Fuel detection in forest environments training deep learners with smartphone imagery

Francesco Pirotti[1,2], Alessandro Carmelo[1], Erico Kutchartt[3,4,1]

[1] Department of Land, Environment, Agriculture and Forestry (TESAF), University of Padova. Via dell’Università 16, 35020 Legnaro, Italy - francesco.pirotti@unipd.it, alessandro.carmelo@studenti.unipd.it
[2] Interdepartmental Research Centre in Geomatics (CIRGEO), University of Padova. Corte Benedettina, Via Roma 34, 35020 Legnaro, Italy
[3] Forest Science and Technology Centre of Catalonia (CTFC). Carretera de Sant Llorenç de Morunys Km 2, 25280 Solsona, Spain – erico.kutchartt@ctfc.cat, erico.kutchartt@unipd.it
[4] Joint Research Unit CTFC – AGROTECNIO. Carretera de Sant Llorenç de Morunys Km 2, 25280 Solsona, Spain

Keywords: Object detection, Yolo8, Forestry, Fuel, Deep Learning, Image Segmentation

Abstract

Unmixing mixtures in images is one of the challenges for extracting information from data. Forest environments are particularly complex due to the relatively irregular structure of trees, shrubs and low vegetation. The amount and condition of vegetation, i.e. thin vs thick branches, trunk vs leaves, understorey and litter provide information to infer the amount of burnable fuel and consequently a key factor for predict fire behaviour. In this work we test a deep learning framework for training and testing the performance of detecting logs and litter of broadleaves and conifers in imagery of forest environments recorded through smartphones. Roboflow and YOLOv8 were employed, using a dataset of forest images manually segmented in four classes: “broadleaf-litter”, “broadleaf-logs”, “conifer-litter” and “conifer-logs”. The results indicate that the "Extra-large Instance Segmentation" model achieved the best performance with F1-score value of 0.79 at a confidence of 0.763 on familiar images in the validation phase with 214 epochs, whereas the "Large Instance Segmentation" model was more effective on new images in the test phase, as expected with a lower F1-score of 0.24 and a confidence value of 0.492. It was observed that this was due mostly to omission errors due to low light conditions in the forestry environment. We conclude that segmenting key elements and including varied images in terms of seasonality and lighting conditions could potentially improve performance. This work lays a useful foundation for refining the use of AI in forest fuel monitoring.

1. Introduction

Forest fires represent a sensitive issue within the forestry sector. They are a difficult phenomenon to investigate in depth, as they depend on ecological, physical and anthropogenic factors that combine to provide a high heterogeneity of variables. They represent natural phenomena in many ecosystems, promoting biodiversity and natural regeneration, although human activity often alters their natural regimes, leading to different scenarios (Chuvieco et al., 2021).

Europe is widely affected by wildfires, to a lesser extent in the north, while Mediterranean areas contribute 94 percent of the total area burned. Moreover, forest fires in southern Europe have largely increased since the second half of the 1900s, partly due to the abandonment of rural areas (Xanthopoulos et al., 2006). This phenomenon involves, among other things, the loss of control over the growth of herbaceous, shrub and tree vegetation and thus over the quantity and quality of forest fuel, which is one of the three components of the fire triangle, along with the igniter. These are the three basic elements that allow a forest fire to occur. Fuel is perhaps the element that can be manipulated the most, but at the same time is difficult to control.

In a forest fire, in fact, the flame front advances mainly through herbaceous or litter fuel. The rate of advance of the flame front is influenced by load, bulk density, size of fuel components (leaf size), caloric content, and extinction moisture (Rothermel, 1972). Changing any of these variables results in changes on the propagation of a possible phenomenon. Delving into their dynamics and behavior, therefore, has become a challenge for research groups that invest economic resources each year with the aim of learning about and then controlling such natural events.

Regarding fuel models, those proposed by Scott and Burgan (2005), provide mathematical models of fire behavior for several possible situations. Aragoneses et al. (2023) took up those of Scott and Burgan (2005), developing a 1-km spatial resolution map on a pan-European scale to identify possible fire behavior in different areas, which was later updated by Kutchartt et al. (2024) at 100 m resolution, providing the dataset in an interactive Web-GIS server.

The H2020 FIRE-RES Project interpolated these data with the other satellite-derived data mentioned earlier, as well as from Corine Land Cover mapping to obtain a raster with spatial resolution of 100m x 100m (i.e., one hectare) that allows the creation of increasingly detailed fire propagation models. The data obtained through these interpolations, however, are certainly not error-free. Certainly, obtaining data over such a small area from satellites is relevant, but still in one hectare of area the actual situation may not coincide with the remote sensed data.

To optimize the accuracy of the information, the H2020 FIRE-RES Project proposes an intuitive and easy-to-use smartphone application, even by inexperienced personnel, to collect fuel data, allowing you to take georeferenced photographs with your device. This application is called the FIRE-RES Geo-Catch app (Kutchartt et al., 2023), to date already used by over two hundred users, collecting more than six thousand images across Europe. These provide data on the plant formations present in a given area, which can be classified, depending on composition,

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-G-2025-649-2025 | © Author(s) 2025. CC BY 4.0 License.

into fuel categories, echoing those proposed by Scott and Burgan (2005) and Aragoneses et al. (2023).

2. Materials and methods

2.1 Study area

The images used in this work are taken from multiple users using their smartphones. Each image catches the forest ecosystem at a specific location, mostly in Europe (Figure 1). Therefore, there is not a specific defined area for this work, but a collection of images that represent very different scenarios. The training was done with a subset of images taken from the online repository that can be seen in Figure 1. This repository collects automatically users’ images that are taken with smartphone imagery and displays them immediately if there is an internet connection when the image is taken. As in forestry field work it is often the case that there is no internet connection, then in this scenario the images are stored on the smartphone memory and uploaded automatically once an internet connection is detected.

2.2 Smartphone imagery FIRE-RES Geo-Catch app

This synched method for collecting images is done via the FIRE-RES Geo-Catch app which was developed into the framework of the H2020 FIRE-RES project. The main objective of developing an app that collects images differently from existing apps for smartphones is to keep a simple interface, with a very flat learning curve and intuitive process. The app automatically embeds key information in the image EXIF tag. This information is the project name, coordinates for geolocation, orientation angles, and accuracy of geolocation, plus the user’s device identification, which keeps the privacy of the users that are providing the information by an unique user id (UID), containing a 10-character alphanumeric string that is assign to the user’s phone and browsers, instead of storing personal data (e.g., e-mails). This keeps the user privacy, while maintaining responsibility for images that get uploaded. If a user uploads images that are not part of any projects, the device will get blacklisted.

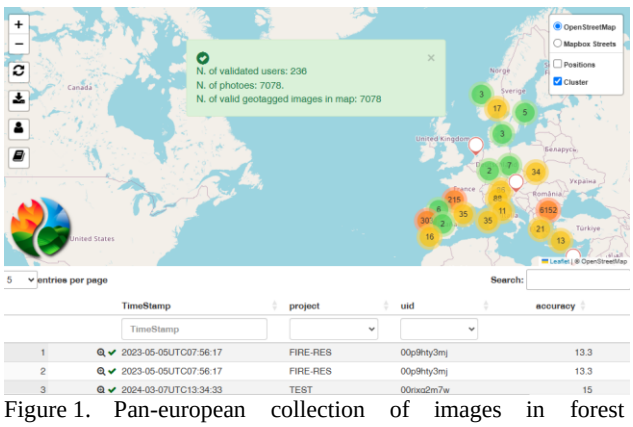

Figure 1. Pan-european collection of images in forest ecosystems used for image analysis. Each point is an image taken with a smartphone with project id, timestamp, location coordinates and orientation (azimuth and zenith angles).

2.3 Machine learning and deep learning

Once a number of images are available, machine learning and deep learning approaches are applied to train models to allow to detect specific elements in the image. Machine Learning is a technique that improves the performance of a system by learning from experience through computational methods (Zhi-Hua, 2021). In computer systems, experience is manifested in the form of data. The main task of machine learning is to develop learning algorithms that build models from data. By providing the learning algorithm with data, it obtains a model capable of making predictions. Deep learning is a branch of machine learning that is composed of several layers, including input and output layers, which allow the realization of multiple stages of nonlinear information processing. These multi-level architectures are used for feature learning and model classification (Alom et al., 2019).

2.4 Roboflow and YOLO-v8

Roboflow is an online platform that allows developers and researchers to train artificial intelligence models from image datasets by segmenting and labeling the elements contained in them. It was here used for carrying out these tasks. It has various applications in a variety of fields, from healthcare, for cancer mass recognition (Al-masni et al., 2018), to security, through intrusion detection (Hanan Ashraf et al., 2022). The main artificial intelligence technologies employed by Roboflow include Machine Learning and Deep Learning models. These models, specifically, are based on convolutional neural networks (CNNs), which are employed for tasks such as object recognition and image classification and segmentation. For example, YOLO (You Only Look Once) represents one available algorithm and is used here. YOLO is composed of convolutional layers that extract features from the image (e.g., edges, textures, shapes, etc.). These are subsequently transformed into one-dimensional vectors (Figure 2) and used as input to fully connected layers that determine the final classification (Hussain, 2023).

There are several CNN architectures for feature extraction, such as AlexNet, VGGNet, GoogleNet and ResNet, each with specific features to improve performance and accuracy in image classification. Object detection can be divided into two main approaches. Two-stage methods initially select a set of possible regions where objects might be located (first stage), and then make a detailed prediction about the selected regions (second stage). One-stage methods simplify the process by turning it into a regression problem that performs region selection and prediction in a single step. This makes them faster and less complex to compute, although with lower accuracy than two-stage methods (Figure 2).

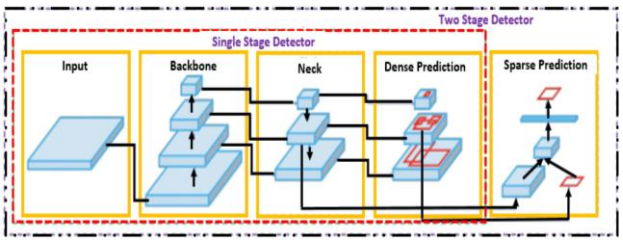

Figure 2. Schema of CNN workflows (Hussain, 2023)

For this work, the Roboflow platform was employed only for the phase of segmentation and label assignment to the image dataset. After creating the personal account on Roboflow, it is possible to start the segmentation phase, uploading the photos to be processed.

2.5 Input imagery for training, validation and test

A dataset consisting of 177 images was used. Of these, 136 images were taken with a smartphone camera by one of the authors and were employed for the training and validation phase: 105 for training 31 for validation and 41 for the testing phase. These 41 images were taken from the images collected in the H2020 FIRE-RES project (Figure 1) and were used for independent testing of the accuracy of results. All the 117 images for training were related to coniferous (*Picea abies*) and deciduous (*Fagus sylvatica*) forests during the late-winter season, which coincides with the period of maximum danger for forest fires for the reference geographic region in which the photographs were taken, corresponding with the Veneto pre-Alpine belt. The test images, on the other hand, were taken from the FIRE-RES Geo-Catch app portal to use completely independent datasets for testing. These two vegetation formations were chosen because they represent the standard of the most widespread forests along the Alps, with fairly homogeneous characteristics.

The photos taken had to be blur-free, with good contrast and had to represent portions of the forest as homogeneous as possible. The following four different classes were used:

- Conifer litter
- Broadleaf litter,
- Conifer stems
- Broadleaf stems

For each image, the different classes were assigned and manually segmented by contouring the elements in the image (Figure 3).

Figure 3. Result of the manual segmentation process of two images: the first (left column) related to the deciduous forest, the second (right column) to the coniferous forest. The process involves cropping the images to 640x640 pixels in Roboflow.

2.6 Training, validation and testing

For the next phases, i.e., training, validation and testing, YOLOv8 was employed in a remote HPC virtual machine with an Nvidia A40 GPU and 36 processors. All segmented images were assigned to the training (105 images), validation (31 images) and testing (41 images) phases. Figure 4 shows the number of segments used to train the model. It can be noted that the number of segments are slightly unbalanced with many belonging to logs, and the litter being under-represented. This is due to the nature of the litter class, which is not found in all the images, whereas logs are always present in the photos.

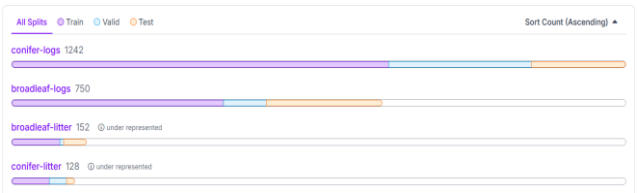

Figure 4. Number of segments per each class.

During the training phase, the model learns through the training data and optimizes the parameters. This phase involves n epochs of training to improve the performance of the model. At each epoch, the model processes all the samples in the dataset, updating the parameters with the aim of improving object recognition and classification capabilities. During each epoch, the model calculates the “loss”, which measures how far the predictions deviate from the correct labels. The goal of training is to gradually reduce this loss to improve model performance. The greater the number of epochs, the more opportunities the model must learn patterns in the data. However, training for too many epochs can lead to a problem known as overfitting, that is, the model learns the details of the training set too well and becomes less effective on test data or unpublished data (Srivastava et al., 2014).

In the validation phase, however, the model is evaluated on a separate dataset to optimize parameters and prevent overfitting. Finally, during the testing phase, the model is tested on completely new data to evaluate its overall performance and ensure that it can predict effectively. This process ensures that the model not only fits the training data well, but also that it can make accurate and reliable predictions on new data while maintaining good prediction ability (Yaseen, 2024).

In order to train the model, once the images have been segmented via the Roboflow interface via manual segmentation (Figure 3), the created dataset was exported in a workstation with a A40 NVIDIA GPU and 30+ cores for processing. In this work, RStudio IDE was used to interact with the necessary Python libraries and scripts for running the commands needed to apply YOLO model training and testing.

3. Results

3.1 Validation

The following table 1 reports on models that were tested and performance indicators for the training and validation phase.

Model	Epochs	F1-Score	Confidence
n-seg	100	0.75	0.342
n-seg	193	0.78	0.577
s-seg	10	0.69	0.902
s-seg	50	0.75	0.315

s-seg	100	0.78	0.478
s-seg	242	0.79	0.436
m-seg	100	0.78	0.559
l-seg	100	0.78	0.557
x-seg	50	0.79	0.535
x-seg	100	0.81	0.407
x-seg	214	0.79	0.763

Table. 1. Performance indicators: the first column shows the type of model used; the second column shows the number of training epochs; the third column shows the score for the F1-score parameter; and the fourth column shows the confidence interval within which the maximum score for the F1-score occurs.

The best performance was achieved by the last training/model strategy, with 214 epochs, employing the extra-large Instance Segmentation model. The 0.79 value for the F1-score is achieved at a confidence of 0.763.

3.2 Test and prediction

The performance results presented in Table 1 reflect the model's ability to accurately predict images contained within the training and validation dataset. If instead we test the model on completely new images, the results change, as highlighted in Table 2.

Model	Epochs	F1-Score	Confidence
n-seg	100	0.17	0.377
n-seg	193	0.15	0.147
s-seg	10	0.09	0.101
s-seg	50	0.16	0.197
s-seg	100	0.18	0.112
s-seg	242	0.15	0.622
m-seg	100	0.23	0.437
l-seg	100	0.24	0.492
x-seg	50	0.22	0.318
x-seg	100	0.22	0.113
x-seg	214	0.23	0.298

Table 2. Results obtained from the testing phase. See caption of table 1 for column definition.

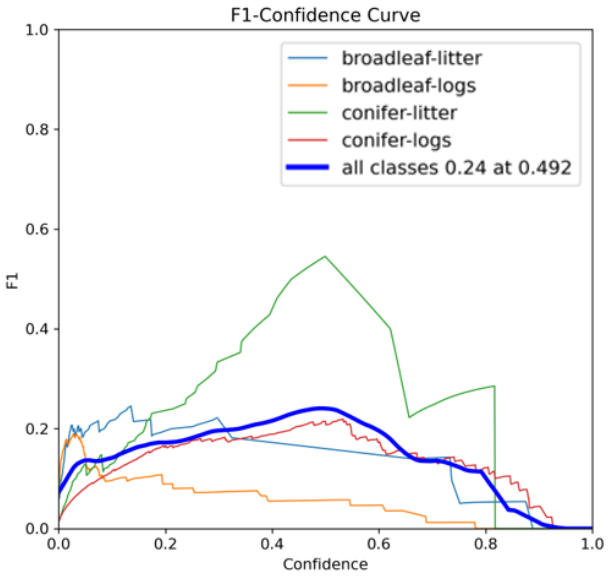

Figure 5. The F1-Score of the most efficient model relative to confidence

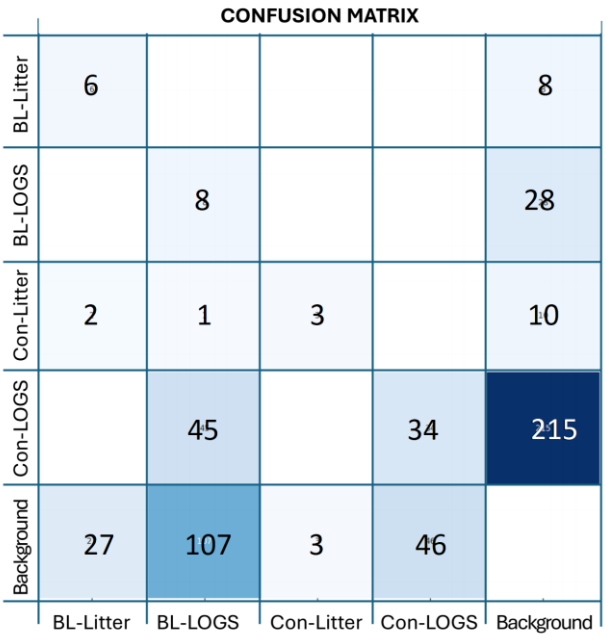

Figure 6. The performance of the most efficient model relative to the F1-Score parameter as a function of confidence (A) and the relative confusion matrix (B). “Con”=Conifers and “BL”=Broadleaves.

Table 2, in contrast to Table 1, shows results where the data used are completely new to the model and independent from the training set. As expected, the values of accuracy have dropped significantly with respect to Table 1.

The model that would seem to perform best with the test images is the large Instance Segmentation (l-seg) with 100 training epochs. The F1-Score index scores 0.24 at a confidence value of 0.492.

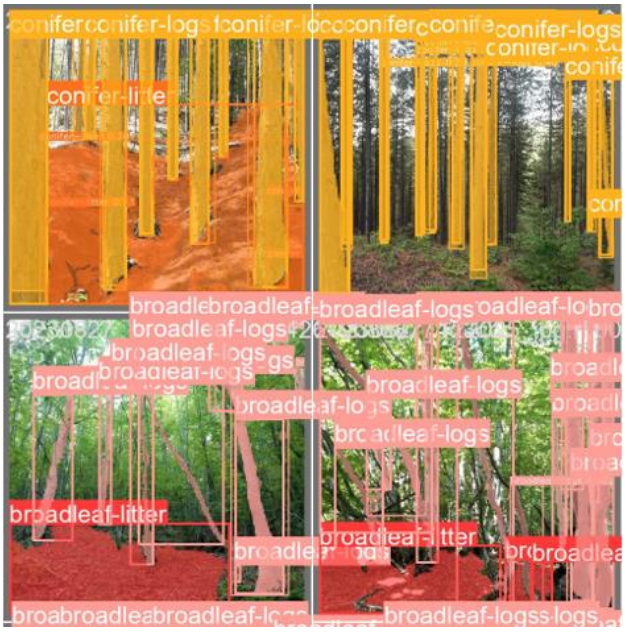

Figure 7. Examples of predicted outputs generated by the YOLOv8 model during the interface process.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-G-2025-649-2025 | © Author(s) 2025. CC BY 4.0 License.

Additional insights into the model’s performance is seen in Figure 5, Figure 6 and Figure 7. First, we can note that a large number of elements were not assigned to any class and were identified as “background.” This misclassification contributes to a high number of false negatives that weight a lot in the final F-score and overall values of accuracy metrics. Furthermore, it can be seen also that 45 of the elements belonging to the class “broadleaf-logs,” were assigned to the class “conifer-logs,” and one element even to the class “conifer-litter.” Only 8 were assigned correctly.

4. Discussion

4.1 Validation images

Based on the results obtained, it was found that the best performance of the model, when assessed during training and validation phases, was achieved, as might be expected, using the Extra-large Instance Segmentation. This model underwent a total of 214 training epochs before reaching its optimal performance (training automatically stopped at 264 epochs, identifying the best performance in the 214th epoch). However, it is important to note that this evaluation was conducted on validation images, which are taken under the same conditions as the training set. In fact, in terms of performance on test images, the model that responded best was the Large Instance Segmentation with 100 training epochs. This is probably since the greater the number of training epochs, the more likely the model is to have the overfitting problem, meaning that the model learns very well how to classify images similar to the training images, struggling to process new images due to low “elasticity” (Srivastava et al., 2014).

Instance Segmentation, therefore, would seem to be an effective system for its intended purpose. Given access to a high-performance remote machine, it was feasible to utilize the “extra-large” model, falling within a maximum training time of ten hours. However, it can be seen in the results that the “large segmentation” model also performed adequately; indeed, with these specific test images it proved to be the most effective model. Looking at the confusion matrices and graphs related to the F1-Score, it appears that the greatest difficulty lies in the correct classification of “broadleaf-logs,” which is almost always confused with “conifer-logs.”

The best performing model on validation images, the Extra-large Instance Segmentation with 214 training epochs, demonstrates strong classification accuracy. Notably, it classifies some conifer-logs as broadleaf-logs but maintains a very high accuracy (low omission error) of broadleaf-logs without confusing them with other classes (true positives). This may be due to several reasons. First, image quality may have affected the correct classification; the presence of shadows, blurring or variations in lighting conditions could have impacted the model’s ability to make correct classifications. Above all, segmenting some images in too much detail, where the model attempted to segment distant or poorly defined objects, potentially reducing classification effectiveness. This effect can be observed in Figure 8, where over-segmentation may have influenced the overall performance of the model.

4.2 Test images

When we use test images, however, most deciduous stems are classified as coniferous (false negatives). This may be due to the fact that the test images are sharper because they were made under more favourable atmospheric conditions (clear day, light high above the horizon), as well as during the growing season. Opposite conditions compared to the images used for training. The false negative rate is the parameter that, for our purpose, is best lowered as much as possible. Classifying a coniferous forest instead of a deciduous forest is a mistake that cannot be afforded to achieve the intended purpose.

Figure 8. Low quality image as it was taken with low light conditions. The model was unable to learn from the color of the trunk, only from the shape. In addition, manual segmentation included plants a little too far apart.

To avoid this, it might be useful to create a dataset with training images that are as varied as possible, both in terms of conditions (seasonality, forest density) and brightness (different weather conditions). In this way, we provide a wide variability of situations, useful for the model to learn as many nuances as possible. However, it is necessary to have a fair number of images for each different situation, otherwise it is possible to incur “underfitting,” a phenomenon whereby the model struggles to learn unambiguous classification patterns due to the large variability of training images (Srivastava et al., 2014).

Looking at the confusion matrices, different elements are often classified as “background.” This means that the model is unable to assign them to any class. This is not a problem for the intended purpose, since it is not necessary for the model to recognize every single element within an image; rather, it is sufficient for those few key elements to be identified in order to determine the prevailing plant formation. For this reason, during the segmentation phase it would be best to highlight close and sharp elements, so that the model learns to recognize these components in detail, with good levels of confidence. For example, in images in Figure 9 below, two situations are depicted in which the number of elements identified by the model is not large, but sufficient to recognize vegetation formations.

Regarding litter classes, the model would seem to be quite reliable, although there are limitations: first, the classes are poorly represented. The other limitation is seasonal variability, especially in sparse forest situations. The vegetation that develops in different seasons can change its appearance considerably. For this reason, identifying litter classes may not be the best way forward; the type of surface fuel could be inferred through the species identified through stem classification.

A different case, however, is the presence of herbaceous litter, which is why it would be interesting to understand how a classification model would respond by employing that class. In this case, the season in which the photos will be taken will prove decisive for proper classification: it will therefore be necessary to include in the dataset both images during the spring-summer period and during the fall-winter season. Linked to the latter element, it would be useful to develop models including the identification of the class “grasslands.” Identification of the latter, in fact, is crucial in determining possible flame front propagation (Rothermel, 1972).

4.3 Image quantity and quality

Currently, the number of participants in image sampling, as well as the total number of photographs collected, is 200+ users and about 10000+ images respectively at the time of the research. Figure 1 shows the number of users and photographs at the time of writing this article, 6 months later, with a definite increase in such numbers. These numbers, which are currently increasing and likely to further increase in the future, support the vision of using AI-based recognition for automatically extracting information from these images, as visual interpretation of thousands of images is not efficient. A validated YOLO model can provide a very valid support to having point-based information on fuels and forest parameters distributed across Europe.

Some research in this sense has been carried out recently. Concerning the detection of certain foliar diseases Zhang et al. (2024) have used photographs for interpretation. To determine the shape, size and contours of flames during fires employing A.I. was the objective in Wang et al. (2024). The results obtained in this latter work in recent literature would seem to be encouraging, returning an accuracy of 91.23 percent (Wang et al., 2024). The use of these modern technologies would seem to be quite reliable. Images can also support remote sensing estimation of biomass in complex scenarios where access to ground sampling is difficult, e.g. mountainous terrain (Kutchartt et al., 2022) or disturbances in forests (Piragnolo et al., 2021, Dalponte et al., 2023) and to validate estimation of bulk densities (Martin-Ducup et al., 2024).

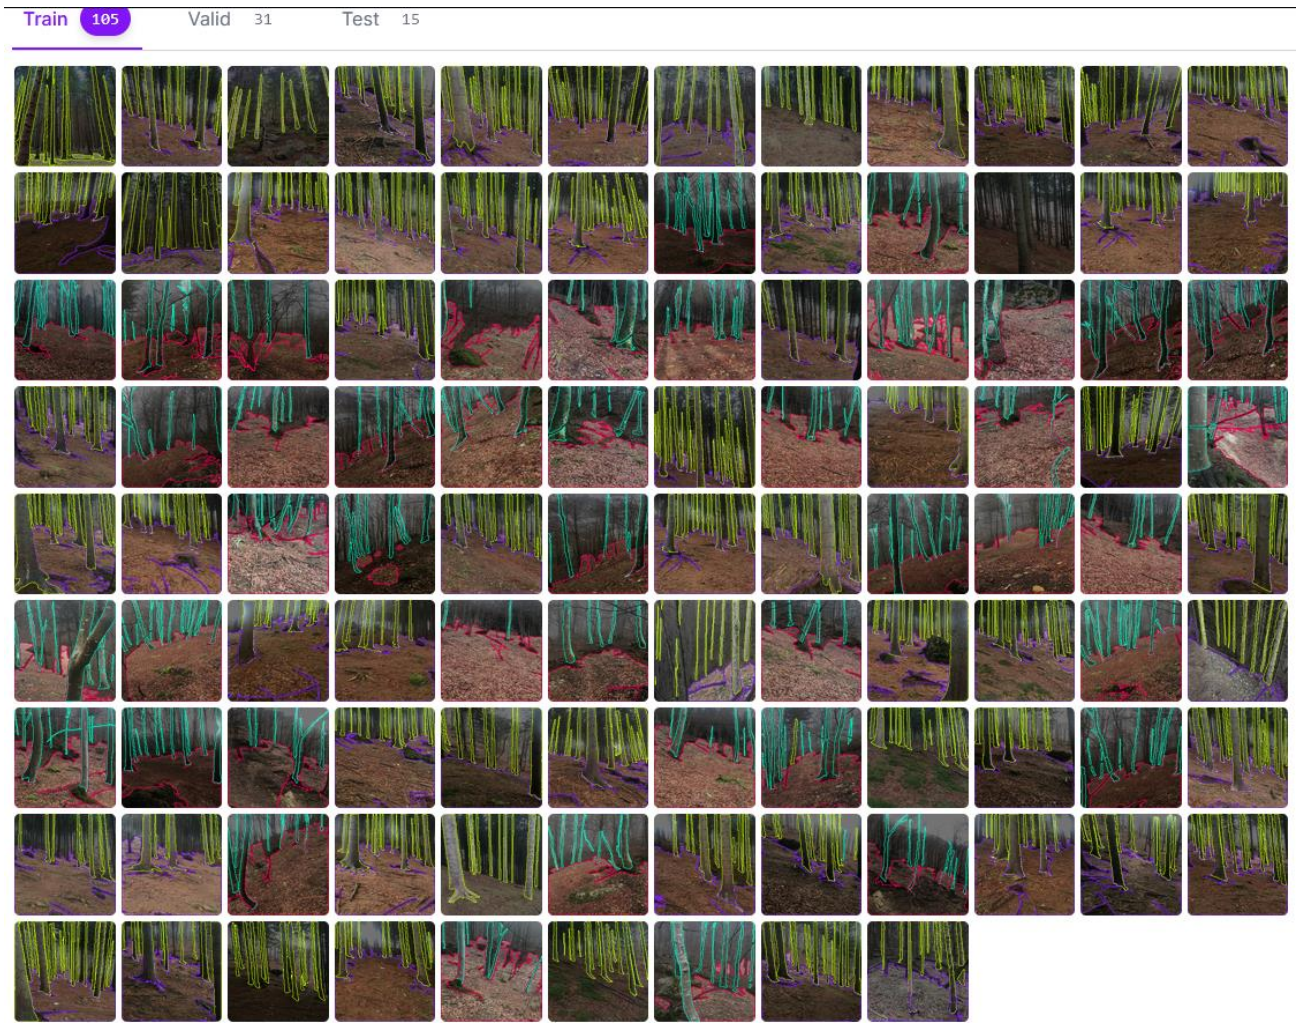

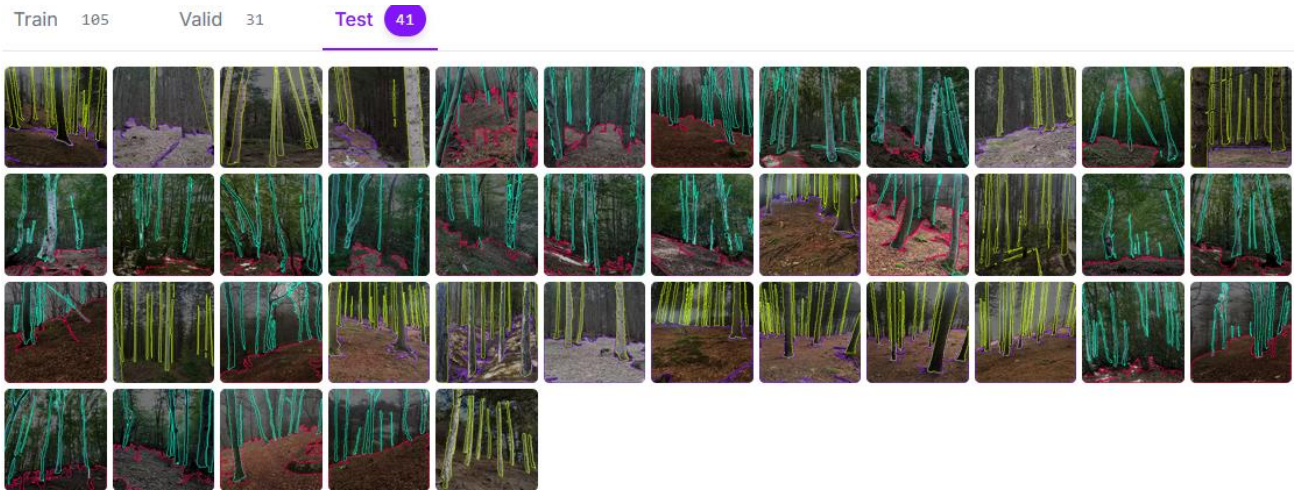

Figure 9. Examples of segmented images in the Roboflow web interface for training (top) and testing (bottom).

5. Conclusions

The work presented here aimed to demonstrate that the use of A.I. can become a useful support for the management and analysis of discrete amounts of images by fulfilling the requirements for effective prevention activities. This, in fact, is one of the tools available to mankind to contain the risk of fire. Using modern technologies in various fields, including forestry, can be a valid support in getting as complete a picture as possible of the current state of forests on a large scale.

In conclusion, it is suggested that forest litter classes be identified based on recognizing the stems of the species that make up the forest, while as for herbaceous areas, these should be determined by training specific models for direct recognition. By being shrewd in taking pictures and segregating the most defined and representative elements, it is probably possible to obtain a good model for the recognition of forest elements.

However, the number of pictures used for the training phase (105) may not be sufficient to obtain a reliable model. In addition, the diversification of situations (brightness, seasonality) is probably an even more important factor in achieving good results.

As a result of these considerations, it would be interesting to implement a classification system of the images segmented by the model based on the abundance and confidence level of the classified objects, so as to assign the correct fuel category with the aim of automatically assigning them to those proposed by Scott and Burgan (2005).

Acknowledgements

This study was supported by the European Union Horizon 2020 Research and Innovation Programme through the project entitled “Innovation technologies & socio-ecological-economic solutions for fire resilient territories in Europe – FIRE-RES”, under grant agreement N°101037419. Dr. Erico Kutchartt was supported by Fondazione Cassa di Risparmio di Padova e Rovigo (CARIPARO). Dr. Francesco Pirotti was supported within the Agritech National Research Center and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR)—MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.4—D.D. 1032 17/06/2022, CN00000022)—task 4.3.1. This manuscript reflects only the authors’ views and opinions; neither the European Union nor the European Commission can be considered responsible for them.

References

Al-masni, M. A., Al-antari, M. A., Park, J.-M., Gi, G., Kim, T.-Y., Rivera, P., Valarezo, E., Choi, M.-T., Han, S.-M., Kim, T.-S. 2018. Simultaneous detection and classification of breast masses in digital mammograms via a deep learning YOLO-based CAD system. Computer Methods and Programs in Biomedicine, 157, 85–94. https://doi.org/10.1016/j.cmpb.2018.01.017

Alom, M. Z., Taha, T. M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M. S., Hasan, M., Van Essen, B.C., Awwal, A.A.S., Asari, V. K. 2019. A State-of-the-Art Survey on Deep Learning Theory and Architectures. Electronics, 8(3), 292. https://doi.org/10.3390/electronics8030292

Chuvieco, E., Pettinari, M. L., Koutsias, N., Forkel, M., Hantson, S., Turco, M. 2021. Human and climate drivers of global biomass burning variability. Science of The Total Environment, 779, 146361. https://doi.org/10.1016/j.scitotenv.2021.146361

Dalponte, M., Cetto, R., Marinelli, D., Andreatta, D., Salvadori, C., Pirotti, F., Frizzera, L., Gianelle, D., 2023. Spectral separability of bark beetle infestation stages: A single-tree time-series analysis using Planet imagery. Ecological Indicators 153. https://doi.org/10.1016/j.ecolind.2023.110349

Hanan Ashraf, A., Imran, M., Qahtani, A. M., Alsufyani, A., Almutiry, O., Mahmood, A., Attique, M., Habib, M. 2022. Weapons Detection for Security and Video Surveillance Using CNN and YOLO-V5s. Computers, Materials & Continua, 70(2), 2761–2775. https://doi.org/10.32604/cmc.2022.018785

Hussain, M. 2023. YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing

and Industrial Defect Detection. Machines, 11(7), 677. https://doi.org/10.3390/machines11070677

Kohavi, R., Provost, F. 1998. Glossary of Terms. Machine Learning, 30(2), 271–274. https://doi.org/10.1023/A:1017181826899

Kutchartt, E., González-Olabarria, J. R., Trasobares, A., de-Miguel, S., Cardil, A., Botequim, B., Vassilev, V., Palaiologou, P., Rogai, M., Pirotti, F. 2023. FIRE-RES Geo-Catch: a mobile application to support reliable fuel mapping at a pan-European scale. iForest – Biogeosciences and Forestry, 16(5), 268-273. https://doi.org/10.3832/ifor4376-016

Kutchartt, E., González-Olabarria, J. R., Aquilué, N., Garcia-Gonzalo, J., Trasobares, A., Botequim, B., Hauglin, M., Palaiologou, P., Vassilev, V., Cardil, A., Navarrete, M. A., Orazio, C., Pirotti, F. 2024. Pan-European fuel map server: An open-geodata portal for supporting fire risk assessment. Geomatica, 76, 100036. https://doi.org/10.1016/j.geomat.2024.100036

Kutchartt, E., Pedron, M., Pirotti, F., 2022. Assessment of Canopy and Ground Height Accuracy from Gedi Lidar Over Steep Mountain Areas. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences V-3–2022, 431–438. https://doi.org/10.5194/isprs-annals-V-3-2022-431-2022

Lang, N., Jetz, W., Schindler, K., Wegner, J. D. 2023. A high-resolution canopy height model of the Earth. Nature Ecology & Evolution, 7(11), 1778–1789. https://doi.org/10.1038/s41559-023-02206-6

Martin-Ducup, O., Dupuy, J.-L., Soma, M., Guerra-Hernandez, J., Marino, E., Fernandes, P.M., Just, A., Corbera, J., Toutchkov, M., Sorribas, C., Bock, J., Piboule, A., Pirotti, F., Pimont, F., 2025. Unlocking the potential of Airborne LiDAR for direct assessment of fuel bulk density and load distributions for wildfire hazard mapping. Agricultural and Forest Meteorology 362, 110341. https://doi.org/10.1016/j.agrformet.2024.110341

Piragnolo, M., Pirotti, F., Zanrosso, C., Lingua, E., Grigolato, S., 2021. Responding to large-scale forest damage in an alpine environment with remote sensing, machine learning, and Web-GIS. Remote Sensing 13. https://doi.org/10.3390/rs13081541

Powers, D. 2008. Evaluation: From Precision, Recall and F-Factor to ROC, Informedness, Markedness & Correlation.

Rothermel, R. 1972. A mathematical model for predicting fire spread in wildland fuels. USDA Forest Service Research Paper INT-115.

Scott, J. H., Burgan R. E. 2005. Standard Fire Behavior Fuel Models: A Comprehensive Set for Use with Rothermel’s Surface Fire Spread Model. Ft. Collins, CO: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station https://doi.org/10.2737/RMRS-GTR-153.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research, 15(56), 1929-1958. http://jmlr.org/papers/v15/srivastava14a.html

Talaat, F. M., ZainEldin, H. 2023. An improved fire detection approach based on YOLO-v8 for smart cities. Neural Computing and Applications, 35(28), 20939–20954. https://doi.org/10.1007/s00521-023-08809-1

Terven, J., Córdova-Esparza, D.-M., Romero-González, J.-A. 2023. A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS. Machine Learning and Knowledge Extraction, 5(4), 1680–1716. https://doi.org/10.3390/make5040083

Turner, M. G., Braziunas, K. H., Hansen, W. D., Harvey, B. J. 2019. Short-interval severe fire erodes the resilience of subalpine lodgepole pine forests. Proceedings of the National Academy of Sciences, 116(23), 11319–11328. https://doi.org/10.1073/pnas.1902841116

Wang, G., Wang, F., Zhou, H., Lin, H. 2024. Fire in Focus: Advancing Wildfire Image Segmentation by Focusing on Fire Edges. Forests, 15(1), 217. https://doi.org/10.3390/f15010217

Xanthopoulos, G., Caballero, D., Galante, M., Alexandrian, D., Rigolot, E., Marzano, R. 2006. Fuels Management-How to Measure Success: Conference Proceedings. USDA Forest Service Proceedings RMRS-P-41. 2006.

Yaseen, M. 2024. What is YOLOv8: An In-Depth Exploration of the Internal Features of the Next-Generation Object Detector. arXiv. Retrieved from http://arxiv.org/abs/2408.15857

YOLO Performance Metrics - Ultralytics YOLO Docs. 2023. Retrieved September 7, 2024, from https://docs.ultralytics.com/guides/yolo-performance-metrics/

Zhang, M., Yuan, C., Liu, Q., Liu, H., Qiu, X., Zhao, M. 2024. Detection of Mulberry Leaf Diseases in Natural Environments Based on Improved YOLOv8. Forests, 15(7), 1188. https://doi.org/10.3390/f15071188

Zhi-Hua, Z. 2021. Machine Learning. Springer Nature Singapore, 458 pp. https://doi.org/10.1007/978-981-15-1967-3