or false negatives FN). Precision instead is the ration of TP over the sum of TP and misclassifications, i.e. false positives (FP), which are objects labelled as trees which are not trees. Recall and precision can be combined with a single metric, the F1-score, which is the harmonic mean of recall and precision. This is a more robust indicator then normal average, as it will be penalized by low scores in either of the two metrics. If one of the two metrics, precision or recall, is equal to zero, the F1 score will be zero.

2.5.1 Strategy 1: For strategy 1 validation was more challenging due to the absence of independent ground truth data such as LiDAR-derived canopy information. As a result, an alternative validation dataset consisting of manually generated masks was used to compare predicted and reference classifications.

2.5.2 Strategy 2: Strategy 2 was tested in an area which had a canopy height model available which was used as ground truth. Each detected tree was considered a TP if its bounding box was covered with 70% or more with a canopy pixel from the CHM. The area itself was divided in three subareas to check for consistency.

3. Results

3.1 Strategy 1

The classes considered in this study were: tree, grass, road, building, water, soil, and other surfaces. The estimated total area covered by the tree class resulted as follows (see also table 1): 30% with i-Tree Canopy – 31% with YOLO – 19% with Detectree. i-Tree Canopy and YOLO are expected to overestimate tree cover by approximately 5%, whereas Detectree tends to underestimate it by about 6%. Efforts are currently focused on improving these estimates through the creation of new datasets and the retraining of both YOLO and Detectree. An alternative approach is also being explored using, U-Net, via the Python library segmentation-models-pytorch. This approach remains under experimentation, and results are not yet available to determine whether repeating the process will yield improved performance. Further updates will be reported in due course. Table 2 and Figure 2 below shows the difference between resulting percentage of urban tree cover for Verona. Figure 3 and Figure 4 also show details of results.

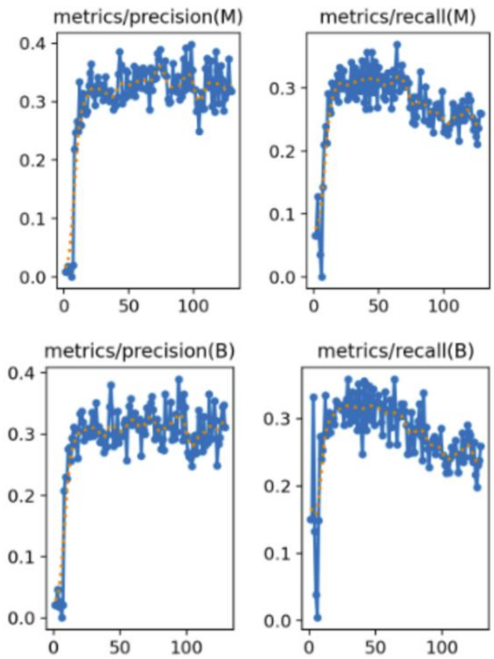

Figure 2. Accuracy of YOLO results; top the precision and recall for masks: bottom the precision and recall for bounding boxes.

Area Type	YOLO	Detectree	iTree Canopy
Hills	58.81%	52.40%	59.87%
City	27.93%	13.54%	26.35%
Periphery	10.33%	0.85%	10.32%
Total Tree Area	**30.83%**	**19.10%**	**30.11%**

Table 2. First result of percentage of tree cover percentage from the three different approaches in the first training strategy.

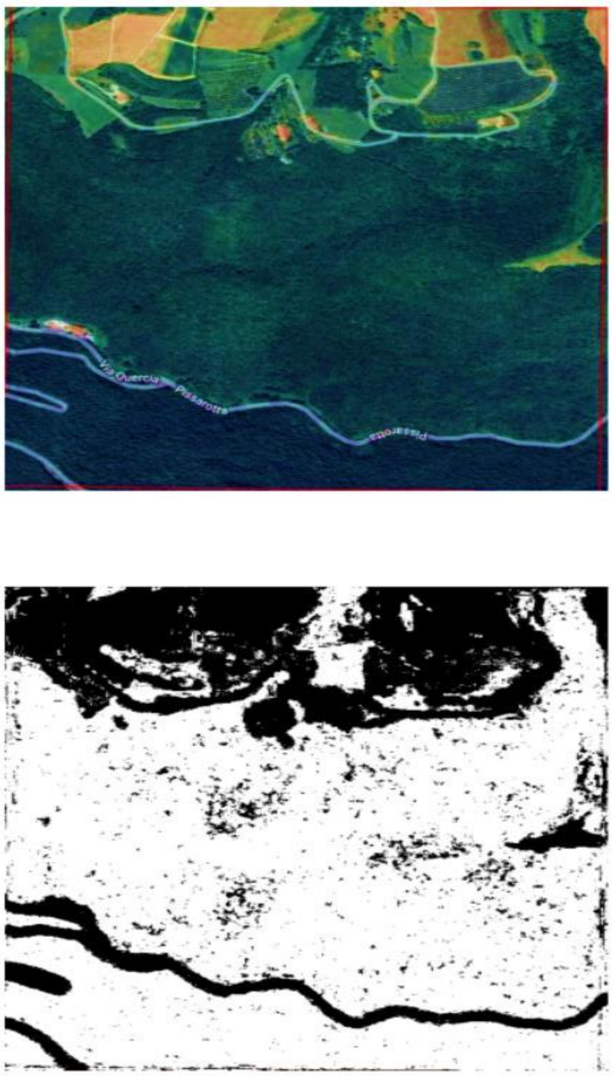

Figure 3. Results for Detectree on image (left) and output mask (right).

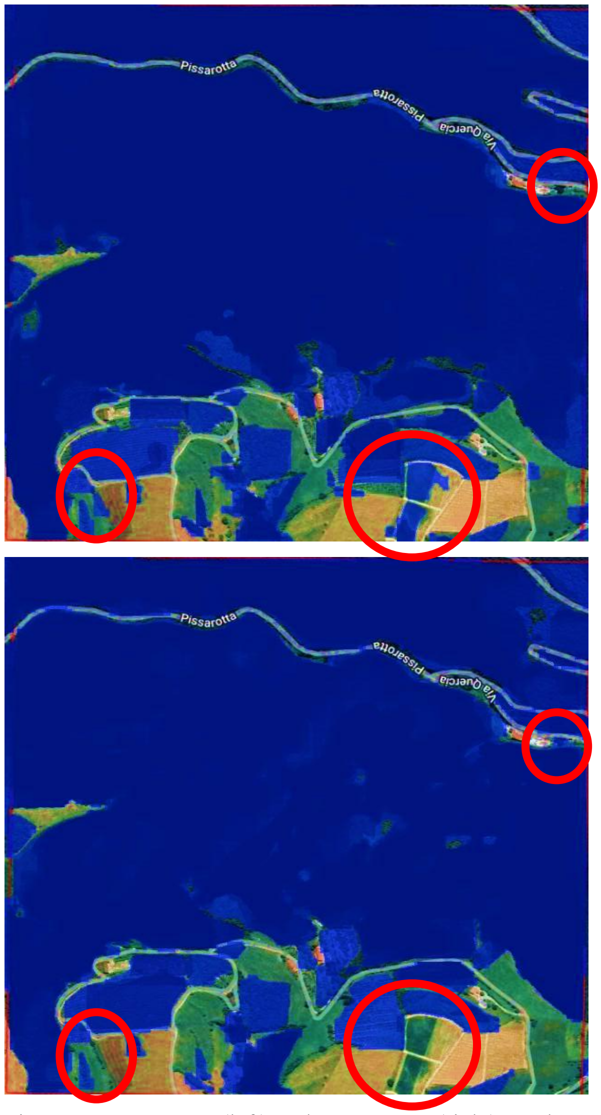

Figure 4. YOLO v8 (left) and YOLO v11 (right) results, with highlighted areas with differences.

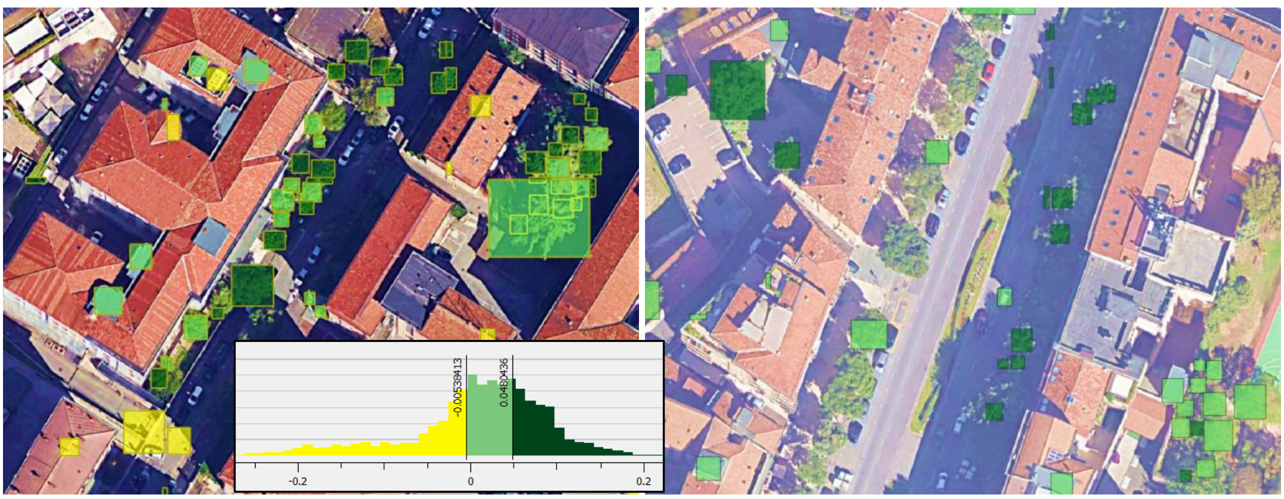

Figure 5. Detected trees from DeepForest algorithm in TreeEyed; on left with colors form orange to green related to the normalized greenness index (histogram showed at bottom right of left image) and on the right only canopies filtered out.

Figure 6. Detected trees from algorithm HighResCanopyHeight in TreeEyed; areas represented are the same as in figure 5 above.

3.2 Strategy 2

For a better understanding of accuracy, a small area with known canopy information was used (Figure 7). This area was divided into three sub-areas to check for possible differences in accuracy and to see if substantial differences were found.

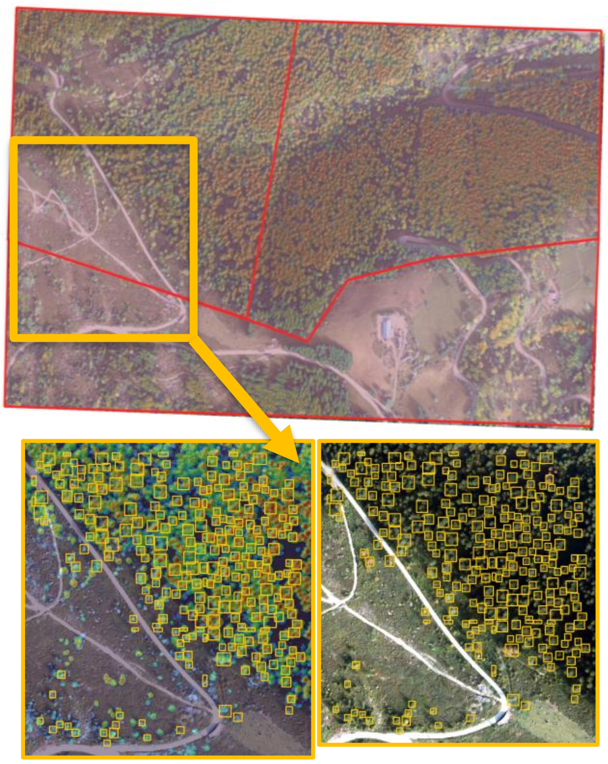

Figure 7. Three areas over the second area investigated and a zoom in on results.

The implementation of the TreeEyed plugin allowed for a comparative visualization of multiple deep learning architectures across the Verona study area. As illustrated in Figures 5 and 6, the output format varies significantly depending on the specific model architecture selected; these outputs typically include bounding boxes for object detection or raster masks for semantic segmentation of the canopy. For instance, the raster masks can be converted into vector polygons for further spatial analysis, as demonstrated in Figure 5. While the HighResCanopyHeight model possesses the unique capability to generate estimated canopy height maps, this specific vertical data product was excluded from the current analysis in favor of horizontal extent metrics. Due to the absence of independent ground-truth LiDAR data for the Verona city center, the initial performance of these models was evaluated through visual inspection. This qualitative assessment identified two primary categories of false positives that impacted the raw results. In some instances, the branches of a single large tree were erroneously detected as multiple individual trees, a phenomenon particularly visible in the upper-right quadrant of Figure 5. Non-vegetative urban features with similar geometric or spectral profiles—specifically building chimneys and parked cars—were occasionally misidentified as tree canopies. To mitigate these errors, a streamlined post-processing workflow was implemented to improve thematic accuracy. This involved the application of a normalized difference greenness index, calculated by subtracting the red channel from the green channel and dividing by their sum. By using the bounding polygons to extract average color statistics, the framework could effectively filter out "spectral imposters". Any bounding box or polygon that failed to meet a specific "greenness" threshold was removed from the final dataset, ensuring that only true vegetative features were included in the total canopy count

Table 3 shows the recall and precision of the different models available in TreeEyed.

	Area 1	**Area 2**	**Area 3**
HighRes CanopyHeight	0.81 | 0.69	0.61 | 0.34	0.90 | 0.34
Mask R-CNN	0.52 | 0.61	0.52 | 0.71	0.18 | 0.22
Deepforest	0.82 | 0.52	0.77 | 0.79	0.62 | 0.73
VHRTrees	0.24 | 0.65	0.85 | 0.42	0.15 | 0.56

Table 3. Accuracy metrics, recall | precision from the four models provided in TreeEyed.

4. Discussion

4.1 Stratification and model paradigms

The training strategy was designed to address the significant spatial heterogeneity of the Verona study area. By stratifying the dataset into three representative zones the framework aimed to understand how the models compare across different land-cover scenarios. The implementation of two distinct deep learning paradigms revealed complementary strengths. YOLO, functioning as an object detection and instance segmentation model, proved superior in identifying and delineating discrete tree crowns. In contrast, Detectree, operating as a pixel-based classifier, was more effective at capturing continuous canopy patches. However, both architectures demonstrated a high sensitivity to the representativeness of the training data. This emphasizes that in complex urban environments, the architectural choice (object-based vs. pixel-based) must be balanced by a rigorous dataset design to prevent systematic biases in canopy estimation.

4.2 Accuracy and methodological differences

The high degree of alignment between the i-Tree statistical baseline and the YOLO results suggests a consensus between point-based sampling and object detection. However, both may slightly overestimate total area due to the merging of overlapping crowns or the inclusion of small gaps within the bounding geometries. Conversely, Detectree's consistent underestimation (19%) highlights a known limitation of pixel-based segmentation in fragmented landscapes; the model often fails to classify isolated street trees or sparse vegetation where the spectral signature is heavily influenced by surrounding gray infrastructure. These discrepancies underscore the "uncertainty of the ground truth" in urban remote sensing, where the lack of standardized reference data complicates the validation of competing GeoAI frameworks.

4.3 The "Pseudo-Ground Truth" Paradox and Spatial Accuracy

The availability of global high-resolution products, such as the 1 m canopy height map by Tolan et al., (2024), presents a compelling opportunity for automated training. However, this study identifies spatial co-registration as a critical constraint.

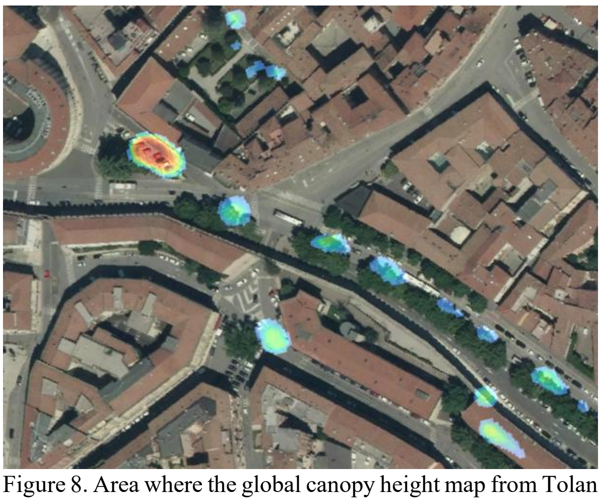

Figure 8. Area where the global canopy height map from Tolan et al., 2024 shows incorrect co-registration with the orthoimage.

As illustrated in Figure 8, global datasets often exhibit geometric shifts when overlaid with local high-resolution orthoimagery. Training a model on misaligned data introduces noise that degrades the precision of the resulting segmentation masks, when not completely mis-training the model. To avoid such cases, this research tested such approach by using the Veneto Region's LiDAR-derived CHM, which offers superior reliability in co-registration. This creates a methodological paradox: the models are being trained to predict a CHM using an existing CHM as ground truth. While this limits immediate replicability to regions with existing LiDAR data (Pirotti et al., 2017), it serves as a vital proof-of-concept for "training with predictions." By using established model outputs as pseudo-ground truth, the Detectree model can potentially be retrained in similar urban areas which also have existing CHMs to leverage the strengths of the YOLO architecture without requiring exhaustive manual annotation. This iterative cross-training approach offers a viable path toward reducing model bias and improving consensus in tree cover estimation.

4.4 Socio-environmental and health implications

The technical accuracy of these mapping frameworks has direct implications for urban policy and public health. Recent epidemiological research, such as the work by Iungman et al., (2023) and Giannico et al., (2024), have quantified the effects of urban trees and their ability to reduce mortality. By providing a high-resolution map of tree visibility and density, this study potentially enables a more granular analysis of how green infrastructure mitigates the Urban Heat Island (UHI) effect in Verona. Preliminary findings suggest that even marginal increases in visible canopy cover can correlate with lower perceived stress levels among residents (Lefosse et al.,, 2025), reinforcing the need for GeoAI tools that can accurately monitor municipal greening targets.

5. Conclusions

The comparative assessment of GeoAI frameworks in Verona demonstrates that deep learning models offer a robust and scalable alternative to traditional, labor-intensive urban tree mapping methods. The study highlights a clear functional divergence between model paradigms: object-based frameworks like YOLO excel at identifying discrete tree crowns and align closely with statistical baselines, while pixel-based models like Detectree are better suited for continuous canopy patches but may significantly underestimate fragmented urban vegetation. Furthermore, the integration of these models into accessible platforms like the TreeEyed QGIS plugin provides a "user-friendly" pathway for municipal planners and researchers to monitor green infrastructure without requiring extensive technical expertise. Ultimately, the accuracy of these mapping tools has profound implications for automatic mapping of urban assets that affect human well-being, such as trees. While the "pseudo-ground truth" paradox—using existing LiDAR-derived data (e.g. Alberti et al., 2013) to train models—presents some limitations for immediate global replicability, it establishes a vital proof-of-concept for iterative model improvement and bias reduction. By providing high-resolution data on tree density and visibility, these GeoAI frameworks enable a more granular understanding of how urban forests mitigate heat islands and reduce resident stress, facilitating more effective, data-driven sustainable urban planning

Acknowledgements

This work was carried out with the support of the èVRgreen project, funded by Fondazione Cariverona, and the European Union—Next Generation EU through the project of national interest (PRIN) “Geo-Intelligence for improved air quality monitoring and analysis (GeoAIr)” 202258ACSL.

References

Afruizh, 2026. https://github.com/afruizh/TreeEyed Accessed on 10 January 2026

Alberti, G., Boscutti, F., Pirotti, F., Bertacco, C., de Simon, G., Sigura, M., Cazorzi, F., Bonfanti, P., 2013. A LiDAR-based approach for a multi-purpose characterization of alpine forests: An italian case study. IForest 6, 156–168. https://doi.org/10.3832/ifor0876-006

Bosch, M 2020. “DetecTree: Tree detection from aerial imagery in Python”. *Journal of Open Source Software*, (550), 2172. doi.org/10.21105/joss.02172

CHM Sette Comuni - Regione del Veneto. https://www.regione.veneto.it/web/agricoltura-e-foreste/chm-sette-comuni. Accessed 25 Mar 2026

De Lucia, A., Donisi, V., Loi, M., Pasini, M., Pogliaghi, S., & Perlini, C. 2025. Urban Green Spaces and Psychological Well-being: Review of the Literature and Preliminary Data from the eVRgreen Study. National Congress of the Research Group in Psychosomatics (RGP) Verona, June 19 and 20. *La Clinica Terapeutica*, *176*(3).

Gazzea, M., Kristensen, LM., et al. 2022. Tree Species Classification Using High-Resolution Satellite Imagery and Weakly Supervised Learning. *IEEE Trans Geosci Remote Sensing* 60:1–11. https://doi.org/10.1109/TGRS.2022.3210275

EVRgreen, 2026. Verso una Verona più Verde e Sostenibile https://evrgreen.it accessed on 10 February 2026

Giannico, O.V., Sardone, R., Bisceglia, L., Addabbo, F., Pirotti, F., Minerba, S., Mincuzzi, A., 2024. The mortality impacts of greening Italy. *Nature Communications*. https://doi.org/10.1038/s41467-024-54388-7

Iungman, T., Cirach, M., Marando, F. et al., 2023. Cooling cities through urban green infrastructure: a health impact assessment of European cities. *The Lancet*, 2023; 401, 577-589

Lefosse, D.C., Duarte, F., Sanatani, R.P., Kang, Y., van Timmeren, A., Ratti, C., 2025. Feeling Nature: Measuring perceptions of biophilia across global biomes using visual AI. npj *Urban Sustain* 5, 4. https://doi.org/10.1038/s42949-025-00192-1

Moussaoui, H., Akkad, N.E., Benslimane, M. et al., 2024. Enhancing automated vehicle identification by integrating YOLO v8 and OCR techniques for high-precision license plate detection and recognition. *Sci Rep* 14, 14389 https://doi.org/10.1038/s41598-024-65272-1

Nowak, D.J., Maco, S., & Binkley, M. 2018. i-Tree: Global tools to assess tree benefits and risks to improve forest management. *Arboricultural Consultant*. 51 4): 10-13., 514), 10-13.

Pirotti, F, Carmelo, A, Kutchartt, E 2025. Fuel detection in forest environments training deep learners with smartphone imagery. *ISPRS Ann Photogramm Remote Sens Spatial Inf Sci* X-G-2025:649–656. https://doi.org/10.5194/isprs-annals-X-G-2025-649-2025

Pirotti, F., Kobal, M., Roussel, J.R., 2017. A Comparison of Tree Segmentation Methods Using Very High Density Airborne Laser Scanner Data. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W7, 285–290. https://doi.org/10.5194/isprs-archives-XLII-2-W7-285-2017

Pirotti, F., Guarnieri, A., Vettore, A., 2013. Vegetation filtering of waveform terrestrial laser scanner data for DTM production. Applied Geomatics 5, 311–322. https://doi.org/10.1007/s12518-013-0119-3

Michelini, D., Dalponte, M., Carriero, A., et al., 2022. Hyperspectral and LiDAR Data for the Prediction via Machine Learning of Tree Species, Volume and Biomass: A Contribution for Updating Forest Management Plans, in: Borgogno-Mondino, E., Zamperlin, P. (Eds.), Geomatics for Green and Digital Transition, Communications in Computer and Information Science. Springer International Publishing, Cham, pp. 235–250. https://doi.org/10.1007/978-3-031-17439-1_17

Redmon, J., Divvala S., Girshick R., Farhadi A. 2016. You Only Look Once: Unified, Real-Time Object Detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR)

Regione del Veneto – Direzione Pianificazione Territoriale. 2026) Servizio WMTS (Web Map Tile Service) dell’Infrastruttura Dati Territoriali (IDT-RV 2.0). accessed on 10 February 2026 Available at: https://idt2.regione.veneto.it/gwc/service/wmts

Regione del Veneto, 2026. Modelli digitali delle chiome https://www.regione.veneto.it/web/agricoltura-e-foreste/modelli-digitali-delle-chiome. Accessed 28 Mar 2026

Ruiz-Hurtado, A.F., Bolaños, J.P., Arrechea-Castillo, D.A., Cardoso, JA 2025. TreeEyed: A QGIS Plugin for Tree Monitoring in Silvopastoral Systems Using State of the Art AI Models. SoftwareX 29:102071. https://doi.org/10.1016/j.softx.2025.102071

Tolan, J., Yang, H.-I., et al., 2024. Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. Remote Sensing of Environment 300, 113888. https://doi.org/10.1016/j.rse.2023

Yang, L., Wu, X., Praun, E., Ma, X., 2009. Tree detection from aerial imagery, in: *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS ’09*. Association for Computing Machinery, New York, NY, USA, pp. 131–137. https://doi.org/10.1145/1653771.1653792