

Multi-model Approach for Tree Detection and Classification in Wallonia Region (Belgium)

Nicola Dimarco¹, Benoît Bartiaux¹, Louis Andreani¹, Romy Schlögel¹

¹ Spacebel sa., Earth Observation Applications, Angleur, Belgium – eoservices@spacebel.be

Keywords: Mixed-Forests, Spatial Deep Learning, Tree Species classification, Object Detection, Convolutional Neural Network

Abstract

Forests play a pivotal role in global ecosystems by sequestering carbon, preserving biodiversity, and providing valuable resources for both humans and wildlife. Monitoring and management of these forests require accurate, up-to-date information on individual trees and species composition—challenges that can be addressed with advanced remote sensing and deep learning. This paper presents a multi-season, multi-year approach to automatic tree detection and species classification in heterogeneous forests. Using over 5,000 high-resolution (0.25 m) RGB orthophoto tiles from the Wallonia region (spanning 2018–2023), we annotated more than 100,000 individual trees representing 14 classes of deciduous and coniferous species. A Faster R-CNN model trained for tree detection achieved a F1 score of 0.828 and a mAP@50 of 0.827, effectively locating tree crowns under varying illumination and phenological conditions. Meanwhile, a convolutional neural network (CNN) for species classification attained an overall accuracy of 0.937, accurately distinguishing most species and age classes. Despite strong performance, limitations persist, particularly in identifying small saplings and visually similar species (e.g., oak vs. beech). These findings highlight the potential of multi-temporal aerial imagery and deep learning to enhance forest inventories, reduce field survey costs, and inform targeted management.

1. Introduction

Forests serve as critical ecosystems, regulating local and global climates through carbon sequestration and sustaining biodiversity by providing habitats for countless species (Harris et al., 2021). As anthropogenic pressures like deforestation, land-use change, and climate variability intensify, precise and up-to-date forest inventories are essential for effective conservation strategies and sustainable resource management. Tree-level information—encompassing species composition, tree dimensions, and health status—offers especially valuable insights into forest dynamics and resilience (Fassnacht et al., 2016). Traditionally, field surveys have constituted the gold standard for such detailed inventories, but their high cost and labour intensity limit their scalability, particularly in large or remote forest expanses.

Against this backdrop, the surge in high-resolution remote sensing data and advanced machine learning techniques has created new avenues for forest monitoring. Many existing methods rely on single-season or single-year imagery, often at coarser resolutions (e.g., satellite platforms) that may be insufficient for delineating individual crowns or differentiating visually similar species (Michez et al., 2016; Thapa et al., 2024). Although these approaches have demonstrated success in detecting broad forest patterns, they often fall short when forest stands are diverse, phenological stages vary significantly, or spatial resolution is inadequate for capturing fine-grained canopy differences (Immitzer et al., 2012; Bolyn et al., 2022). Additionally, single-date analyses fail to account for how tree appearance and canopy reflectance may shift across phenological phases, such as the transition from leaf-off to leaf-on conditions in deciduous species.

To address these challenges, we propose a multi-season and multi-year tree detection and classification framework rooted in deep learning with high-resolution aerial orthophotos. By integrating imagery from 2018 to 2023, our method captures subtle spectral-textural variations that arise under changing

phenological states and lighting conditions—a key gap in many prior studies focused on single-season acquisitions. Specifically, we combine a Faster R-CNN architecture (Ren et al., 2015) for robust tree-crown delineation with a specialized convolutional neural network for classifying 14 tree classes, including mature and young coniferous and deciduous species. This approach provides comprehensive tree-level data in heterogeneous forest parcels, a notoriously difficult environment where overlapping canopies and variable crown shapes frequently hinder traditional classification models.

In doing so, our framework offers three principal contributions. First, we assemble a large, annotated dataset (over 100,000 labelled trees) reflecting diverse stand types and phenological stages across multiple years, ensuring broad coverage of the Wallonia region in Belgium. Second, we show how multi-year imagery can alleviate species confusion by capturing seasonal colour changes, facilitating clearer distinctions between visually similar species (e.g., oak vs. beech). Finally, we present an in-depth evaluation on large, real-world forest parcels, highlighting not only the overall accuracy gains but also the persistent challenges in dense canopies and small saplings that future research must address. By advancing tree-level detection and classification in multi-season contexts, this work lays the groundwork for more accurate, scalable forest inventories and targeted management interventions.

The contributions of this work are as follows:

1. Development of a high-resolution tree detection framework using Faster R-CNN, trained on extensive annotated datasets.
2. Implementation of a species classification model capable of distinguishing between 14 tree classes, including young and mature trees.
3. Analysis of multi-seasonal and multi-year imagery to assess the impact of temporal variations on model performance.

- Validation of the proposed approach on diverse forest parcels across the Wallonia region to assess its applicability in real-world forestry management.

2. Methodology

2.1 Dataset Creation

We built our dataset using high-resolution (0.25 m) RGB orthophotos of Wallonia. A total of 5,098 non-overlapping tiles, each measuring 50×50 meters (covering 1,274.5 ha), were selected to represent diverse forest conditions across multiple seasons from 2018 to 2023. Expert annotators used QGIS to draw bounding boxes around individual tree crowns (in mono-species parcels provided by local stakeholders), assigning each to a predefined species label. To reduce inter-annotator variability, multiple reviewers cross-verified ambiguous cases. The final dataset comprised over 100,000 labelled instances, each reflecting a distinct tree. We then split the data into training (70%), validation (15%), and test (15%) sets, ensuring all species appeared in each subset and maintaining geographical diversity where possible.

2.2 Data Augmentation Techniques

2.2.1 Tree Detection: For the Faster R-CNN tree detection model, random horizontal and vertical flips were applied, with bounding box coordinates updated accordingly. Colour jitter was also introduced by adjusting brightness, contrast, hue, and saturation by $\pm 10\%$, helping the model learn to handle variations in canopy appearance due to changing light conditions.

2.2.2 Tree Species Classification: The model employed two key augmentations: random brightness and random contrast, both at $\pm 10\%$. Although hue and saturation adjustments can be useful, preliminary experiments indicated these two augmentations were sufficient to address typical variations seen in aerial imagery. By applying these transformations, we mitigated biases related to illumination and positional differences, ultimately enhancing the model's generalization across different seasons.

2.3 Model Architecture and Training

2.3.1 Faster R-CNN for Tree Detection: The Faster R-CNN model comprises three key components: a convolutional feature extractor, a Region Proposal Network (RPN), and a classification head. We used a ResNet-50 backbone pretrained on ImageNet for feature extraction, as it balances accuracy and computational efficiency. The anchor sizes and aspect ratios in the RPN were chosen to accommodate typical crown sizes observed in our dataset (e.g., small saplings to large mature trees), and images were resized to a fixed height and width 200×200 pixels while maintaining aspect ratio.

We conducted an iterative hyperparameter tuning process to determine the optimal learning rate, weight decay, and batch size, experimenting with values that minimized validation loss without overfitting. Table 1 presents the final selected hyperparameters.

Hyper-parameter	Selected values
Batch size	1
Epochs	150
Learning rate	0.0001
Optimizer	SGD (Stochastic Gradient Descent)
Weight decay	0.0005

Table 1. Hyper-parameters of the Faster R-CNN

2.3.2 CNN for Species Classification: A sequential CNN was implemented in Keras (Chollet, 2015). Each input image is cropped to $48 \times 48 \times 3$ pixels. Data augmentation (random brightness and contrast at 10%, as well as horizontal/vertical flips) further increases robustness against illumination changes and viewpoint variations.

The network consists of four convolutional blocks, each featuring:

- A convolution layer (using a ReLU activation)
- A max-pooling layer
- Dropout (e.g., 0.2)
- Batch normalization

After flattening, four fully connected layers—with additional dropout and batch normalization—help the network generalize and avoid overfitting. The final output layer uses a SoftMax activation to produce class probabilities for our 14 total classes (Spruce, Douglas, Larch split into mature and young sub-classes). The model summary is presented in Table 2.

Layer (Type)	Output Shape	Parameters
RandomBrightness	(48, 48, 3)	0
RandomContrast	(48, 48, 3)	0
Conv2D + ReLU	(48, 48, 32)	896
MaxPooling2D	(24, 24, 32)	0
Dropout + BatchNorm	(24, 24, 32)	128
Conv2D + ReLU	(24, 24, 64)	18,496
MaxPooling2D	(12, 12, 64)	0
Dropout + BatchNorm	(12, 12, 64)	256
Conv2D + ReLU	(12, 12, 128)	73,856
MaxPooling2D	(6, 6, 128)	0
Dropout + BatchNorm	(6, 6, 128)	512
Conv2D + ReLU	(6, 6, 256)	295,168
MaxPooling2D	(3, 3, 256)	0
Dropout + BatchNorm	(3, 3, 256)	1,024
Flatten	(2304)	0
Dropout	(2304)	0
Dense + ReLU	(256)	590,08
BatchNormalization	(256)	1,024
Dense + ReLU	(128)	32,896
BatchNormalization	(128)	512
Dense + ReLU	(64)	8,256
Dropout + BatchNorm	(64)	256
Dense (Softmax)	(14)	910
Total Parameters	—	1,024,270

Table 2 Summary of the Sequential model architecture used in this study. The model consists of four convolutional blocks with increasing depth, interleaved with pooling, dropout, and batch normalization layers, followed by fully connected layers for final classification into 14 classes.

We train the model using the Adam optimizer and a categorical cross-entropy loss. Evaluation metrics include accuracy and weighted accuracy to address class imbalances. Table 3 summarizes the key hyperparameters.

Hyper-parameter	Selected values
Batch size	32
Epochs	200
Epochs	200
Learning rate	0.001
Optimizer	Adam

Table 3. Hyper-parameters for the tree classification model

2.4 Error analysis

To gain deeper insights into model behaviour, we plotted learning and loss curves for each training run and compared them across different seasons. This allowed us to identify points of overfitting (e.g., diverging validation and training loss) and assess seasonal differences. For example, performance in spring imagery sometimes dipped due to partial leaf coverage, whereas summer imagery provided more distinct canopy features. By comparing error metrics (precision, recall, F1-score) in these seasonal subsets, we gleaned where our models excel and where further improvements may be necessary.

A summary of the methodology is described in Figure 1.

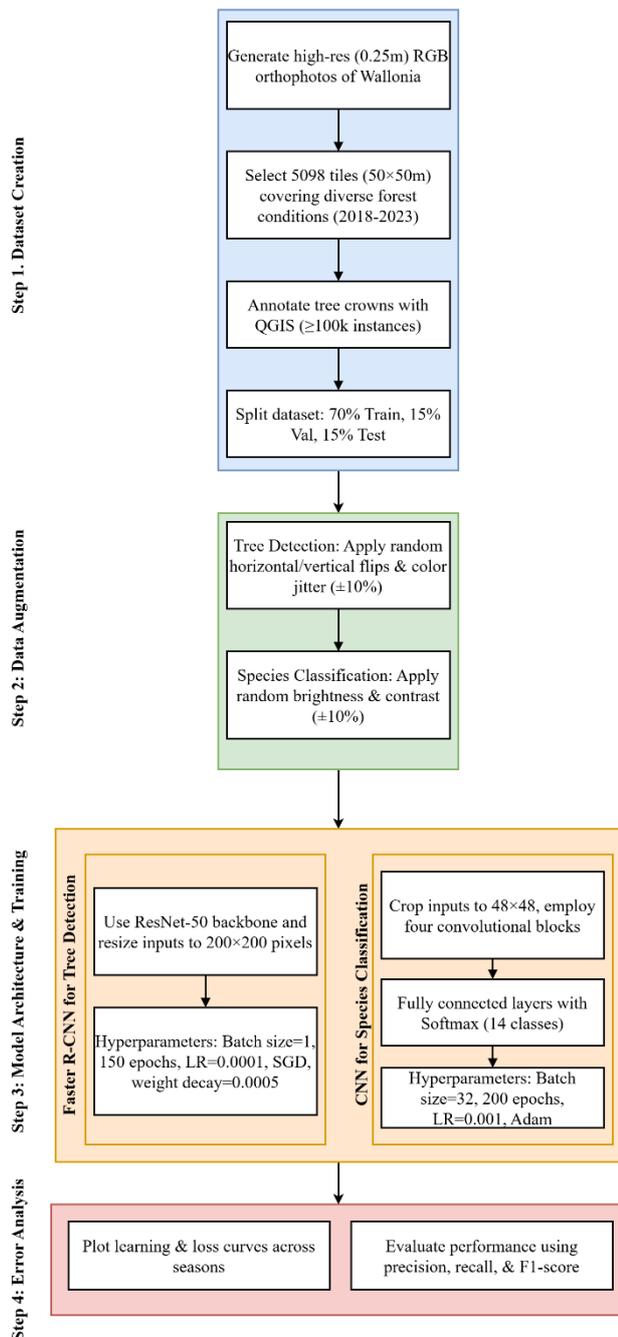


Figure 1. Graphical illustration of the methodology used for the multi model training for tree species classification.

3. Results

3.1 Tree Detection performance

The tree detection model, based on the Faster R-CNN architecture (Ren et al., 2015) and trained using the hyperparameters in Table 1, achieved precision of 0.796, recall of 0.862, an F1 score of 0.828 (at IoU@0.5) illustrated in Figure 2, and a mean Average Precision (mAP@50) of 0.827 as shown in Figure 3. These metrics were derived from the validation set using a confidence detection threshold of 0.1.

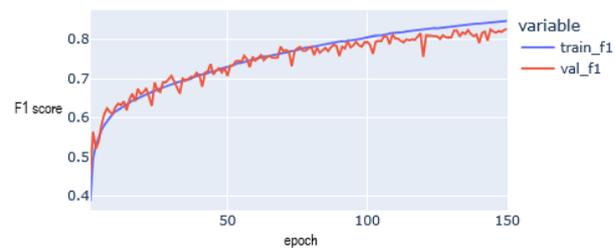


Figure 2. Faster RCNN F1 score on train and validation

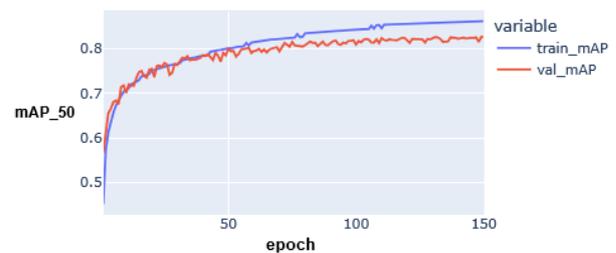


Figure 3. Faster RCNN mAP_50 score on train and validation.

Figure 4 illustrates sample detection outputs on a full parcel where trees are outlined with bounding boxes.



Figure 4. Results of the Faster R-CNN predictions on a parcel near Arlon, Wallonia Belgium.

The Figure 5 and Figure 6, show the tree detection in different scenarios and contexts:

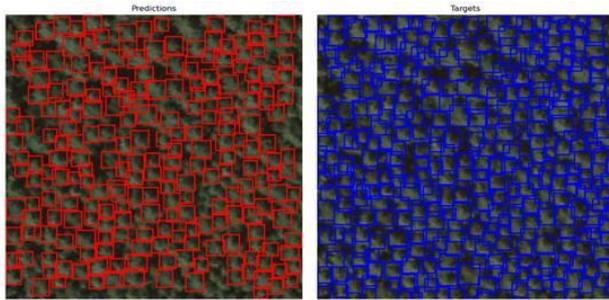


Figure 5. Tree detection results on the validation set in a dense parcel with young spruce.



Figure 6. Tree detection results on the validation set in a low populated tile with mature trees.

3.1.1 Confidence Threshold and Detection Strategy: A low threshold (0.1) was deliberately chosen to ensure a high recall and reduce the risk of missing smaller or more ambiguous trees. This approach prioritizes a minimal omission error, which is critical when comprehensive forest inventories are desired. In scenarios where users desire fewer false positives (e.g., precise tree counts in urban parks), additional filtering can be applied post-inference based on detection score, bounding box dimensions, or object shape ratios.

3.1.2 Error Patterns and Future Considerations. False positives tended to occur in dense canopies or at forest edges where shadows and overlapping tree crowns complicated detection. While the false negatives primarily involved small, closely spaced saplings or areas with homogeneous canopy cover lacking distinct boundaries. These results confirm that the chosen architecture (Faster-RCNN) and hyperparameters, described in the Table 1, effectively locate individual trees under diverse conditions and across different year acquisitions.

3.2 Tree species Classification and Validation Results

Using the CNN described in Section 3.3.2, the global validation accuracy across all 14 classes was 0.937 as shown in Figure 7. This represents a high level of performance in distinguishing both mature and young trees for a broad range of species. Figure 2 provides an example classification map, where each detected tree is color-coded by its predicted species.

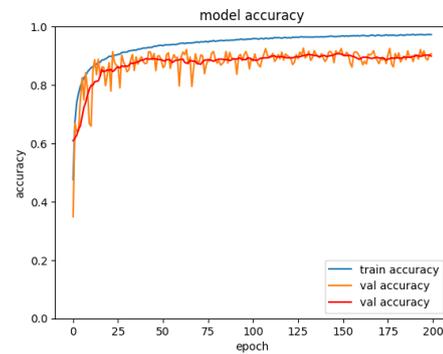


Figure 7. Sequential model accuracy on the train and validation sets.

Table 4 lists the recall for each species, highlighting per-class performance. Recall values above 90% are observed for most classes, with some variability across different species.

Class	Precision	Recall	F1-score
Birch	83,86	94,25	88,75
Oak	93,24	88,24	90,67
Douglas Fir	96,25	87,19	91,5
Young Douglas Fir	98,64	97,98	98,31
Beech	87,52	87,07	87,29
Wild Cherry	81,33	95,31	87,77
Larch	92,26	87,82	89,99
Young Larch	99,87	98,5	99,18
Poplar	92,97	92,68	92,82
Scots Pine	96,93	94,19	95,54
Spruce	94,49	98,09	96,25
Young Spruce	98,88	99,58	99,23
Maple	95,65	100	97,78

Table 4. Recall per class for the species classification model

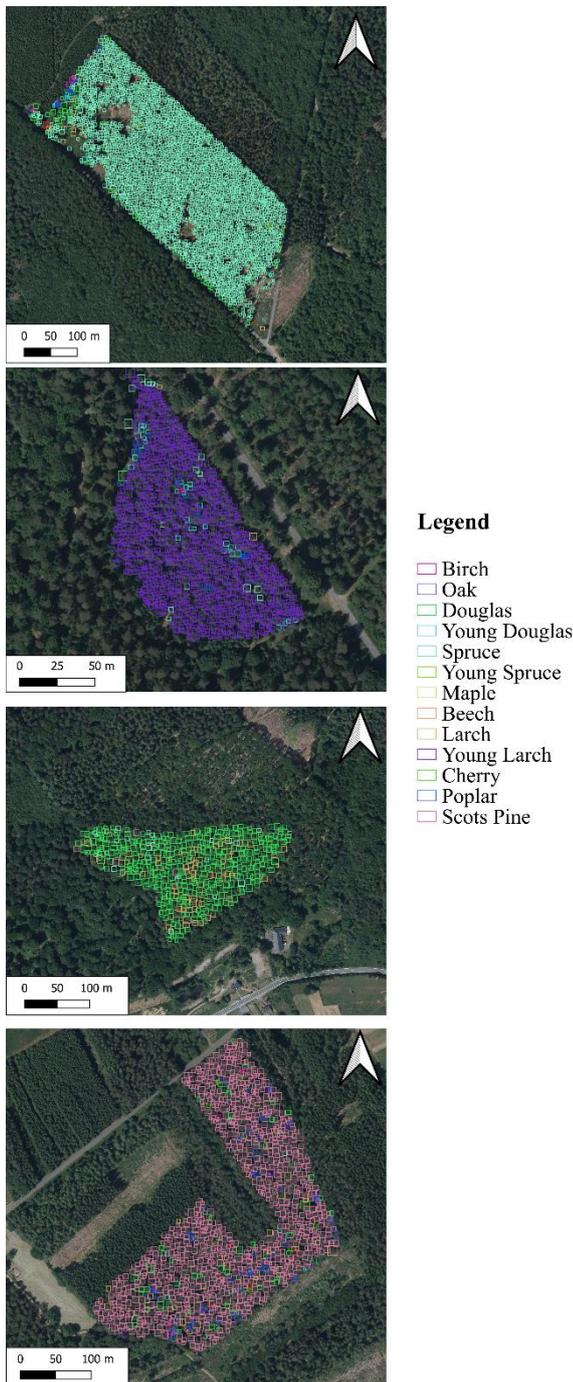


Figure 8. Results of the tree species classification map using the sequential model

Figure 9 shows the confusion matrix for our 14-class species classification, capturing both correctly identified classes (diagonal entries) and misclassifications (off-diagonal entries). Overall, most classes exhibit high diagonal counts—for instance, oak (2,645), Douglas fir (2,569), and spruce (2,863)—indicating strong performance in correctly recognizing these species. Young conifers (e.g., young Douglas, young larch, and young spruce) also display notably large correct predictions (1,784; 2,665; and 1,774, respectively), reflecting the model’s ability to handle different growth stages. However, there are some key misclassifications. For instance, oak occasionally appears as beech or other species, suggesting that subtle canopy similarities or seasonal variations might cause confusion. Douglas fir and

larch similarly overlap, particularly when younger trees share comparable coloration or crown textures. Although these off-diagonal values are relatively small compared to the main diagonal counts, they underscore the challenge of distinguishing visually similar species—especially under changing phenological states.

Actual Species	Unclas.	Other spp.	Birch	Oak	Douglas Fir	Young Douglas	Beech	Cherry	Larch	Young Larch	Poplar	Scots Pine	Spruce	Young Spruce	Maple
Unclas.	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Other spp.	3	553	0	7	0	1	14	3	5	0	0	0	0	0	0
Birch	5	1	747	5	0	3	2	9	0	0	10	3	0	0	0
Oak	40	67	11	2645	19	0	91	37	6	0	12	11	1	0	3
Douglas Fir	41	14	5	25	2569	7	27	5	100	0	3	66	42	3	0
Young Douglas	6	0	2	0	2	1784	0	3	2	3	49	0	0	4	0
Beech	34	5	29	231	1	0	1699	28	7	0	17	6	0	0	5
Cherry	2	32	4	11	0	0	18	694	3	0	3	3	0	0	0
Larch	27	4	14	11	131	5	48	10	1085	0	3	12	7	3	2
Young Larch	5	2	3	0	2	7	1	0	0	2665	12	3	1	0	0
Poplar	2	2	27	7	0	2	0	9	0	0	1130	0	0	0	0
Scots Pine	24	17	62	52	40	4	4	7	3	0	5	2452	3	0	0
Spruce	15	0	6	6	43	2	9	3	2	0	4	9	2863	1	0
Young Spruce	1	0	0	0	0	0	1	0	0	0	0	0	0	1774	0
Maple	4	0	0	2	0	0	1	0	0	0	0	0	0	0	242
Predicted Species	Unclas.	Other spp.	Birch	Oak	Douglas Fir	Young Douglas	Beech	Cherry	Larch	Young Larch	Poplar	Scots Pine	Spruce	Young Spruce	Maple

Figure 9. Confusion matrix showing the number of classified samples for each tree species pair. Rows represent actual species and columns represent predicted species. Values along the diagonal indicate correctly classified instances.

3.2.1 Performance analysis: According to the species types and their level of maturity, various performances are observed. Maple, young spruce, and young larch stand out with recall values near or exceeding 99%. These results suggest that the CNN effectively captures the visual cues (colour, texture, canopy structure) for these species, even across different seasonal image sets. Beech (87.52%) and wild cherry (81.33%) show comparatively lower recall. This may be due to leaf-off conditions in some orthophotos for deciduous species like beech, or because larch needles and canopy structure share similarities with other conifers in certain imagery.

Misclassifications might also stem from the temporal variability (e.g., early spring vs. mid-summer) that alters leaf colour and canopy density.

3.2.2 Service operationalization and reporting: To streamline the dissemination of model outputs across multiple parcels, we developed an automated reporting pipeline. This system generates comprehensive PDF reports containing key performance metrics, data visualizations (e.g., detection plots and classification maps), and summary tables. By automating these processes, results can be systematically and consistently shared with stakeholders, facilitating faster, data-driven decision-making in diverse forest management contexts.

4. Discussion and conclusion

Our results suggest that the proposed deep learning framework for tree detection and species classification can serve as a valuable resource for forest stakeholders—ranging from regional forestry managers to researchers studying forest dynamics. By leveraging high-resolution orthophotos collected over multiple

seasons and years, we captured variations in canopy appearance and structure, thereby enhancing the temporal robustness of our model. The Faster R-CNN (Ren et al., 2015) achieved commendable precision and recall, affirming its suitability for identifying trees in diverse canopy conditions. Meanwhile, our species classification CNN showed high overall accuracy, demonstrating promise as a rapid and scalable approach to mapping forest composition.

Despite these achievements, our investigation revealed several important limitations:

4.1 Smooth Canopies and Dense Tree Cover

In areas with uniform or smooth deciduous canopies, such as mature beech stands with little shadow contrast, the detection model struggled to correctly isolate individual tree crowns. Dense forest blocks often led to overlapping bounding boxes or omission errors.

This issue underscores the need for multi-scale approaches or additional data sources (e.g., LiDAR) to better delineate tree crowns, especially in scenarios where aerial imagery alone provides insufficient texture cues.

4.2 Species-Level Classification Variability

While the classification model exhibited high accuracy for several species, there were discrepancies when applied to new study areas. Lower performance likely stems from factors such as overfitting to local conditions, seasonal variability, or different image acquisition geometries. Conifers generally performed better than deciduous species, potentially due to distinct canopy textures. However, early growth stages of conifers, especially small saplings, were often misclassified as other conifers. This highlights the need for more training examples capturing different developmental stages.

4.3 Underrepresentation of Small Trees

Trees with crown diameters under 2 meters frequently went undetected or unclassified, a limitation with implications for sustainable forest management if regenerating populations are overlooked. Addressing this would require higher spatial resolution imagery, improved annotation protocols, or complementary data from ground surveys.

4.4 Confusion Among Similar Species

Oak and beech were sometimes mistaken for one another, presumably due to their comparable leaf coloration and canopy shape in aerial imagery. Such confusions reduce the reliability of species-specific metrics and highlight the potential benefit of spectral data for better species separation.

Despite these challenges, the overall performance of the detection and classification models is promising for large-scale, near-real-time forestry applications. When integrated into forestry management workflows, these models can streamline inventory processes, track changes in species composition over time, and inform more targeted interventions (e.g., pest management, planting efforts).

5. Outlook

5.1 Multi-Spectral and Hyperspectral Data

Incorporating near-infrared (NIR) bands, or even hyperspectral images, can significantly improve species discrimination, particularly for deciduous species that appear visually similar in RGB imagery. NIR reflects leaf chemistry differences, aiding in separating oak from beech or identifying subtle stress indicators.

5.2 Enhanced Model Architectures

Employing more complex or specialized deep learning architectures—such as Vision Transformers or Swin Transformers—could capture more nuanced canopy features and better handle complex forest structures. Increasing the model's capacity (e.g., deeper networks, attention mechanisms) may help with challenging tasks like detecting very young saplings or distinguishing species under variable seasonal conditions.

5.3 Tree Height and Structural Information

Integrating LiDAR or photogrammetric data would provide 3D information on canopy height and structure, enabling more robust tree detection. This is particularly beneficial for underrepresented small trees, whose limited vertical extent might go unnoticed in 2D imagery alone.

5.4 Health and Stress Indicators

Including Normalized Difference Vegetation Index (NDVI) or similar spectral indices can offer timely insights into tree health, water stress, or pest infestations. Early detection facilitates proactive management by allowing stakeholders to intervene before irreversible damage occurs.

Integrating these health indicators directly into the model (e.g., multi-task learning approaches that predict both species type and health status) could further streamline decision-making.

5.5 Scalability and Transferability

Future efforts should focus on replicating these methods in other forest regions with different species compositions or management practices. Demonstrating generalization across geographies would bolster confidence in the model's utility and adaptability.

6. References

- Bolyn, C., Lejeune, P., Michez, A., & Latte, N. (2022). Mapping tree species proportions from satellite imagery using spectral-spatial deep learning. In *Remote Sensing of Environment* (Vol. 280, p. 113205). Elsevier BV. <https://doi.org/10.1016/j.rse.2022.113205>
- Chollet, F., & others. (2015). Keras. Retrieved from <https://keras.io>
- Fassnacht, F. E., Latifi, H., Stereńczak, K., Modzelewska, A., Lefsky, M., Waser, L. T., Straub, C., & Ghosh, A. (2016). Review of studies on tree species classification from remotely sensed data. In *Remote Sensing of Environment* (Vol. 186, pp. 64–87). Elsevier BV. <https://doi.org/10.1016/j.rse.2016.08.013>

Fassnacht, F. E., Latifi, H., Stereńczak, K., Modzelewska, A., Lefsky, M., Waser, L. T., Straub, C., & Ghosh, A. (2016). Re-view of studies on tree species classification from remotely sensed data. In *Remote Sensing of Environment* (Vol. 186, pp. 64–87). Elsevier BV. <https://doi.org/10.1016/j.rse.2016.08.013>

Harris, N. L., Gibbs, D. A., Baccini, A., Birdsey, R. A., de Bruin, S., Farina, M., Fatoyinbo, L., Hansen, M. C., Herold, M., Houghton, R. A., Potapov, P. V., Suarez, D. R., Roman-Cuesta, R. M., Saatchi, S. S., Slay, C. M., Turubanova, S. A., & Tyukavina, A. (2021). Global maps of twenty-first century forest carbon fluxes. In *Nature Climate Change* (Vol. 11, Issue 3, pp. 234–240). Springer Science and Business Media LLC. <https://doi.org/10.1038/s41558-020-00976-6>

Immitzer, M., Atzberger, C., & Koukal, T. (2012). Tree Species Classification with Random Forest Using Very High Spatial Resolution 8-Band WorldView-2 Satellite Data. In *Remote Sensing* (Vol. 4, Issue 9, pp. 2661–2693). MDPI AG. <https://doi.org/10.3390/rs4092661>

Oghaz, M.M.D., Saheer, L.B., Zarrin, J. (2022). Urban Tree Detection and Species Classification Using Aerial Imagery. In: Arai, K. (eds) *Intelligent Computing. SAI 2022. Lecture Notes in Networks and Systems*, vol 507. Springer, Cham. https://doi.org/10.1007/978-3-031-10464-0_32

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ... Chintala, S. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library (Version 1). <https://doi.org/10.48550/ARXIV.1912.01703>

Ren, S., He, K., Girshick, R., & Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks (Version 3). doi.org/10.48550/ARXIV.1506.01497

Sibanda, M., Mutanga, O., & Rouget, M. (2015). Examining the potential of Sentinel-2 MSI spectral resolution in quantifying above ground biomass across different fertilizer treatments. In *ISPRS Journal of Photogrammetry and Remote Sensing* (Vol. 110, pp. 55–65). Elsevier BV. <https://doi.org/10.1016/j.isprsjprs.2015.10.005>

Thapa, B., Darling, L., Choi, D. H., Ardohain, C. M., Firoze, A., Aliaga, D. G., Hardiman, B. S., & Fei, S. (2024). Application of multi-temporal satellite imagery for urban tree species identification. In *Urban Forestry & Urban Greening* (Vol. 98, p. 128409). Elsevier BV. <https://doi.org/10.1016/j.ufug.2024.128409>