RESEARCH ARTICLE
doi:10.59236/geomorphica.v2i1.47

GEOMORPHICA

Automated instream large wood detection and load estimation applying machine learning on high-resolution aerial imagery

Janbert Aarnink [*1], **Gabriele Consoli** [1,2], **Bryce Finch** [1], **Marc O'Callaghan** [1], **Ivan Pascal** [1], **Samuel Wiesmann** [3], **Virginia Ruiz-Villanueva** [1,2]

[1]Institute of Earth Surface Dynamics, Université de Lausanne, Lausanne, Switzerland, [2]Institute of Geography, University of Bern, Bern, Switzerland, [3]Swiss National Park, Zernez, Switzerland

Author contributions: *Conceptualization*: J. Aarnink, G. Consoli, V. Ruiz-Villanueva. *Methodology*: J. Aarnink, G. Consoli, V. Ruiz-Villanueva. *Software*: J. Aarnink. *Formal Analysis*: J. Aarnink. *Resources*: V. Ruiz-Villanueva, S. Wiesmann, I. Pascal, B. Finch, M. O'Callaghan, I. Pascal. *Writing - original draft*: J. Aarnink, G. Consoli, V. Ruiz-Villanueva. *Visualization*: J. Aarnink. *Supervision*: V. Ruiz-Villanueva. *Project administration*: V. Ruiz-Villanueva. *Funding acquisition*: V. Ruiz-Villanueva.

Abstract Together with the flow and sediment, instream large wood (LW; fallen trees, branches, and roots within river corridors) sustains the physical and ecological integrity of rivers. Still, how much wood rivers store remains largely unknown. Quantifying LW is crucial for river management, restoration, and flood mitigation, but remains challenging due to its high spatio-temporal variability. Field surveys are labour-intensive, while manual mapping of aerial imagery is also time-consuming and constrained by the observer's experience. This study presents an automated approach to detect and measure stationary LW using high-resolution imagery and convolutional neural networks. Detection locates wood with bounding boxes, while segmentation outlines its size. Two models (YOLOv10 for detection and YOLOv8 for segmentation) were trained with data from eight rivers in the Alps and Andes, and tested on independent data. The method detected 6.60 m^3 of wood, while 7.36 m^3 were estimated from field surveys, identifying up to 97% of large pieces and 66% of all pieces at 0.3 confidence. Segmentation achieved a mean Average Precision of 70%. Although diameters were underestimated and jam sizes differed from observations, spatial distributions were reliably captured. This scalable approach overcomes key survey limitations, offering an efficient tool for monitoring LW.

Non-technical summary Wood in rivers plays a crucial role in maintaining healthy ecosystems by influencing sediment movement and providing habitats for aquatic life. Measuring how much wood is stored in a river is essential for restoring habitats, managing rivers, and reducing flood risks. However, counting and measuring wood in the field is slow and can only cover small areas, while analysing aerial images manually takes a lot of time and effort. We developed a new automated method using high-resolution drone images and artificial intelligence to solve this. We trained computer models to detect and measure wood in rivers by analysing images from eight rivers in the Swiss Alps and Argentinean Andes. The models were tested on a separate river to ensure accuracy. Our method identified over 97% of the largest wood pieces and accurately estimated wood volumes, comparable to field measurements. While there were inaccuracies with certain wood shapes, the results were promising for capturing the overall distribution of wood in rivers. This automated method makes it possible to monitor wood in rivers quickly and on a larger scale. It offers a useful tool for scientists, conservationists, and river managers to better understand and manage river ecosystems.

Keywords Remote Sensing, Aerial Imagery, Wood Quantification, Machine Learning

Received:
December 13, 2024
Accepted:
August 26, 2025
Published:
October 6, 2025

Editor-in-Chief:
Alice Lefebvre
Associate Editors:
Roberto Fernández
Serina Diniega
Managing Editor:
Xun Yang
Production Editor:
Nathaniel Klema

Reviewers:
Daniel Scott
Alejandro Tenorio
Anonymous Reviewer 3

1 Introduction

Instream large wood (LW), which refers to fallen trees, branches and rootwads lying within the river corridor, is a crucial element in river ecosystems (Verdonschot and Verdonschot, 2023) that supports greater ecosystem complexity (Gurnell et al., 2005) and promotes biodiversity, increasing resilience to disturbances like floods and droughts (Wohl et al., 2019). Quantifying the volume of LW in rivers is important for effective river management, habitat restoration, and flood mitigation strategies. However, LW storage quantification is challenging due to its significant spatial and temporal variability, which is influenced by multiple factors (Wohl et al., 2017, 2019; Iroumé et al., 2020). Several studies have shown that LW storage is controlled by basin-scale variables like bioclimatic regime, precipitation, land use and forest characteristics and reach-scale variables, like river morphology and its spatial heterogeneity and wood characteristics and recruitment processes (Scott and Wohl, 2018; Wohl et al., 2018). Effective management must account for the fluctuations in LW storage and balance them with flood mitigation goals (Iroumé et al., 2020). The natural wood regime, which includes LW supply, transport and storage, varies spatially and temporally, necessitating region-specific

*Corresponding author: janbert.aarnink@unil.ch

management strategies that consider local dynamics (Wohl et al., 2019). While traditional river management has commonly resulted in reduced LW storage (Wohl et al., 2017), the amount of LW stored in rivers is largely unknown. This study focused on these challenges and aimed to advance the detection and quantification of LW volumes in river environments.

Traditionally, field surveys have been employed to manually measure LW, its size and orientation (Wohl et al., 2010; Eaton et al., 2012; Merten et al., 2010; Galia et al., 2018). A survey team walks through the river's section of interest and tag all pieces of large wood after gathering relevant information. The length of the piece, as well as the diametre, usually measured at one or two locations chosen to best represent its overall size, using tape measures or range finders, are measured, and information on the orientation of the pieces with respect to the flow direction of the river and the decay was stored, as well as information on which pieces were in a wood jam with other pieces. These common field survey techniques have been described in detail (Máčka et al., 2011). These methods are labour-intensive, and outcomes are often temporally and spatially limited. Recent advances in remote sensing technology, the increased availability of satellite data and particularly the use of Uncrewed Aerial Vehicles (UAVs) (see Table S2) have changed how we monitor river environments (Sendrowski and Wohl, 2021; Piégay et al., 2019), offering high-resolution and cost-effective alternatives for large-scale assessments (Manfreda et al., 2018). UAV-derived RGB orthomosaics provide detailed spatial data that enable precise quantification of forested areas and fluvial features. However, the complexity and variability of natural environments present challenges in accurately identifying and measuring LW from these datasets (Sanhueza et al., 2022; Jutras-Perreault et al., 2023; Hess et al., 2024). Table S1 expands on a previous review (Buscombe et al., 2024; Sendrowski and Wohl, 2021) and shows a list of studies using aerial imagery to detect large wood in rivers. These studies demonstrate the potential of UAV imagery in capturing detailed and accurate data on LW. UAVs have also been used to explore LW recruitment and retention in braided rivers and dynamic river systems impacted by volcanic activity (Bertoldi et al., 2013; Ulloa et al., 2015). These studies emphasize the importance of UAVs for capturing spatial variability and wood deposition patterns, particularly in response to natural disturbances. One of the most critical challenges for detecting LW from UAVs is the interference from riparian vegetation and canopy cover, which can obscure wood (Hess et al., 2024), particularly in forested environments, leading to significant detection and measurement errors. Additionally, the similarity in colour between sediments and wood can make it difficult to differentiate them in orthomosaics (Sanhueza et al., 2022). The porosity in wood jams and background materials like rocks and organic debris further complicates volume calculations (Wohl et al., 2019). While UAVs provide a cost-effective means to estimate LW volume and distribution with high accuracy (Sanhueza et al., 2022, 2019), the process of manually identifying, measuring, and analysing LW from UAV imagery can still be time-consuming and labour-intensive. Ground truthing involves validating UAV data with field measurements, which adds to the workload. The manual annotation of images and subsequent volume calculations necessitate significant human resources, making the process less efficient and prone to human errors (Hess et al., 2024; Sanhueza et al., 2019). Also, the detection and identification of LW orientations have yet to be automated.

Many studies summarized in Table S1 effectively detected and monitored LW but did not quantify wood volume automatically, still requiring manual calculations. However, quantifying the LW volume is critical for understanding fluvial processes, ecosystem dynamics, and flood risk management (Wohl et al., 2019; Ruiz-Villanueva et al., 2014, 2016). Estimating LW volume accurately remains a major challenge due to the complexity of wood jams, variability in orientation, and the limitations of existing automatic detection methods. To address these challenges, scientists have been conducting manual surveys. Given these challenges, there is a clear need to automate the analysis of UAV imagery to enhance efficiency. Automation can significantly reduce the human effort required for such studies, allowing for more rapid, standardized, repeatable, and reliable data processing (Tassielli et al., 2024). Integrating remote sensing techniques can facilitate long-term studies of wood dynamics. Airborne LiDAR Scanner (ALS), which creates 3D models of topography and vegetation, has shown significant potential in detecting and mapping instream wood on a large scale. Another advancement was the development of an automated framework that uses ALS data to map LW in small coastal streams (Kuiper et al., 2023). This method, which integrates point cloud filtering and skeletonization techniques, demonstrated detection accuracy ranging from 37% to 87%, depending on factors such as canopy cover and wood submergence. Similarly, Structure from Motion (SfM) photogrammetry has been used to assess wood accumulations in the field, offering a new approach to detecting LW, even though they are limited in spatial scale and focus on a limited amount of structures (Ortega-Terol et al., 2014; Spreitzer et al., 2019; Sanhueza et al., 2019). Similarly, several studies (Haschenburger and Rice, 2004; Lassettre et al., 2008; Tamminga et al., 2015) relied on manual or semi-automated approaches that involved significant manual validation steps, which limited their scalability and automation potential. An early attempt showed that computer vision techniques can automatically measure wood volumes (Correia et al., 1993). This approach could further enhance the accuracy of LW detection and quantification by minimizing manual errors and increasing efficiency. The use of photo analytical methods was also explored with a focus on stacked wood, which has implications for solid wood content estimation (Pásztory and Polgár, 2016). The methods achieved higher accuracy than traditional techniques, emphasizing the role of image processing in wood detection. Also, even though not created explicitly for image processing, the XGBoost tool effectively classified LW in aerial imagery (Liang et al., 2022).

By leveraging advanced machine learning techniques, such as convolutional neural networks (CNNs), it is possible to develop models that can automatically detect, segment, and measure LW in UAV imagery. This approach not only streamlines the workflow but also minimizes the potential for human error, ensuring more reliable and reproducible results (Manfreda et al., 2018; Kamilaris and Prenafeta-Boldú, 2018). CNNs have shown promise in various object detection tasks, offering a potential solution to enhance the automated detection and measurement of LW (Kamilaris and Prenafeta-Boldú, 2018; Schmidhuber, 2015; Rouge, 2022; Reymond, 2022; Buscombe et al., 2024). The prior research highlights how automated deep-learning-based image segmentation and object detection can accurately identify and measure wood features in high-resolution orthoimages. Most of these studies, however, used the same data for training and testing. Therefore, the method might overfit and be accurate for one specific site instead of being globally applicable. Furthermore, although effective in wood detection and storage estimation, Sendrowski and Wohl (2021) and Grimmer et al. (2025) relied on the use of multispectral data, which is not a common tool that is available with all aerial observation devices.

While CNNs have been successfully used for wood detection, the key challenge remains accurately estimating wood volume from UAV imagery. This study aims to bridge that gap by leveraging CNNs to not only detect LW but also quantify its volume, providing a scalable solution for river monitoring. We present a novel application of CNNs to analyse UAV orthomosaics. By leveraging CNNs' ability to learn and generalize from high-dimensional RGB image data, we aim to improve the accuracy and efficiency of LW monitoring. Our goal is to provide scientists and river managers with a tool to track wood volume over time and assess whether critical amounts of stored wood could pose a risk to infrastructure during floods. Automating the detection and measurement process can also facilitate large-scale studies and continuous monitoring, providing valuable insights into temporal changes in LW storage and its spatial distribution. It can also turn data gathered for different purposes into a new source of wood storage estimation. Understanding these changes is crucial for assessing how LW dynamics influence river characteristics like sediment transport, habitat formation, and flood risks, ultimately aiding in effective river management. This can be particularly beneficial for river management and restoration efforts, where timely and accurate data are important for informed decision-making (Hess et al., 2024; Sendrowski and Wohl, 2021). By implementing CNNs for volume estimation, this study provides a scalable and repeatable method for quantifying LW storage in rivers.

2 Methods

The goal of this study is to detect LW in UAV-derived orthomosaics and to estimate its volume. To achieve this, our approach combines two sequential processes. At first, object detection using a convolutional neural network (CNN) is performed to locate wood pieces and delineate bounding boxes. Thereafter, a wood segmentation model estimates the wood surface area within these bounding boxes, which is then used to infer LW diametre and volume.

2.1 Data Acquisition

High-resolution data from 7 rivers in the Swiss Alps and the Argentinean Andes, along with an additional satellite image of the Hoh River in the United States, were used to build a diverse training dataset. A separate dataset from the Avançon de Nant in the Swiss Alps was exclusively used for independent testing to evaluate model performance. Table S2 shows the different locations at which the data was collected. Furthermore, a series of satellite images in which LW was visible at the Hoh River around coordinate (47.806721, -124.053095) WGS84 in the United States was taken from Google Maps and added to the training database. In machine learning, having training data from diverse sources enhances the model's ability to generalize across different conditions, improving robustness (Goodfellow et al., 2016). The UAV data were converted into high-resolution georectified orthomosaics, the primary data source for this analysis. The orthomosaics were cropped to include only the river and the immediate vicinity, excluding surrounding forests or other non-relevant areas. The selected rivers offer a range of environmental conditions. The Swiss rivers are high-elevation mountain streams with narrower channels and predominantly coniferous riparian forests. The Argentinean rivers are wider, and characterized by deciduous vegetation. While overall LW storage was comparable, slightly higher volumes were observed in the Argentinean rivers. The inclusion of the Hoh River in the training dataset added valuable diversity, particularly in terms of larger wood pieces. Despite being derived from satellite imagery, the resolution resulted in wood pieces occupying similar pixel dimensions to the UAV-acquired images. In addition to detecting LW and estimating the volume of the single pieces and at the river reach scale, the LW piece orientation was also identified. To do so, a shapefile (.shp) containing a line that follows the river's thalweg is required, which, for this study, we visually drew based on the orthomosaic. The main line of the shapefile needs to run from upstream to downstream and follow the river's thalweg. Secondary lines, like tributaries, can be included in the shapefile. The algorithm determined the orientation with respect to the closest point on any of the lines.

2.2 Model training and validation

Two convolutional neural networks were trained using different versions of the YOLO model (Varghese and M., 2024) from Ultralytics (Jocher et al., 2023). The YOLO models used in this study are open-source and freely available under the GPL-3.0 license, making them accessible for both academic and commercial applications. The first model used the 10th version of the YOLO architecture and was designed to detect LW in orthomosaics by locating and drawing bounding boxes

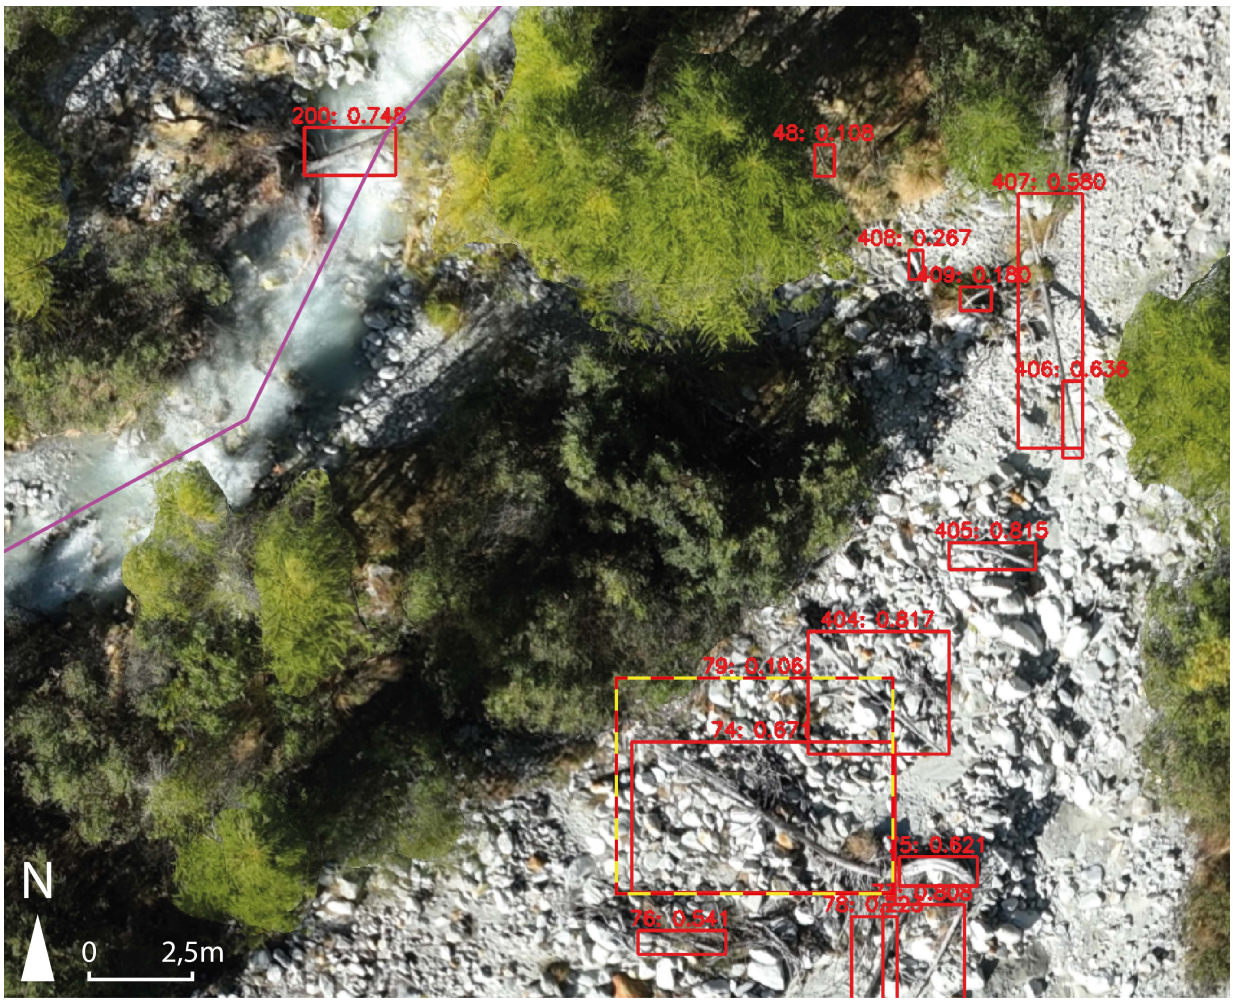

Figure 1 Example of the detection on an orthomosaic at the Navizence (Canton Valais, Switzerland), on a section of the river that was not used as training data. The red bounding boxes indicate the detection of a piece of instream wood. Its ID number is indicated above (left), and so is the confidence score of the CNN (right). The purple line shows the thalweg of the river. The bounding box indicated with a yellow-red dashed line is a false detection and has a confidence score of 10.6%.

around objects (detection). Segmentation, used in the second model, outlines the exact pixels belonging to each detected object. YOLO was selected because we have a large, labelled image database specifically for LW detection in rivers, which allows us to train and fine-tune the model effectively. Additionally, YOLO is a well-established object detection framework with known performance in both research and industry (Sharma et al., 2021). Its flexibility enables easy adaptation for UAV-derived detection purposes. The model comes with different complexities that influence its performance and processing time. Larger models have more links and nodes and can handle more complex data. For this study, the small (17 MB), medium (34 MB), balanced (42 MB) and large (52 MB) versions of the detection model were trained, all with 960 pixels (960p) input image size. Training time ranged from 2 to 6 hours, depending on the model size, with larger models requiring longer training but offering marginal accuracy gains, although this can vary based on the hardware used. Different colour spaces, as proposed by Liang et al. (2022), who found that detections in the YCbCr colour space were more effective, were also considered when training the detection model. Subsequently, the best-performing model was chosen for the detection step in the method. The gathered datasets were cropped into tiles of 960x960 pixels. Twenty percent of the tiles were randomly copied, after which the data was manually labelled by drawing bounding boxes around the visible pieces of LW. The training dataset consisted of 226 labelled images extracted from the orthomosaics of the 8 rivers, with 205 images used for training and 21 for validation. These were randomly selected from a larger set of cropped tiles to ensure diverse training input. These images include a range of lighting conditions and shadow scenarios, such as shaded sections under canopy and varying illumination conditions, to help the model generalise across real-world variability.

The most recent YOLO model (v10) did not yet include a feature for segmentation. Therefore, for the second step, the 8th version of the YOLO model (Jocher et al., 2023) was trained to segment wood pixels within detected bounding boxes. The wood segmentation algorithm was developed to estimate the wood surface area within each detection bounding box, from which the diametre and volume of the wood pieces were calculated. To optimize the model's performance, we tested the 'medium,' 'large,' and 'extra-large' versions of the YOLOv8 model at 160 (160p) and 320 pixels (320p) input image resolutions. For this dataset, labelled boxes from

the wood detection training data were cropped from the orthomosaics described in Table S2. On these cropped boxes, polygons were drawn around the pixels that represent wood. A total of 266 cropped bounding boxes were labelled, with 247 used for training and the rest for validation.

The images were manually annotated by three individuals with experience in surveying LW in rivers, and all labels were subsequently reviewed and verified by the first author to ensure consistency and accuracy. Although inter-annotator agreement was not formally assessed, this multi-review process aimed to minimise inconsistencies in the training data. The full set of annotated images used for training, validation and testing is publicly available (see data availability section) to support reproducibility and verification.

2.3 Detection and segmentation process

The proposed method consists of wood detection and wood segmentation steps. During the detection steps, the orthomosaic was divided into 480x480 pixel tiles with an added border layer of 240 pixels (called padding) on all sides, resulting in 960x960 pixel tiles. The wood detections were performed on these tiles, retaining only bounding boxes with their centres in the middle 480x480 region. Detected bounding boxes were stored with their location, size, and confidence score in the coordinate system of the orthomosaic. Qualitative analysis where we visually inspected the detection data has indicated that the algorithm can detect the same piece of wood multiple times. Therefore, we implemented a safeguard approach to address this by retaining only one detection if the Intersection over Union (IoU) between overlapping bounding boxes exceeds 0.7. Nonetheless, multiple detections may occur frequently for larger pieces.

The detections were then filtered by confidence score at intervals of 10%. Per confidence score, which functions as a threshold, separate text files were created for each detection in the form of bounding boxes, which are boxes that envelope the detected log. The file includes Detection ID, the confidence score, the centre of the X and Y coordinates, the minimum X and Y coordinates of the bounding box, the maximum X and Y coordinates of the bounding box, and the diagonal length of the bounding box (which represent the length of the log). Also, per confidence score, orthomosaics include only the bounding boxes of the detections with a confidence score larger than the threshold. The orthomosaics, of which Figures 1 & 2 show examples, also include the detection ID and confidence score per bounding box and can aid in the decision of what confidence threshold to use for the reduction of false positives. For example, in the case shown in Figure 1, a threshold of 20%, instead of 10% would be advised because detection number 79 (indicated with a yellow-red dashed bounding box) is a clear false positive with a confidence score of 10.6%.

Each detection was cropped once a confidence level was chosen (which was 30% for all orthomosaics except for the Navisence, where we used a 20% confidence level based on visual inspection), and the wood segmentation model was applied. The segmentation model created a mask polygon, a polygon line that borders all the pixels segmented as LW. The number of segmented pixels was determined by counting the segmented pixels included in the mask polygon. The length of the wood piece was determined by measuring the diagonal of the bounding box. This approach assumes the LW piece lies in a straight line, which is not always the case for bent or irregularly shaped wood. Such simplifications can introduce inaccuracies in length estimates. The total surface area was calculated by multiplying the pixel count by the size of each pixel. To counteract pixels of overlapping bounding boxes (for example, in the case of multiple LW pieces within a jam) being counted multiple times, the total surface area was corrected based on an analysis of the number of unique pixels versus the number of total pixels per jam. The diametre was calculated by dividing the surface area by the length, assuming the wood piece is cylindrical. This approach simplifies the geometry of the wood piece, disregarding irregularities, and assumes a uniform, circular cross-section along its entire length. The number of unique pixels segmented per jam was divided by the total pixels detected as wood for that jam. This ratio adjusts the detected pixels for each piece in the jam, ensuring accurate volume calculation. The volume was then calculated using this diametre and length. Edge detection filters were applied to analyse the orientation of each piece of LW within the bounding box using the Sobel operator, which computes intensity gradients to detect edges and has been adapted here for diagonal edge detection using customized Sobel filters (Gonzalez and Faisal, 2019). The function estimates the predominant orientation of elongated objects like wood by applying these filters in two diagonal orientations (North-West South-East and North-East South-West). The angle of the LW relative to the closest point on the river line shapefile (i.e., thalweg) was calculated, as well as the longitudinal distance along the thalweg line from upstream to downstream.

2.4 Performance metrics

Several performance metrics were used to evaluate the effectiveness and accuracy of the trained machine-learning models. Precision, recall, mean Average Precision (both mAP50 and mAP50-95) are the most common metrics used to evaluate object detection and segmentation models and often refer to the concept of Intersection over Union (IoU) (Lin et al., 2014). IoU, also known as Jaccard's Index, measures the overlap between the predicted bounding box inferred by the model and the ground truth bounding box from the validation dataset. Mathematically, IoU is defined as the ratio of the area of overlap between the predicted and ground truth bounding boxes to the area of their overlap. A higher IoU score indicates better alignment between the prediction and the actual object, with a value of 1 representing perfect overlap. Subsequently, precision refers to the proportion of correctly detected objects out of all detections made, while recall measures the proportion of actual objects correctly detected by the model. Precision

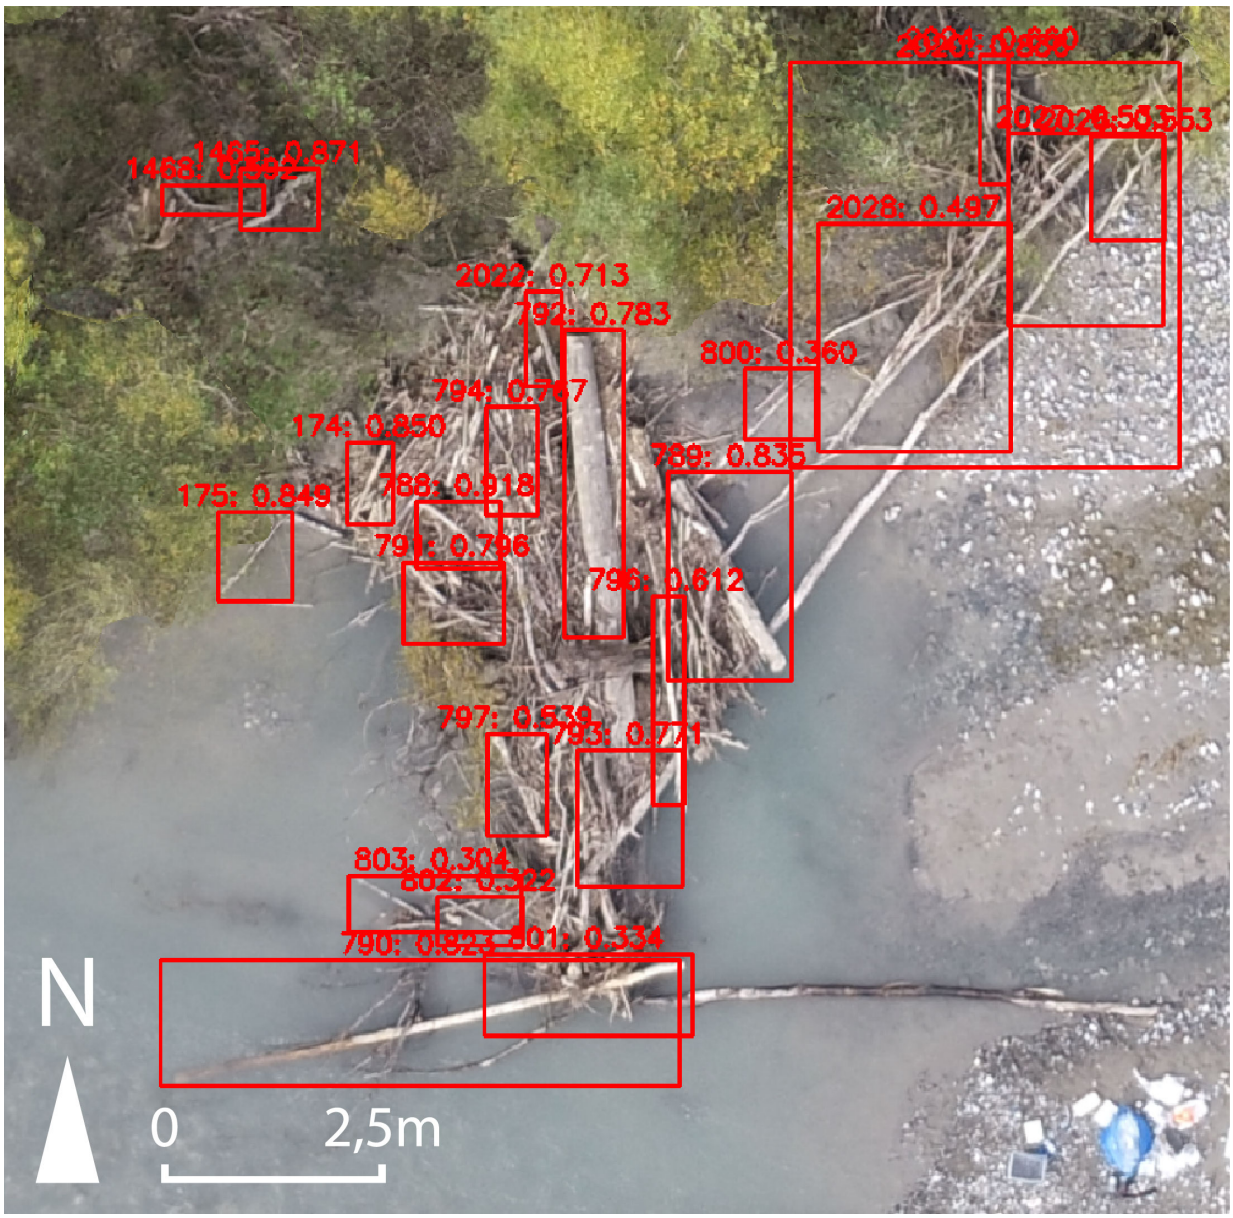

Figure 2 Example of the detection of LW accumulated in a jam, on an orthomosaic from the Kander River (Canton Bern, Switzerland), on a section of the river that was not used as training data. The red bounding boxes indicate the detection of a piece of LW. Its ID number is indicated above (left), and so is the confidence score of the CNN (right).

specifically assesses how well the model distinguishes actual objects from false positives, providing insights into the model's ability to make positive predictions that are indeed accurate. While recall assesses the model's completeness in identifying objects of interest, a high recall score indicates that the model effectively identifies most of the relevant objects in the data. Average Precision (AP) and mean Average Precision (mAP) incorporate a trade-off between precision and recall and are calculated for a given threshold. For example, the mAP50 (mean Average Precision at IoU threshold 0.50) is the average precision when detection is correct if the Intersection over Union (IoU) between the predicted and ground truth object is at least 0.50. The mAP50–95 provides a more comprehensive evaluation by averaging precision across IoU thresholds from 0.50 to 0.95 in steps of 0.05, reflecting the model's performance at varying levels of localization accuracy (i.e., thresholds).

In our study, the primary objective was not to achieve a perfect fit of bounding boxes but to maximize detection regarding the number of LW pieces and their estimated size and volume. Given this focus, relying on the mean Average Precision (mAP) alone as the key performance metric may not fully capture the model's effectiveness in meeting these objectives. Therefore, we prioritized the recall over the mAP.

2.5 Test site: Avançon de Nant

To test the method's performance in a dataset not used for training or validation, we applied our models to a test dataset from the Avançon de Nant stream in the Vallon de Nant (VDN) region in the Swiss Alps. This separation ensures that the testing evaluates the model's performance on completely unseen data. The VDN is a protected area in the Alps of Canton Vaud, in Switzerland's upper Rhône River basin. The drainage area is around 13 km^2, and the elevation ranges from the Pont

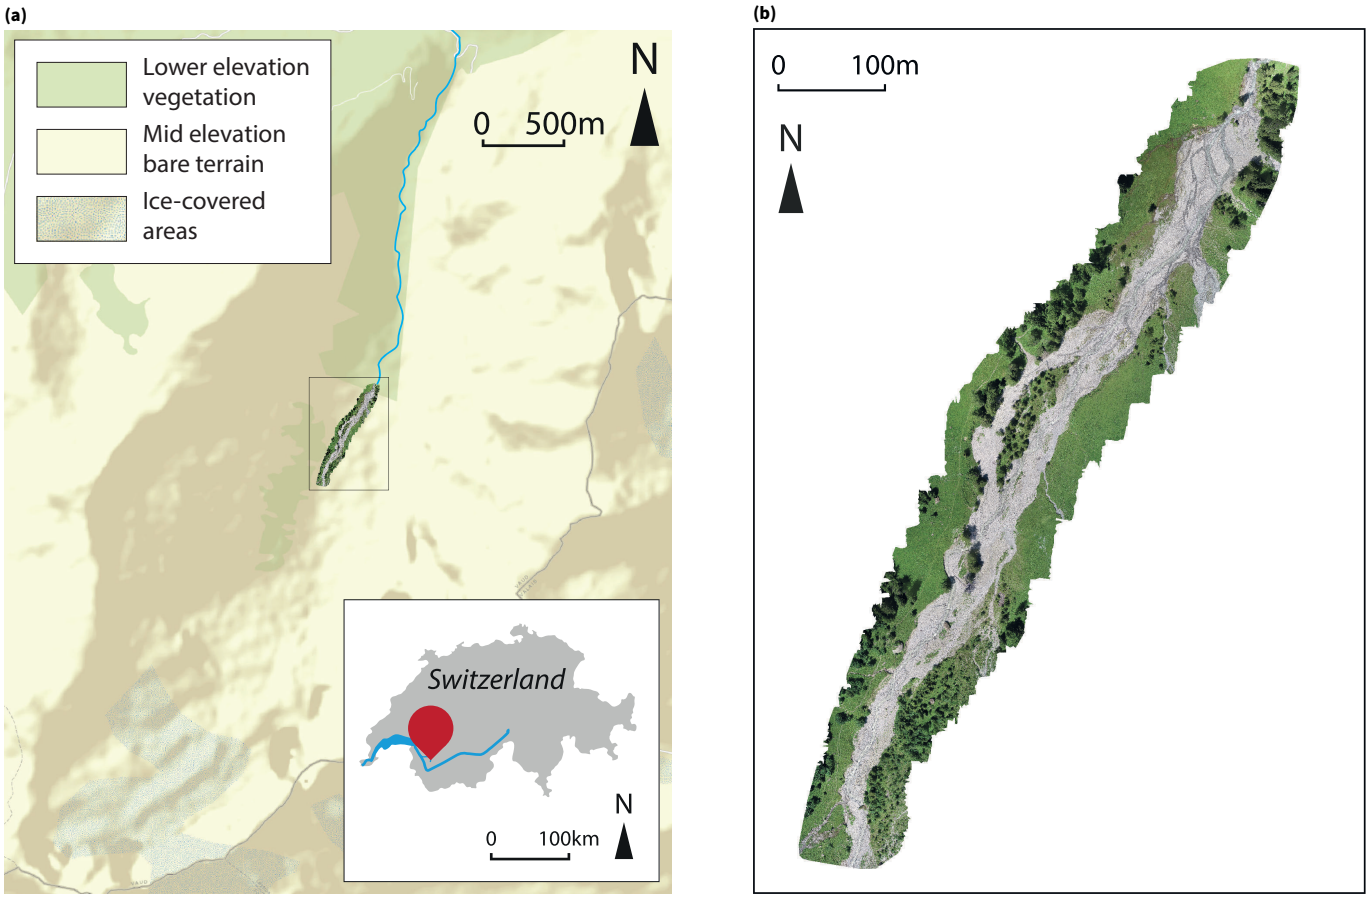

Figure 3 (**a**) The catchment of the Avançon de Nant with the acquired orthomosaic overlaid in the centre and (**b**) a closeup of the orthomosaic. Background map source: Esri Terrain Map.

de Nant at 1253 to 3100 metres above sea level. With a partially glacial and snowmelt dominated regime, the Avançon de Nant stream (see Figure 3) runs through a relatively steep, partially forested valley (average gradient is around 17% according to Antoniazza et al., 2022), varying in morphology from a braided multi-thread reach in the upper part to a more confined step-pool lower reach. For this research, we focused on the upper forested multi-thread channel section of the valley between 1470 and 1530 metres in altitude, covering about 1 kilometre of the river.

RGB aerial imagery of a 950-metre section of the Avançon de Nant River was captured using a DJI Phantom 4 Pro drone on the 25th of July, 2024. The imagery was processed using WebODM photogrammetry software, producing an orthomosaic with a spatial resolution of 2 cm per pixel.

As part of an extensive fieldwork campaign in spring-summer 2024, the location, length, diametre, and orientation of all LW pieces in the river were surveyed. The aerial survey was conducted approximately one month after this field survey (see Figure 4), resulting in an orthomosaic of the same river section. This overlap provides an opportunity to directly compare the algorithm's detection performance against field-measured values for each piece of LW. However, due to the wet summer conditions and frequent high flows, some wood pieces were transported and moved from their initial location within the month between the fieldwork and the aerial survey. Figure 4 shows that 3 out of 4 of the largest floods of the season took place in this period. The limited GPS signal during the survey made it challenging to match individual LW pieces between the survey data and the orthomosaic in locations with multiple LW. Consequently, only 29 individual wood pieces visible in the orthomosaic could be reliably matched to the surveyed logs. These were selected for precise, one-to-one comparison with the model's results. The remaining surveyed LW could not be confidently located in the imagery due to occlusions, spatial shifts, or ambiguous matching.

Several analyses were performed to evaluate the method's effectiveness and accuracy for these individual LW pieces. LW lengths, diametres, and volumes were compared between the survey data and the algorithm's detections, along with assessing the accuracy of orientation detection. In addition, the spatial distribution of the volume of wood for both the ground truth and the detected data was compared by analysing the measured and obtained wood volumes in the longitudinal direction. The orientation of detected LW was grouped into three classes: 0–22.5 degrees (parallel), 22.5–67.5 degrees (oblique), and 67.5–90 degrees (perpendicular). The frequency distributions of the volume and length of wood pieces were also compared, as well as volumes per wood jam and distribution of the orientations. While the detailed testing was conducted on individual LW pieces, a broader comparison of wood jam

sizes was also performed. The total volume of wood in jams was assessed by summing the estimated volumes of detected LW grouped into jams, allowing for a general evaluation of jam sizes despite the challenges in direct validation.

To compare the distribution of the length, diametre, and volume of LW pieces between field survey measurements and the outputs of the detection algorithm, a Kernel Density Estimation (KDE) was performed. KDE is a non-parametric method for estimating the probability density function of a continuous variable. It visualizes the underlying distribution of data without assuming a specific distribution shape. KDE was used to smooth the distribution of LW characteristics from the survey and algorithmic data. The density plots provide a visual comparison, highlighting areas where the algorithm's estimates closely resemble the survey measurements. The method facilitates a clear and intuitive comparison between the two datasets, offering insights into their respective distributions.

A direct comparison was conducted between survey measurements and algorithm estimations to assess the accuracy of the wood detection algorithm. Scatter plots were used to visualize the relationship between observed and estimated values for key parametres, including LW length, diametre, and volume. The first set of comparisons was performed between survey and algorithm-derived lengths, diametres, and volumes for a subset of 29 LW pieces. The second one compares length versus diametre for all LW, using survey data and algorithm detection. These visualizations allowed for an evaluation of systematic biases, over- or underestimation tendencies, and the overall performance of the detection method in capturing LW dimensions.

3 Results

3.1 Wood detection performance

To show the wood detection performance, we compared the validation data (randomly cut out orthomosaic sections of 480x480 pixels with manually labelled bounding boxes) of the training process to the detections made by the model. This data included both individual pieces of wood and pieces in wood jams. From the tested models, the small YOLOv10 model yielded the best results. When analysing detection rates of different size thresholds (regardless of the LW size), detection rates vary across wood size classes (i.e., bounding box areas divided into five equal-sized bins based on pixel ranges). The model identified 66% of all LW pieces and performed particularly well for larger sizes (the 10% largest pieces in the dataset, here between 3.4 and 10.2 metres in length), achieving a detection rate of 97% (see Figure 5). The wood detection model achieved a mean Average Precision of 46% at an intersection over union (IoU) of 0.5, a recall of over 50% and a mAP-95 approaching 22% (Figure 6). Qualitative analysis showed that the same piece was often detected multiple times (see detection numbers 406 and 407 in Figure 1).

The size classes in Figure 5 correspond to bounding box pixel areas and were divided into five equal-sized bins to analyse performance across LW sizes. For reference, these pixel areas approximately correspond to LW pieces with lengths ranging from 43 cm to 10.2 m, assuming an average pixel size of 2 cm (see Table S2). While the class divisions are fixed for this analysis, we acknowledge that they depend on image resolution. For example, a LW piece with the same physical dimensions would fall into a different class if captured at a different resolution.

The training and validation metrics for the small version of the YOLOv10 model are displayed in Figure 6, which shows the performance metrics throughout 200 epochs. The precision and recall metrics significantly improved during the training process, with precision stabilizing around 0.5 and recall peaking at 0.45. These metrics indicate that the model performed reasonably well at identifying wood pieces despite room for improvement. The mAP50 for bounding box predictions achieved a peak of 0.46. While this value may seem modest relative to standard benchmarks (e.g., the COCO dataset; Lin et al., 2014), it reflects the challenges of detecting irregular, overlapping wood pieces in heterogeneous river environments, and we consider it a satisfactory result in this context. This value indicates the model's ability to accurately detect LW when considering a moderate overlap (IoU of 0.5) between predicted bounding boxes and ground truth labels. The stricter mAP50-95 metric, which averages performance across IoU thresholds between 0.5 and 0.95, achieved a maximum score of 0.22. This indicates that while the model performed well at moderate IoU thresholds, it still has difficulties finding the exact extent of the pieces of wood. This reflects the challenge of precisely localizing irregularly shaped wood objects and finding their limits. When training a similar small YOLOv10 detection model with 960p input data using the YCbCr colour space as suggested by Liang et al. (2022), the best mAP50 and mAP50-95 it produced were 27.4% and 12.6% respectively.

3.2 Wood segmentation performance

The segmentation model generally effectively determines the wood surface area within the detection bounding boxes, providing a solid basis for subsequent LW diametre and volume calculations. Figure 7 shows four detections segmented by the model. It shows that the algorithm performs well on individual pieces of wood but slightly follows an overlapping piece in Figure 7c. In Figure 7d, the algorithm follows the contours of a single LW piece, even though the bounding box is incorrectly showing multiple pieces of wood. Unlike the detection model, the segmentation model's metrics continued to show a gradual upward trend toward the end of the 200 epochs (Figure 8), suggesting that further training could result in additional incremental improvements.

The extra-large version of the YOLOv8 segmentation model at 320p yielded the best results, balancing detection accuracy and computational efficiency. The performance of the segmentation model is visualized in Figure 8, with key metrics including precision, re-

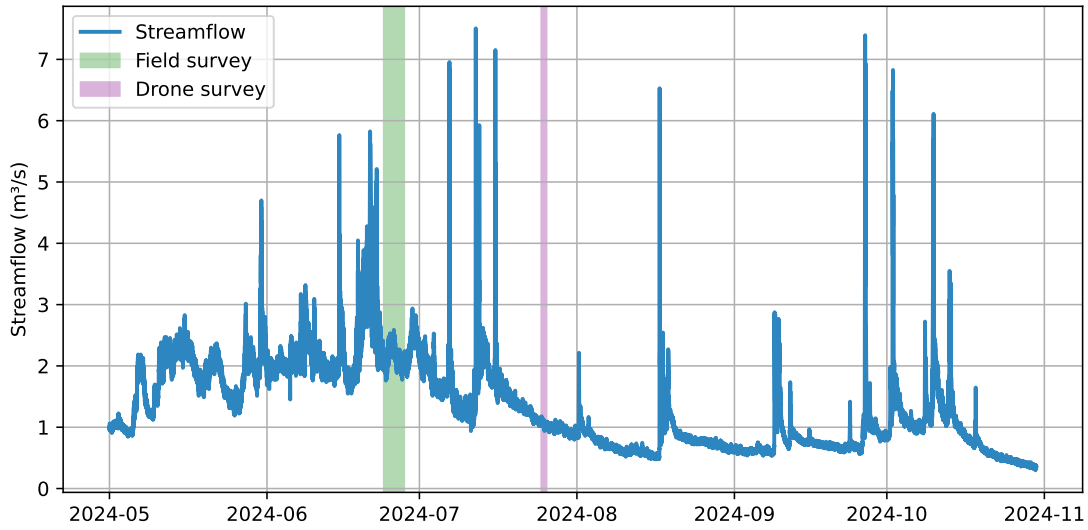

Figure 4 Streamflow at the gauging station at the outlet of the catchment. With field survey and drone survey periods indicated in green and purple.

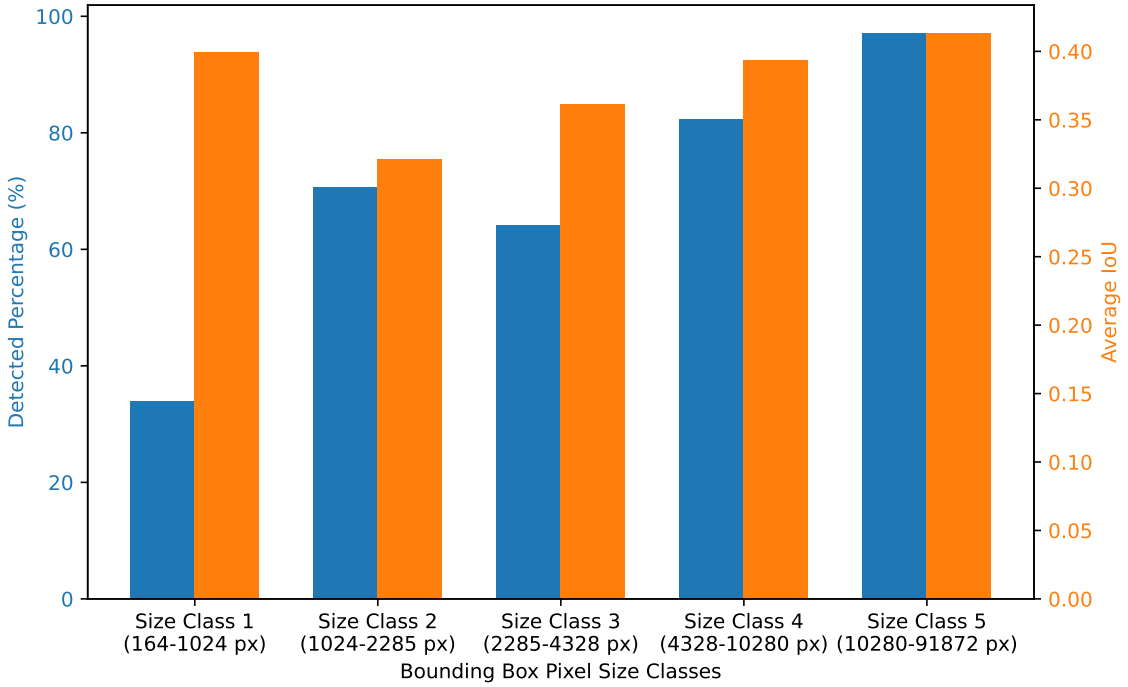

Figure 5 Percentage of detected pieces of wood in the validation dataset in blue at a confidence threshold of 0.3, and the average Intersection over Union (IoU) of the detected pieces in orange. The LW pieces are binned into brackets in terms of pixel size, representing the following approximate length ranges: Class 1: 43–108 cm, Class 2: 108–161 cm, Class 3: 161–221 cm, Class 4: 221–341 cm, Class 5: 341–1020 cm.

call, mAP50, and mAP50-95. The model's precision increased steadily throughout training, reaching approximately 0.8. The recall metric improved to around 0.65, demonstrating that the model successfully captured most of the wood pixels with a low number of false negatives. The mean Average Precision at an IoU threshold of 0.5 for segmentation masks reached values around 0.7, suggesting a substantial overlap between predicted masks and the ground truth at this IoU level. For the more strict mAP50-95 metric, which averages performance over IoU thresholds ranging from 0.5 to 0.95, the model achieved scores between 0.25 and 0.3. While this indicates a lower accuracy across varying overlap thresholds, it highlights the challenges of segmenting complex wood shapes in natural environments. The combination of high precision, recall, and robust mAP50 scores suggests that the extra-large YOLOv8 model at 160p is well-suited for wood segmentation tasks in river environments. However, this segmentation model is only executed on wood pieces that were detected in the first step of the method.

3.3 Algorithm Accuracy Assessment

To evaluate the model on unseen data, we used the Avançon de Nant river, which was not included in the training dataset and serves as an independent test site. All results presented in this section are based solely

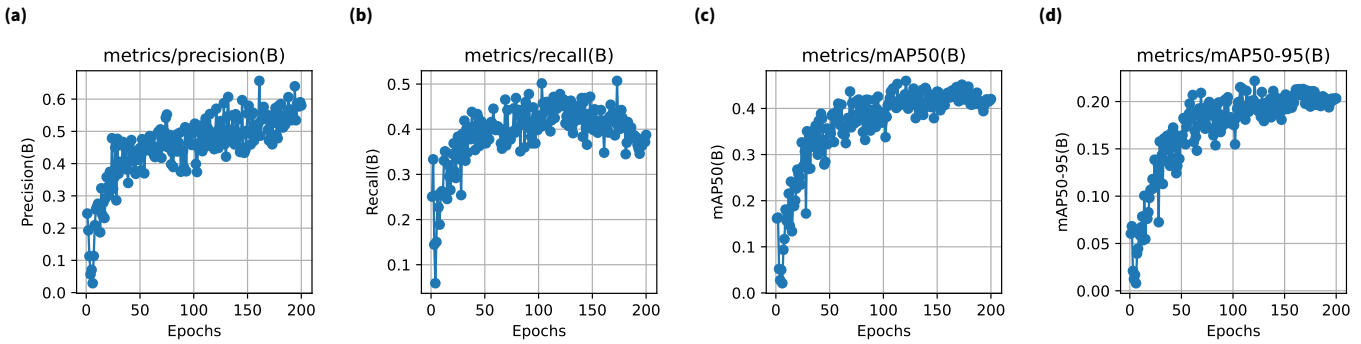

Figure 6 Results training the detection model (YOLOv10, small, 960p). **(a)** precision, which indicates how many of the LW pieces were detected; **(b)** recall, which indicates whether a detected box was a Lw piece; **(c)** mean Average Precision at an Intersection over Union (IoU) of 0.5; **(d)** mAP for IoUs ranging from 0.5 to 0.95. All values on the y axes are dimensionless scalars. An epoch in machine learning is one complete pass through the entire training dataset, during which the model updates its weights to improve learning.

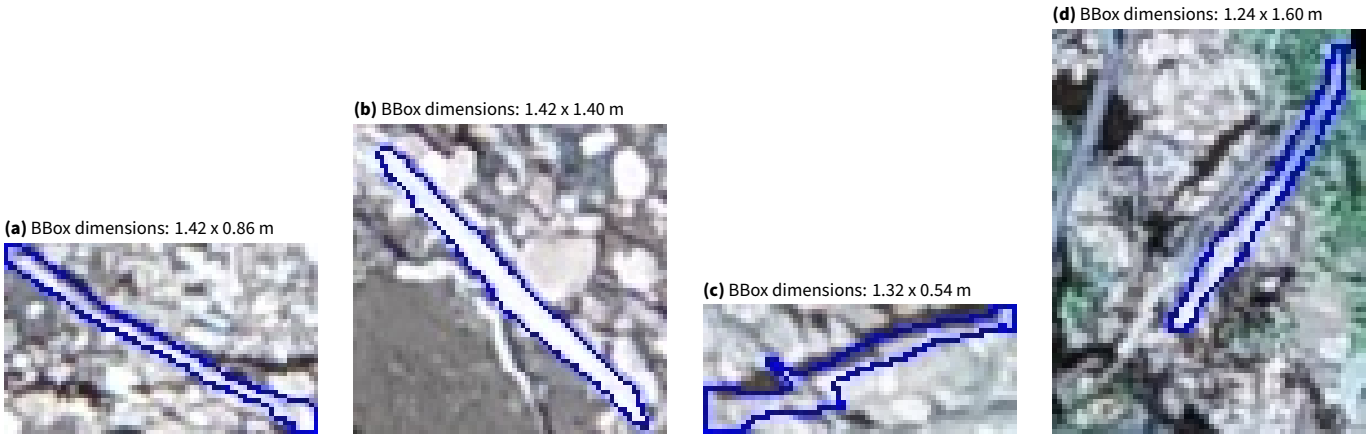

Figure 7 Four segmented detections from the Avançon de Nant dataset, **(a)**: detection 179, **(b)**: detection 391, **(c)**: detection 450, **(d)**: detection 515. The blue line indicates the polygon around all the pixels classified as wood. **(b)** and **(c)** are examples of wood pieces in jams.

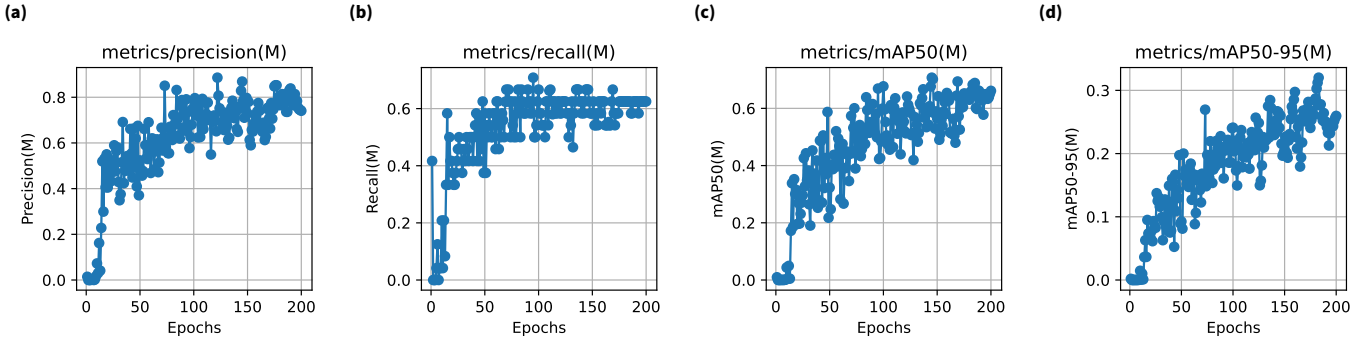

Figure 8 Segmentation model (YOLOv8, extra-large, 320p) training results. **(a)** precision, which indicates how many of the segmented pixels were correct; **(b)** recall, which indicates whether the segmented pixel was wood; **(c)** mean Average Precision at an Intersection over Union (IoU) of 0.5; **(d)** mAP for IoUs ranging from 0.5 to 0.95. All values on the y axes are dimensionless scalars. All values on the y axes are dimensionless scalars. For each epoch, the training process has used all the training data to improve performance.

on detections from this river. This approach ensures that the evaluation reflects the model's generalization to new environments and image conditions.

3.3.1 Individual wood pieces

This section presents an evaluation of detection and volume estimation accuracy for individual LW pieces. While segmentation provides the basis for estimating surface area, volume estimation involves further calculations that combine the segmented area with bounding box length, and therefore is analysed separately from segmentation performance. The comparison between the surveyed individual wood pieces and the detected ones shows a generally high level of accuracy in detecting the correct orientation of the wood pieces, with the algorithm correctly identifying the orientation in 83% of the cases. This means it is a useful tool for analysing spatial arrangement. Table 1 compares 29 LW pieces surveyed during fieldwork that were also de-

Table 1 Comparison between individual LW measurements in surveys (Sur.) and estimates from the algorithm (Alg.). The second and third columns show the surveyed lengths and diametres, where the sixth and seventh column shows the LW lengths and diametres as analysed by the algorithm.

Survey Number	Length Sur. (m)	Diam. Sur. (m)	Vol. (m^3)	Algorithm Number	Length Alg. (m)	Diam. Alg. (m)	Vol. Total (m^3)	Correct Orientation?	Vol. Diff. (m^3)
N17	4.1	0.12	0.046	**333**	4.0	0.14	0.060	no	-0.013
N13	7.3	0.35	0.702	**682**	7.8	0.37	0.836	yes	-0.133
N19	3.0	0.10	0.024	**158**	3.6	0.11	0.036	yes	-0.013
N10	3.0	0.12	0.034	**297**	3.2	0.16	0.065	yes	-0.032
1005	1.9	0.22	0.072	**141**	2.0	0.17	0.046	yes	0.027
15	2.7	0.10	0.021	**294**	2.4	0.12	0.028	yes	-0.007
22	3.5	0.50	0.687	**632**	3.4	0.23	0.146	yes	0.541
14	1.6	0.15	0.028	**654**	1.0	0.14	0.017	yes	0.011
2623	2.5	0.10	0.020	**488**	2.7	0.17	0.059	yes	-0.039
29	15.5	0.35	1.491	**102 & 103***	11.3 & 11.8	0.17 & 0.23	0.470	yes	1.021
N2	2.6	0.12	0.029	**268**	2.7	0.20	0.084	yes	-0.054
8	2.8	0.12	0.032	**617**	2.5	0.12	0.030	no	0.002
6	2.6	0.18	0.066	**611**	2.2	0.12	0.025	yes	0.041
34	2.8	0.17	0.064	**79**	2.7	0.15	0.047	yes	0.016
2	1.6	0.12	0.018	**422**	1.7	0.15	0.029	no	-0.011
66	2.6	0.12	0.029	**412**	2.7	0.13	0.037	yes	-0.008
82	1.7	0.10	0.013	**63**	2.0	0.13	0.027	no	-0.014
78	2.2	0.18	0.056	**590**	3.2	0.23	0.136	yes	-0.080
N42	1.6	0.10	0.013	**225**	1.6	0.11	0.016	yes	-0.003
N54	2.8	0.10	0.022	**199**	2.2	0.06	0.007	yes	0.015
105	2.3	0.39	0.275	**379**	1.1	0.08	0.005	no	0.270
106	1.5	0.12	0.017	**198**	1.5	0.12	0.016	yes	0.001
N45	1.0	0.14	0.015	**378**	1.4	0.09	0.009	yes	0.006
N51	2.7	0.11	0.026	**532**	1.4	0.12	0.015	yes	0.011
103	3.0	0.13	0.040	**362 & 363***	1.5 & 3.2	0.03 & 0.10	0.027	yes	0.013
N50	1.8	0.15	0.032	**535**	1.7	0.17	0.036	yes	-0.004
N49	2.0	0.10	0.016	**181**	2.2	0.14	0.032	yes	-0.016
95	3.4	0.13	0.045	**9**	2.9	0.13	0.037	yes	-0.009
127	2.2	0.10	0.017	**12**	1.7	0.11	0.017	yes	0.000
Sum			3.951	**Sum**			2.394	83%	1.54
								Abs. Diff.	2.41

*Some LW pieces were detected with multiple bounding boxes.

tected by the wood detection algorithm in the orthomosaic. When looking at the detected length of the pieces, the table shows that the method performs well with a median error of 10%. The median difference between the measured and algorithmically estimated lengths was 0.3 m with an interquartile range of 0.4 m. The errors increased up to a median of 20% when looking at the detected diametre of the LW pieces. The median difference between the measured and algorithmically estimated diametres was 0.02 m with an interquartile range of 0.04 m. Regarding the volume, the errors compounded with a median error of 40%. The median difference between the measured and algorithmically estimated diametre was 0.014 m^3 with an interquartile range of 0.031 m^3. This analysis is restricted to 29 LW pieces that were identified in both the survey and the orthomosaic. These logs form the basis for length, diametre, and volume comparisons.

These findings are underlined by the data in Figure 9. The three scatterplots compare the survey data for LW lengths, diametres, and volumes to the estimates made by the method. Ideally, all data points should align with the 1:1 line, representing perfect agreement between the observations and detections. However, in each plot the data points deviate from this ideal line, particularly for larger wood pieces. The largest pieces had the most impact on the estimation of the total volume. On average, the largest four pieces of wood are underestimated by 0.42 m^3, whereas the other 25 pieces are slightly overestimated by 0.0056 m^3 on average. The median uncertainty for the volume estimation based on all 29 pieces is 40%. This discrepancy is most evident in the diametres and, as a result, in the volume, where the algorithm tends to underestimate these values for larger pieces. The consistent underestimation of diametre is a key factor leading to the overall underestimation of volume on the reach scale, as volume is directly related to the length and quadratically related to the diametre (considering the wood pieces as cylinders). This pattern is captured in Table 1.

3.3.2 River reach scale

Unlike the individual comparison in the previous section, this reach-scale analysis includes all LW pieces surveyed and all LW detected, even if individual one-to-one

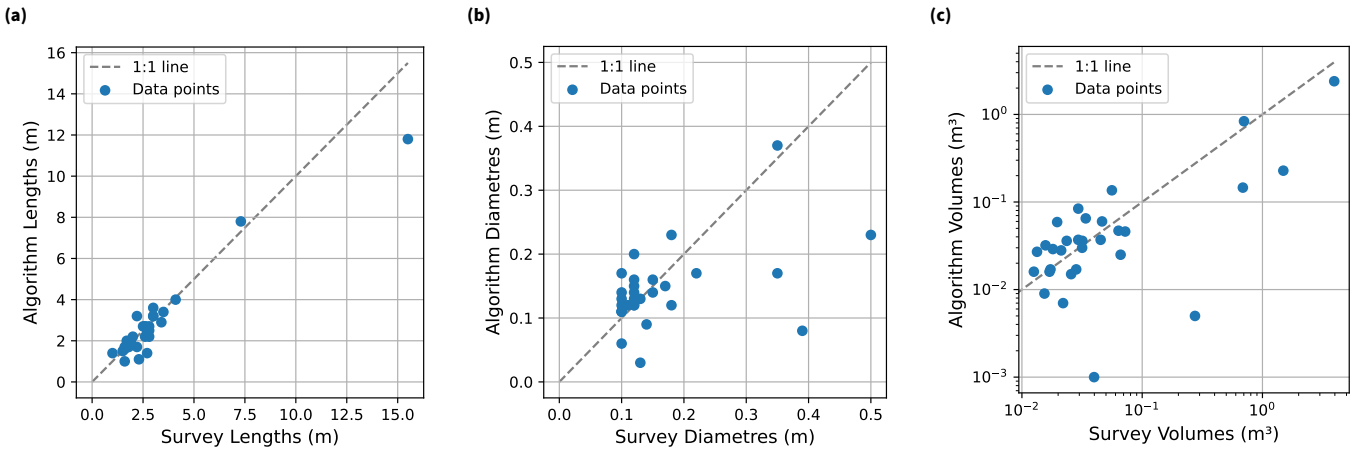

Figure 9 Algorithm data versus survey lengths, diametres and volumes. The 29 LW pieces from Table 1 compared. Survey data vs algorithm data, **(a)**: lengths, **(b)**: diametres, **(c)**: volumes. The algorithm would work perfectly if all the dots were on the 1:1 line.

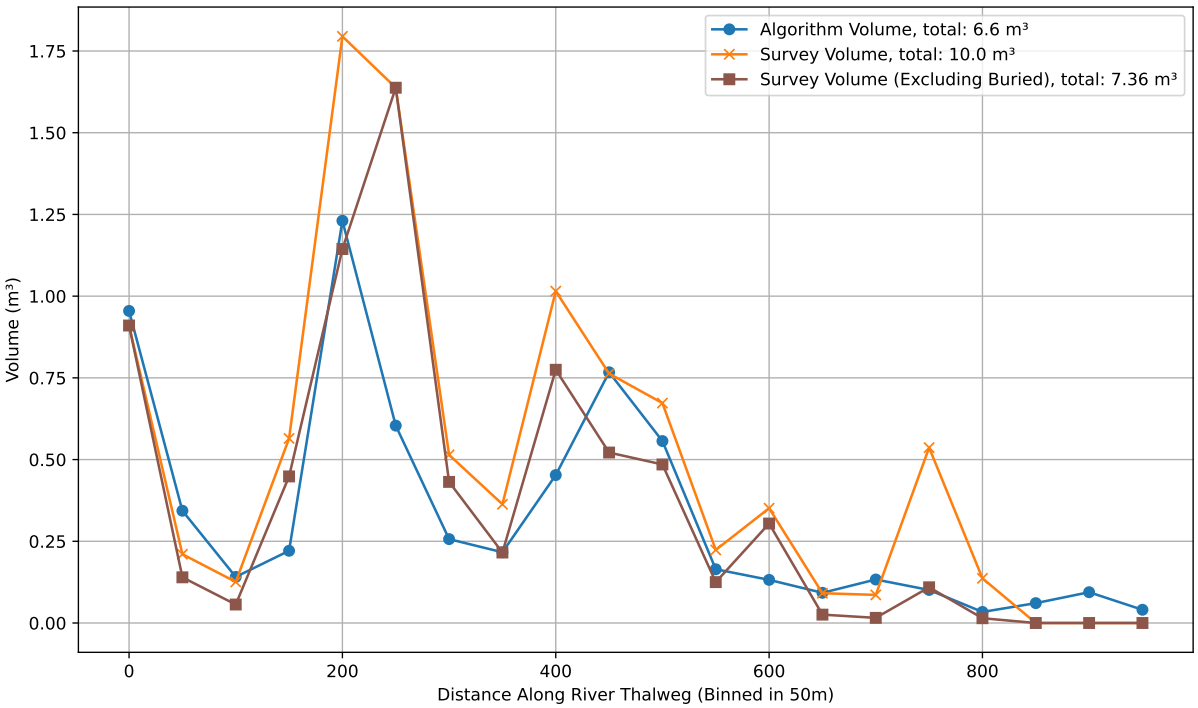

Figure 10 Comparison of volumes along the thalweg in sections of 50 metres, starting upstream. Algorithm vs Survey vs Survey excluding buried pieces.

matching was not possible. The total number of LW pieces surveyed was 135. At the same time, the wood detection algorithm identified 221 pieces in the same area. According to the algorithm, 118 meet the criteria of large wood (longer than 1 m and larger than 0.1 m diametre). The remaining 17 were either too short or too narrow to qualify as large wood based on the defined size thresholds. These smaller pieces are not necessarily false positives. When looking at the larger scale wood storage estimation, Figure 10 shows the volume per 50 metres along the river as detected by the algorithm in blue versus the data from the survey in orange and the data from the survey excluding the buried LW in green. Even though the two datasets were acquired one month apart from each other, and after a visual inspection showed changes in LW deposited on the bed, the peaks in wood storage demonstrate the concentrations of wood along the river bed and indicate correctly where the largest jams were located. The signal of the wood stored along the river generally agrees between the survey and the algorithm. The method estimated the wood volume equal to 6.60 m^3, whereas the total volume of unburied wood found during the survey was 7.36 m^3. The discrepancy can partially be attributed to underestimating the number of pieces in large jams.

Figure 11a compares wood orientations as derived by the algorithm and measured during the field survey. LW pieces were classified into three orientation categories: Orientation 1 (0–22.5°), Orientation 2 (22.5–67.5°), and Orientation 3 (67.5–90°). The algorithm's orientation predictions generally align with the survey data and show more pieces with orientation two than with the other orientation. The algorithm also analyses a similar amount of LW with orientation 3. However, there is

a noticeable discrepancy in LW pieces with orientation 1, where the algorithm predicts fewer than the survey data.

Figure 11b compares the total wood jam volumes of the survey and the algorithm, where the survey data was gathered using traditional measurements. A discrepancy is observed when estimating the largest wood jams, where the algorithm underestimates the total volume. This underestimation is due to the algorithm's difficulty in accurately determining the volume of the larger individual wood pieces. The largest observed wood jam consists of multiple pieces, including piece number 29. This piece (see Table 1, also shown in the middle part of Figure 12) was measured to have a volume of 1.49 m^3 in the survey versus 0.47 m^3 as estimated by the algorithm. Furthermore, a detailed quantitative analysis revealed that the mean absolute error (MAE) between the algorithm's estimated jam volume and the actual volume measured during the survey was 0.10 m^3. To gain further insight into the algorithm's accuracy across different jam sizes, the jams were segmented into two groups: the six largest jams and all other jams. The results indicated that the mean absolute difference for the six largest jams was 0.20 m^3, suggesting a higher volume estimation discrepancy for larger jams. In contrast, the mean absolute difference for the remaining jams was significantly lower at 0.04 m^3, indicating a more precise volume estimation for smaller jams. These findings suggest that while the algorithm performs well overall, its accuracy tends to decrease for larger jams, possibly due to increased complexity in their volume estimation.

3.4 Efficiency of the proposed method

As shown in Table 2, our method significantly reduces the time required for wood volume mapping compared to traditional approaches. A traditional field survey requires approximately 6 work days for a 15-hectare area, while UAV imagery with manual mapping reduces this to 4 work days. In contrast, our automated approach completes the task in only 1.5 work days, enabling a much faster and more scalable solution.

4 Discussion

4.1 Strengths, novelty, and challenges

As summarised in Table S1, few prior studies have proposed an automated approach comparable to this work, but they differ significantly in methodology. For instance, Sendrowski and Wohl (2021) employed both object-based and pixel-based classifications using supervised support vector machines (SVM) and unsupervised ISO clustering on multispectral satellite imagery, classifying units like water, sediment, vegetation, and wood. Their performance, evaluated by correctly classified reference samples, showed overall accuracies between 65% and 99%. Another study by Sendrowski et al. (2023) applied a binary (wood/no wood) deep learning approach to pixel-wise classification. Similarly, Buscombe et al. (2024) used a deep learning model with distinct training strategies and higher-resolution imagery, achieving segmentation accuracies ranging from 0.49 to 0.98 for wood detection. However, these studies did not calculate metrics such as volume, size, or orientation of individual pieces of wood. Liang et al. (2022) did, however, compute volumes. Regarding their results (Table 5 in Liang et al., 2022), their approach overestimated the total LW volume by 43% (22.65 m^3 in the survey and 32.45 m^3 using the automatic approach). However, when comparing volumes per block in their approach to our results in Table 1, we see similar R^2 values (0.53 and 0.48, although we are comparing block volumes to individual piece volumes). Furthermore, changes in colour space increased the performance of their methods, something we did not find.

The automated detection method presented in this study introduces advancements in LW monitoring by overcoming the limitations of previous approaches. While earlier methods focused primarily on either the detection (Rouge, 2022; Reymond, 2022) or segmentation (Pásztory and Polgár, 2016) of wood from aerial or UAV imagery, this novel approach goes beyond identification. Our method detects individual LW pieces and provides detailed measurements of metrics such as length and diametre. Furthermore, it distinguishes whether the wood is part of a jam and determines the orientation of each piece, which might influence the tendency to move (Merten et al., 2010), aspects that have not been addressed in prior work. The ability to quantify these attributes is a substantial improvement because it is the same as the data typically obtained through traditional fieldwork.

LW length, diametre, orientation, and identifying wood jams are important for understanding wood dynamics in river systems. Collecting such data in the field has traditionally been labour-intensive (Galia et al., 2018), often requiring several days of work across challenging terrain. In contrast, this automated approach allows for the rapid collection of comparable data using UAVs within a fraction of the time (Sendrowski and Wohl, 2021; Hess et al., 2024). A drone survey can cover a large river reach in a matter of hours, providing high-resolution orthomosaics that enable measurements without manual intervention. The data derived from the method are getting close to the data quality collected manually, thus facilitating new initiatives for large-scale monitoring of river systems.

Another key strength of this method is scalability. Unlike field surveys, which are limited by personnel and time constraints, using UAV imagery combined with machine learning allows for monitoring LW over large and remote areas. This is particularly advantageous in regions where fieldwork is impractical or access is restricted due to environmental conditions. The utilization of advanced object detection algorithms (e.g., YOLOv8 and YOLOv10, Jocher et al., 2023) ensures that the system is accurate and efficient, capable of detecting a high percentage of LW pieces with minimal error. The capacity to automate such a process brings the promise of near-real-time monitoring, providing river managers and researchers with up-to-date information that can be critical for river restoration, flood mitigation, and river management.

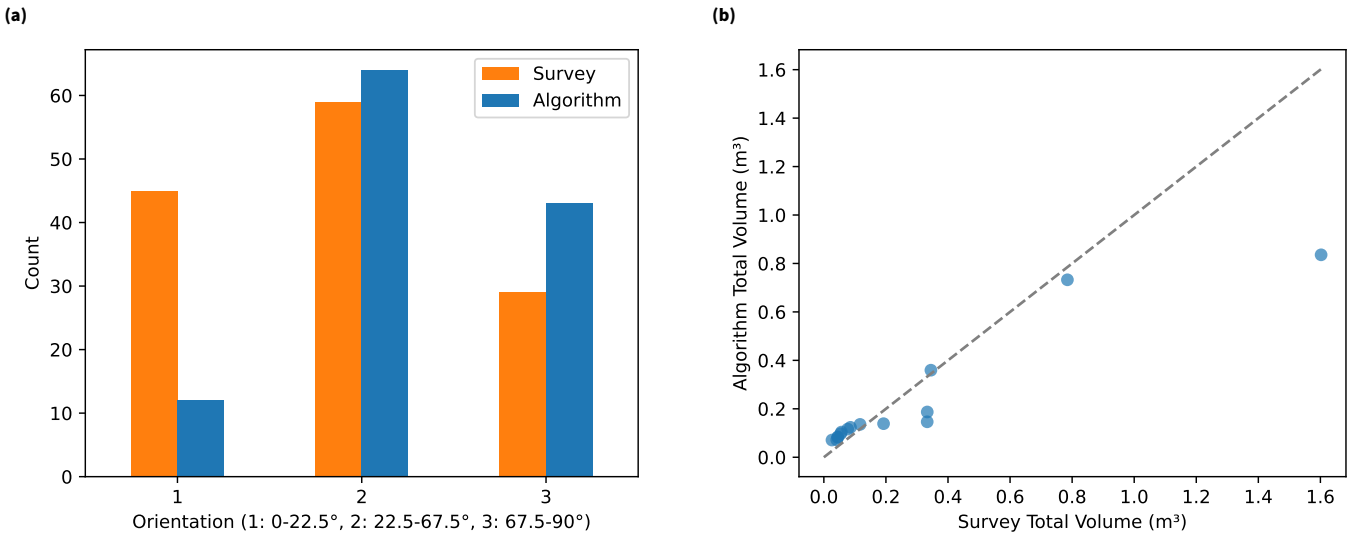

Figure 11 Comparison of **(a)** orientations and **(b)** jam volumes between the algorithm data in blue and the survey data in orange. Only jams with at least a volume of at least 0.025 m^3 were plotted.

Understanding camera parametres, particularly aperture and shutter speed, is important when capturing images in environments like river systems. Automatic white balance adjustments in cameras tend to neutralize colours towards a mid-tone grey (Zapryanov et al., 2012), which can be problematic in areas with sediment bars and vegetation. Sediment bars can appear washed out and overexposed, leading to a loss of detail. This overexposure can obscure important surface features for image analysis. Reducing the shutter speed could help capture colour information more effectively, preserving detail and contrast (James et al., 2017). Optimizing exposure settings is, therefore, important for improving the accuracy of detection and classification models from aerial imagery.

Another limitation is the assumption that all wood pieces have a cylindrical shape. In practice, wood pieces may have irregular forms with several branches, which can affect both field measurements and algorithm-derived estimates (Hortobágyi et al., 2024a). This simplification likely contributes to an error in the estimation of volume, especially for larger or more complex pieces. A useful future development could involve validating the volume estimates of irregularly shaped logs with high-precision field measurements. By comparing detailed field-derived volumes of complex-shaped wood pieces with those estimated by the current method, it may be possible to quantify the extent of over- or underestimation. If a systematic discrepancy is found, the method could be enhanced by adding a shape-classification step to identify cylindrical and irregular wood pieces. Different volume estimation algorithms could then be applied depending on the LW piece geometry. In a similar way, the lengths of in-field pieces can be compared to the field-derived lengths, and the extra step could include a better estimation of length.

Also, a potential improvement for volume estimation would be to split the segmented shape along its main axis and estimate local diametres for each segment, summing the volumes to obtain a more accurate total. This approach could better capture irregular or tapered LW pieces and is a promising direction for future refinement. As field-derived estimations are also prone to simplifications with these added steps, the method could potentially become more accurate than field measurements.

A limitation of our accuracy assessment is that the test dataset was based entirely on a single river, the Avançon de Nant. This river's specific characteristics, such as typical wood dimensions, channel morphology, and riparian forest type, may have influenced detection performance. As a result, the model's accuracy in this study may not fully reflect how it would perform in rivers with very different environmental settings or wood recruitment processes. Expanding field validation to additional rivers with varying geomorphic and ecological contexts is an important next step to assess model generalizability and robustness.

4.2 Wood detection

The wood detection algorithm effectively detects LW (see Figure 1). Using distinct datasets for training/validation and testing strengthens the robustness of the results. When using the CNN with a confidence threshold of 0.3 (detecting pieces of which the network has at least a 30% confidence), 66.2% of all pieces in the validation data (regardless of the large wood criteria) and 97.1% of LW in the largest size bracket (between 3.4 and 10.2 metres, see Figure 5, bracket 5) were detected. The mean average precision (mAP) was 46%. The modest mAP score can be attributed to the detection task's complexity and the mAP metric's intrinsic limitations. Detecting LW in natural environments is inherently challenging due to the variability in shape, size, and environmental factors such as shadows, rocks, and vegetation. Despite these challenges, although the algorithm underestimates the total size of some jams and overestimates others when looking at the total amount of

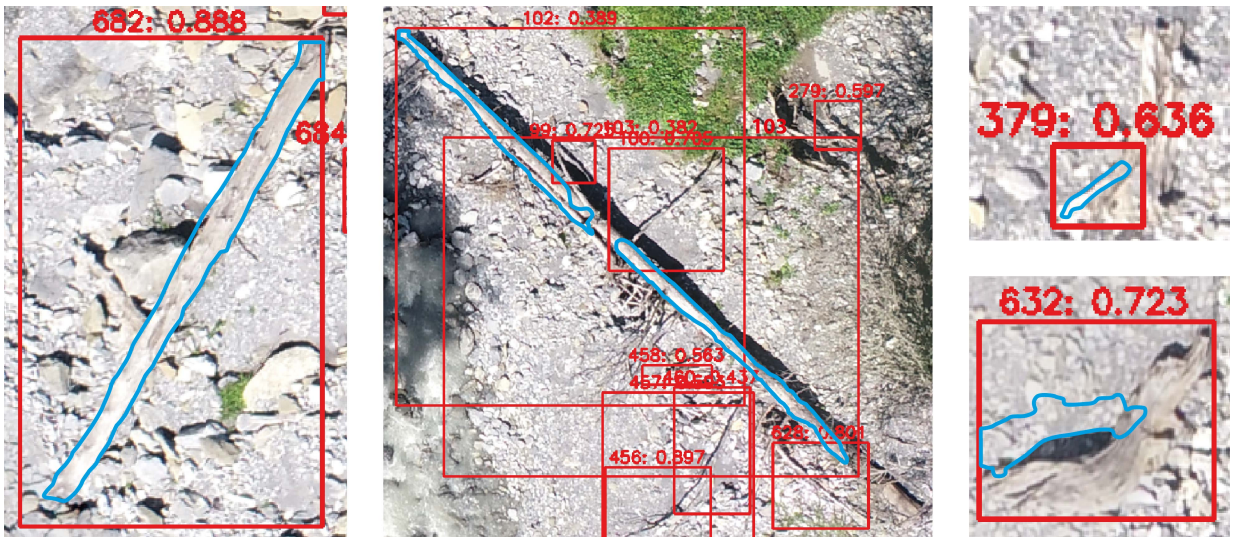

Figure 12 Four largest LW diametre from Table 1 segmented. A blue line is drawn around the segmentation mask and thus indicates the pixels which the segmentation algorithm identifies as wood. Above each Bounding Box, the detection ID and confidence score are shown in red.

Table 2 Comparison of time and effort required for wood detection methods (for 15 ha, based on our own findings, doing surveys, labelling the data and running our algorithm). Processing time includes pre-processing (e.g., cropping the river section, defining thalweg shapefiles), running the detection and segmentation algorithms, and post-processing (e.g., volume and orientation analysis). It excludes model training, which is performed separately and reused across datasets.

Method	Field Time (work days)	Processing Time (work days)	Total Time (work days)
Traditional Field Survey	5 (2 people)	1/2	6
UAV Imagery & Manual Mapping	1	3	4
UAV Imagery & Proposed Approach	1	1/2	1.5

wood stored in the Avançon de Nant test case, the algorithm detects 6.60 m^3 where the infield survey found 7.36 m^3 (89.7 percent) of unburied wood pieces. Unlike what was found by Liang et al. (2022), training detection in the YCbCr colour space was not found to yield better results as compared to RGB. The level of accuracy demonstrates the method's potential for estimating visible wood volume, though limitations remain, such as the underestimation of diametres for thicker pieces (see Table 1 and Figure 9b). Despite these constraints, the results suggest that the method can contribute to LW monitoring by aiding in the detection and quantification of unburied wood. The mAP metric, though widely used, is highly sensitive to deviations in bounding box placement and size. This can be attributed to the stringent Intersection over Union (IoU) threshold of 0.5. Minor misalignments, which may be inconsequential to human observers, can significantly lower the mAP score. The IoUs shown in Figure 5 are lower than the standard 0.5 thresholds as used by the calculation of the mAP; the nature of the data makes achieving higher IoUs particularly challenging. LW in orthomosaics often lack clearly defined boundaries, which introduces subjectivity in labelling, even among human annotators. Thus, although the IoU scores do not often exceed 0.5, the high detection rates for larger LW pieces (the most important for LW volume) demonstrate that the algorithm effectively captures the estimation of the total wood volume. As a result, while the algorithm effectively detects larger, more prominent pieces of wood critical for volume estimation, the score does not fully reflect its practical effectiveness. This highlights the need to consider different quantitative metrics together with practical outcomes carefully when evaluating performance in complex environments, such as the detection of wood in rivers. While mAP provides insights as to how well the algorithm labels the data in the same way as human annotators would, it may not fully capture the algorithm's efficacy in detecting all LW. Future work should prioritise expanding the training dataset and refining the model to improve performance (Zhao et al., 2019; Kamilaris and Prenafeta-Boldú, 2018) so that the model would increase its detection rate on even the smallest pieces.

Furthermore, the results indicate that the classification algorithm has difficulties with identifying shorter, thicker LW pieces. By comparing the percentage of pixels classified as wood in each cropped bounding box to the LW length (detailed in the supplementary material), we observed that examples, where the algorithm underperformed (LW pieces shorter than 5 metres and thicker than 27 centimetres), fell outside the envelope of ratios represented in the training data. This indicates that the training data lacks sufficient examples of short, thick LW, likely contributing to the algorithm's reduced performance on LW pieces with these dimensions. The model was trained on a relatively small dataset of only 206 labelled images. This limited dataset likely contributes to the model's lower precision, particularly for smaller or irregularly shaped wood pieces, which may be underrepresented in the training data. Expanding the dataset with more labelled images, especially of

complex and varied wood structures, would improve the model's generalisation and overall precision. With a larger, more diverse training dataset, the algorithm would likely achieve a higher mAP, reflecting more accurate detections across different LW shapes and sizes.

4.3 Wood segmentation

It was shown that from the 29 LW pieces found in both the survey and by the algorithm (see Table 1), the average error in length was 15.6% and in diametre was 25.7%. However, this dataset is not representative of all the data because the data mainly included pieces that were not in wood jams. This is because, during fieldwork, for each wood jam only one location was acquired instead of a location for each individual LW piece. This made it difficult to determine definitively which LW piece in the survey corresponded to the LW piece detected. Therefore, the error regarding length and diametres for LW in jams might differ. Furthermore, these statistics do not consider the LW that were not detected by the initial algorithm, which is governed by the performance of the detection algorithm.

Figure 12 shows the 4 LW pieces with the largest diametres from Table 1 and their mask. Apart from detection 682, the figure demonstrates a structural underestimation of the number of pixels segmented as wood. In the case of detected piece number 379, not only did the segmentation algorithm perform substandardly, but the detection model also did not detect the entire LW piece, adding to the error.

One approach to address the underestimation of the diametre would be to apply a correction factor. By analysing the percentage of underestimation relative to the bounding box size, it would be possible to devise a scaling factor to multiply the algorithm-derived diametres. This method, however, would require site-specific estimations and, therefore, might not necessarily be a robust solution. A more effective solution would be to improve the algorithm's training data. As stressed in other studies, uncertainties will inevitably arise when correcting such data, as aerial images cannot fully replace field surveys (Galia et al., 2023; Hortobágyi et al., 2024b).

The training data revealed a gap in examples where the mask covers a high percentage of the pixels, indicating thick LW pieces relative to the bounding box diagonal, which corresponds to the length of the piece (see Figure 13). The survey data points in the top left corner of the figure highlight five LW pieces from Table 1 with diametres over 37 cm and lengths shorter than 5 metres. These LW pieces have been manually labelled for this analysis. The training data lacks examples in the range of thick LW, which limits the algorithm's ability to segment such pieces accurately. To address this limitation, labelling additional thick, short LW and retraining the model would enable the algorithm to better generalize to complex shapes and provide more precise estimates of both diametre and volume. Prioritizing these improvements in future algorithm versions would lead to a more robust and reliable solution.

4.4 Wood size and volume at the reach scale

Although effective in identifying the presence of wood, the trained algorithm tends to underestimate the diametre of individual pieces (Figure 13). This underestimation, although relatively small in absolute terms, has a substantial impact on the final LW volume estimation due to the mathematical nature of the calculation. Since the volume of a cylindrical object is determined by squaring the diametre, even minor inaccuracies in diametre measurement are amplified significantly when calculating the final volume. As a result, this discrepancy leads to noticeable differences between the algorithm's predicted volumes and those derived from traditional survey-based wood measurement methods, potentially affecting the accuracy and reliability of inventory assessments or resource estimations.

The results show a discrepancy between the number of LW pieces surveyed and the number detected by the algorithm, which can be explained by several factors. First, the data were collected on different dates, and variations in storage conditions were observed, meaning we are not analysing exactly identical datasets. Consequently, there may have been more or less LW present during the drone survey. Also, the field survey only accounted for LW pieces with a diametre greater than 0.1 metres and a length exceeding 1 metre, whereas the detection algorithm identified all wood pieces regardless of size. Although post-processing removed all pieces shorter than 1 metre, we did not exclude those smaller than 0.1 metres in diametre due to the potential underestimation of diametre and disregarding actual large wood. However, the proposed detection method does not experience a significant increase in processing time when including smaller pieces, suggesting that these can reasonably be considered.

In Figure 10, the comparison of volumes between the algorithm and survey data clearly shows a divergence. The algorithm systematically underestimates the volume compared to the survey, with peaks in the survey volume data being noticeably higher than predicted. This discrepancy is particularly evident in stretches with larger and more irregularly shaped wood, suggesting that the segmentation algorithm has difficulties accounting for such variations. Figure 13 further supports this by illustrating the relationship between length and diametre, with the algorithm consistently predicting smaller diametres than observed in the survey. The figures emphasize the gap between the two, with the survey data showing a generally larger diametre for a given length. This discrepancy directly translates into the underestimations in Figure 10.

When comparing the lengths and diametres from the survey with those detected by the algorithm, Figure 13 illustrates the lengths of the 135 surveyed LW pieces versus the 221 detected ones, along with their diametres. The lengths between the two datasets correspond reasonably well, although the detected dataset includes more small pieces, likely because those smaller pieces were not surveyed due to their size. However, the diametres show a notable discrepancy: while the surveyed dataset includes several LW pieces with diame-

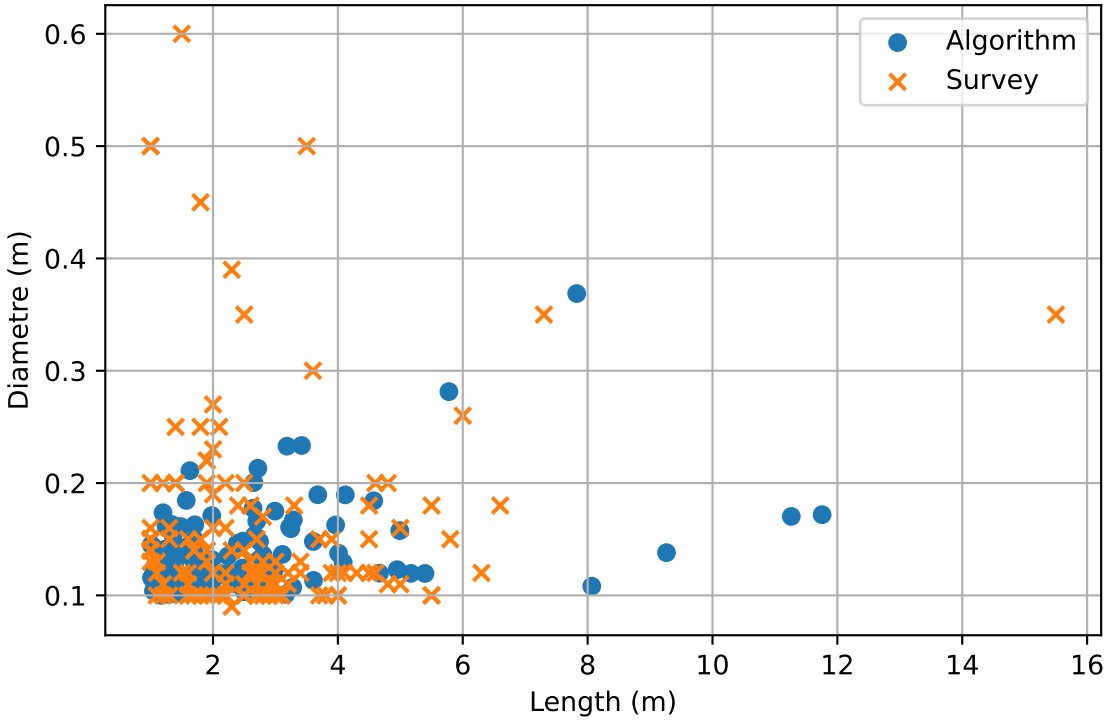

Figure 13 Lengths vs diametres of all the LW pieces that meet the large wood criteria in the survey data in orange and the algorithm data in blue.

tres over 0.25 metres, the algorithm does not estimate any LW piece under a 5-metre length to have a diametre over this threshold. This underestimating diametres leads to a more significant underestimation of volumes, given the quadratic relationship in volume calculations. It suggests that the algorithm may be overfitting to thin pieces, which are overrepresented in the training dataset. Therefore, it is advised to either use different segmentation models for different size bounding boxes to increase performance, or to use a correction factor for volume estimation.

We selected all surveyed LW with a diametre greater than 37 cm and a length shorter than 5 metres. These LW pieces are shown in Figure 14. Among them, three LW pieces (ID 64, 12, and N8) are not visible in the orthomosaic. This may be due to burial by sediment, vegetation cover, or displacement that occurred between the time of the field survey and the drone flight. LW piece number 64 could be number 413 in the detections; however, this is unclear from the data. Also, the segmentation algorithm did not indicate any pixel to be classified as wood in the cropped-out bounding box. The survey indicates that the particular piece is partially buried, which suggests the lack of effectiveness of the method regarding buried or partially buried LW.

The detection algorithm had difficulties detecting thicker LW pieces. LW piece 3, 50, N24 and N46 are visible in the orthomosaic but are not detected by the algorithm. This could be explained by the small number of thicker, shorter LW pieces in the training data. Also, the segmentation algorithm seems to suffer from this lack of training data for pieces with similar characteristics. LW piece 632 and 379 are (partly) detected but improperly segmented (see Figure 12). Again, this is the case because most of the training data includes thin pieces. A solution to this limitation would be to extend the training database of both the detection and segmentation algorithms to include more short, thick examples.

When looking at the comparison between the distribution of the orientations between the survey and the detections in the algorithm (see Figure 11a), the results can be unreliable because of the tendency of the algorithm to infer multiple detections on the same wood piece. When double detections are made, even though the orientation analysis might be correct, the algorithm still counts the piece with that orientation double. This could lead to inaccuracies when analysing the different jams and their volumes (see Figure 11b). The algorithm can indicate a single LW piece as a wood jam if detected multiple times. Therefore, as large jams tend to be underestimated in volume whilst smaller jams tend to be overestimated, with the current performance of the algorithm, the wood orientation and precise jam volume estimation remain challenging. Detecting wood jams presents several challenges due to the complex nature of the jam. One primary difficulty arises from the multiple layers of wood that make up a jam, which can obscure individual pieces. Additionally, wood located beneath a canopy or buried adds more complexity, as it is often hidden from view in aerial or satellite imagery. A further limitation in the detection process of the method is related to the resolution and size of the detection boxes. The orthomosaic is cropped into boxes of 960x960 pixels, with the test data having a resolution of approximately 2 cm per pixel. As a result, the maximum length of a detectable wood piece is theoretically capped at 27.15 metres. However, this limit varies depending on the pixel size of the orthomosaic being

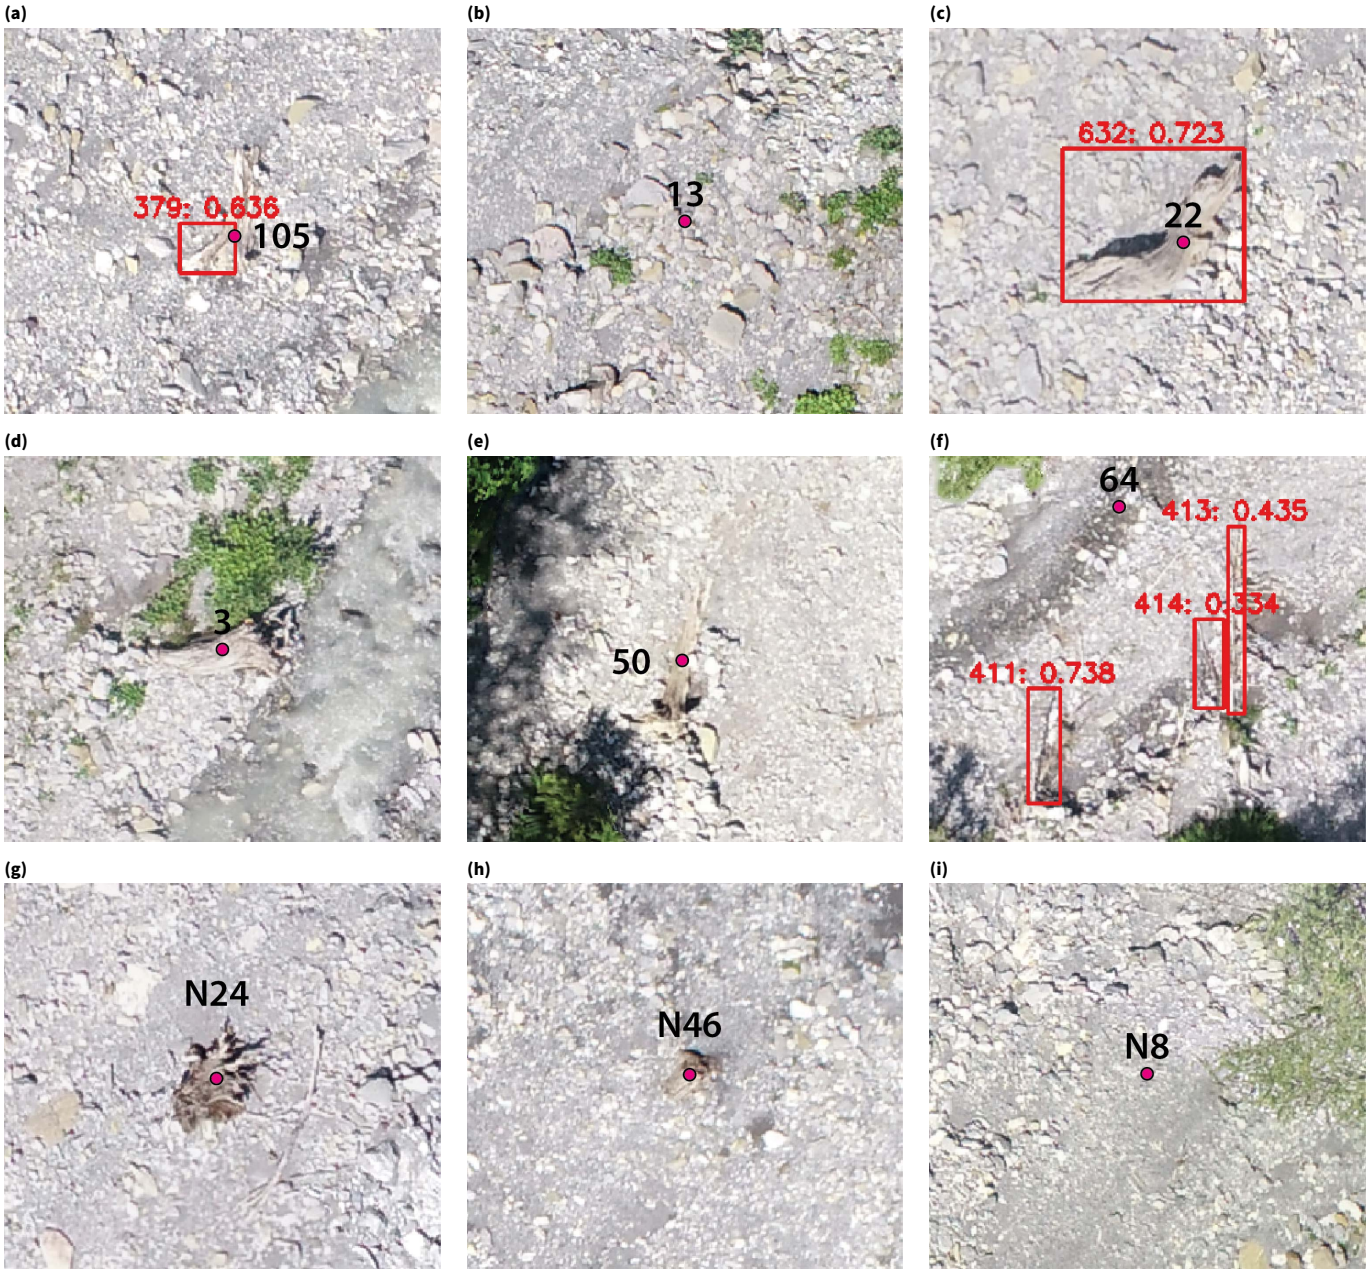

Figure 14 Examples of short LW pieces with large diametres. For different reasons, the LW pieces are not correctly analysed. **(a)** LW piece 105: partially detected. **(b)** LW piece 3: no visual. **(c)** LW piece 22: detected. **(d)** LW piece 3: not detected. **(e)** LW piece 50: not detected. **(f)** LW piece 63: buried. **(g)** LW piece N24, rootwad, not detected. **(h)** LW piece N46, stump, not detected. **(i)** LW piece N8: no visual.

analysed, which might differ in resolution from the test data. Finally, the method also relies on the availability of high-quality orthomosaics to ensure accurate detection and analysis. Lower-quality or less detailed images can hinder the ability to detect wood jams effectively. The resolution is a well-known problem in remote sensing of wood (Sendrowski and Wohl, 2021), and a correction has been proposed (Hortobágyi et al., 2024b) to account for the differences in the number of detected pieces due to the different resolutions of aerial images. These authors showed a linear relationship between image resolution and the number of wood pieces detected, as well as the jam number and the number of pieces in a jam. Similar corrections could be made to minimise this error, but they were not made in this study.

5 Conclusions

This study demonstrates the potential of convolutional neural networks (CNNs) for detecting and quantifying LW volumes in river environments using UAV-derived orthomosaics. The developed models successfully identified large wood and estimated its volume. However, challenges remained due to limitations in training data diversity, particularly for shorter, thicker LW pieces, and the complexity of wood jams. Increasing both the quantity and variety of training data could improve object detection and segmentation performance, further advancing this automated LW quantification approach.

The results revealed specific limitations, notably in segmenting irregular wood pieces and accurately estimating wood diametres, leading to volume underestimations, especially for shorter, thicker LW pieces. Expanding the training dataset with a broader range of wood examples and applying a correction factor for systematic underestimation of diametres could enhance accuracy.

While the current models show promise, further refinements are necessary to improve their accuracy and reliability. Key areas for enhancement include increasing training data diversity and refining segmentation techniques. These improvements will strengthen the models' robustness and applicability for large-scale river ecosystem monitoring and management.

To advance this field and enable fair comparisons of automated wood detection and quantification methods, we encourage the scientific community to collaborate on developing a standardized benchmark dataset for large wood detection and segmentation. Such a dataset would be a valuable resource for optimizing and validating machine learning and traditional computer vision models. Additionally, a comprehensive database of annotated large wood could support future advancements in LW quantification techniques.

Acknowledgements

The data from the Sense River was gathered by Bryce Finch and labelled by Christophe Reymond (University of Lausanne). The data from the Spöl River has been made possible by the Swiss National Park and partly labelled by Florent Rouge (University of Lausanne). The data from the Navizence River was gathered in the context of the Climact project NaviSedWood (University of Lausanne, École Polytechnique Fédérale de Lausanne, and the collaboration of CREALP) and the master thesis of Thomas Krieger (University of Lausanne). The field data in the Vallon de Nant was acquired with the help of Esteban Arcos and Pau Wiersma (University of Lausanne). We thank the comments from the associate editor, three anonymous reviewers, Prof. Bruce MacVicar, Prof. Stuart Lane and Prof. Pascale Biron on an early version of the manuscript, who helped us significantly improve the paper.

Land Recognition

We acknowledge that part of this research (the acquisition of the 'Casa de Piedra' and the 'Huemul' datasets) was conducted on the traditional lands of the Mapuche people.

Data availability

Data used in this manuscript is available at https://doi.org/10.5281/zenodo.13844170

Code availability

Code used to generate results is available at https://doi.org/10.5281/zenodo.17202014

Funding Statement

This work has been supported by the Swiss National Science Foundation project PCEFP2186963, the Leading House for Latin America (ECG2202) and the Universities of Lausanne and Bern (Switzerland), with the collaboration of the Swiss Federal Institute of Aquatic Science and Technology (Eawag).

Conflict of Interest Disclosure

There are no conflicts of interest.

Permission to Reproduce Material from Copyrighted Sources

No copyrighted material has been reproduced in this manuscript.

Declaration on Artificial Intelligence Use

AI-based tools were used to assist in coding and in language editing. The authors take full responsibility for the final content.

References

Antoniazza, G., Nicollier, T., Boss, S., Mettra, F., Badoux, A., Schaefli, B., Rickenmann, D., and Lane, S. N. Hydrological Drivers of Bedload Transport in an Alpine Water-

shed. *Water Resources Research*, 58(3), 2022. doi: https://doi.org/10.1029/2021WR030663.

Bertoldi, W., Gurnell, A., and Welber, M. Wood recruitment and retention: The fate of eroded trees on a braided river explored using a combination of field and remotely-sensed data sources. *Geomorphology*, 180-181:146–155, 2013. doi: https://doi.org/10.1016/j.geomorph.2012.10.003.

Buscombe, D., Warrick, J., Ritchie, A., East, A., McHenry, M., McCoy, R., Foxgrover, A., and Wohl, E. Remote Sensing Large-Wood Storage Downstream of Reservoirs During and After Dam Removal: Elwha River, Washington, USA. *Earth and Space Science*, 11, 08 2024. doi: https://doi.org/10.1029/2024EA003544.

Correia, B., Davies, R. L., Carvalho, F. D., and Rodrigues, F. C. Computer vision system for the automatic measurement of volumes of wood. In *Other Conferences*, 1993, https://api.semanticscholar.org/CorpusID:110473817.

Eaton, B. C., Hassan, M. A., and Davidson, S. L. Modeling wood dynamics, jam formation, and sediment storage in a gravel-bed stream. *Journal of Geophysical Research: Earth Surface*, 117(4): 1–18, 2012. doi: https://doi.org/10.1029/2012JF002385.

Galia, T., Ruiz-Villanueva, V., Tichavský, R., Šilhán, K., Horáček, M., and Stoffel, M. Characteristics and abundance of large and small instream wood in a Carpathian mixed-forest headwater basin. *Forest Ecology and Management*, 424(May):468–482, 2018. doi: https://doi.org/10.1016/j.foreco.2018.05.031.

Galia, T., Škarpich, V., Vardakas, L., Dimitriou, E., Panagopoulos, Y., and Spálovský, V. Spatiotemporal variations of large wood and river channel morphology in a rapidly degraded reach of an intermittent river. *Earth Surface Processes and Landforms*, 48 (5):997–1010, 2023. doi: https://doi.org/10.1002/esp.5531.

Gonzalez, R. C. and Faisal, Z. Digital Image Processing Second Edition. PDF uploaded to ResearchGate by Zahraa Faisal, 2019, https://www.researchgate.net/publication/333856607_Digital_Image_Processing_Second_Edition.

Goodfellow, I. J., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, Cambridge, MA, USA, 2016. http://www.deeplearningbook.org.

Grimmer, G., Wenger, R., and Chardon, V. RiverDetectWood: A tool for automatic classification and quantification of river wood in river systems using aerial imagery. *SoftwareX*, 29, 01 2025. doi: https://doi.org/10.1016/j.softx.2025.102042.

Gurnell, A. M., Tockner, K., Edwards, P. J., and Petts, G. E. Effects of deposited wood on biocomplexity of river corridors. *Frontiers in Ecology and the Environment*, 3(7):377–382, 2005. doi: https://doi.org/10.1890/1540-9295(2005)003[0377:EODWOB]2.0.CO;2.

Haschenburger, J. K. and Rice, S. P. Changes in woody debris and bed material texture in a gravel-bed channel. *Geomorphology*, 60:241–267, 2004. doi: https://doi.org/10.1016/j.geomorph.2003.08.003.

Hess, J. W., Pavlowsky, R. T., and Dogwiler, T. UAV Imagery for Measuring Large Wood Volume and Distribution: Ground Truthing Results and Sources of Error. *Southeastern Geographer*, 64: 11–29, 2024. doi: https://doi.org/10.1353/sgo.00001.

Hortobágyi, B., Milan, D., Bourgeau, F., and Piégay, H. How quickly does wood fragment in rivers? Methodological challenges, preliminary findings, and perspectives. *Earth Surface Processes and Landforms*, 2024a. doi: https://doi.org/10.1002/esp.5877.

Hortobágyi, B., Petit, S., Marteau, B., Melun, G., and Piégay, H. A high-resolution inter-annual framework for exploring hydrological drivers of large wood dynamics. *River Research and Applications*, 40(6):958–975, 2024b. doi: https://doi.org/10.1002/rra.4242.

Iroumé, A., Cartagena, M., Villablanca, L., Sanhueza, D., Mazzorana, B., and Picco, L. Long-term large wood load fluctuations in two low-order streams in Southern Chile. *Earth Surface Processes and Landforms*, 45(9):1959–1973, 2020. doi: https://doi.org/10.1002/esp.4858.

James, M., Robson, S., d'Oleire Oltmanns, S., and Niethammer, U. Optimising UAV topographic surveys processed with structure-from-motion: Ground control quality, quantity and bundle adjustment. *Geomorphology*, 280:51–66, 2017. doi: https://doi.org/10.1016/j.geomorph.2016.11.021.

Jocher, G., Chaurasia, A., Borovec, J., and et al. YOLO by Ultralytics. https://github.com/ultralytics/ultralytics, 2023.

Jutras-Perreault, M.-C., Gobakken, T., Næsset, E., and Ørka, H. O. Comparison of Different Remotely Sensed Data Sources for Detection of Presence of Standing Dead Trees Using a Tree-Based Approach. *Remote Sensing*, 15(9), 2023. doi: https://doi.org/10.3390/rs15092223.

Kamilaris, A. and Prenafeta-Boldú, F. X. Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147:70–90, 2018. doi: https://doi.org/10.1016/j.compag.2018.02.016.

Kuiper, S. D., Coops, N. C., Jarron, L. R., Tompalski, P., and White, J. C. An automated approach to detecting instream wood using airborne laser scanning in small coastal streams. *International Journal of Applied Earth Observation and Geoinformation*, 118: 103272, 2023. doi: https://doi.org/10.1016/j.jag.2023.103272.

Lassettre, N. S., Piegay, H., Dufour, S., and Rollet, A. Decadal changes in distribution and frequency of wood in a free meandering river, the Ain River, France. *Earth Surface Processes and Landforms*, 33:1098–1112, 2008. doi: https://doi.org/10.1002/esp.1605.

Liang, M.-C., Tfwala, S. S., and Chen, S.-C. The Evaluation of Color Spaces for Large Woody Debris Detection in Rivers Using XGBoost Algorithm. *Remote Sensing*, 14(4), 2022. doi: https://doi.org/10.3390/rs14040998.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., and Dollár, P. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014. doi: https://doi.org/10.1007/978-3-319-10602-1_48.

Manfreda, S., McCabe, M. F., Miller, P. E., Lucas, R., Pajuelo Madrigal, V., Mallinis, G., Ben Dor, E., Helman, D., Estes, L., Ciraolo, G., Müllerová, J., Tauro, F., De Lima, M. I., De Lima, J. L. M. P., Maltese, A., Frances, F., Caylor, K., Kohv, M., Perks, M., Ruiz-Pérez, G., Su, Z., Vico, G., and Toth, B. On the Use of Unmanned Aerial Systems for Environmental Monitoring. *Remote Sensing*, 10(4), 2018. doi: https://doi.org/10.3390/rs10040641.

Merten, E., Finlay, J., Johnson, L., Newman, R., Stefan, H., and Vondracek, B. Factors influencing wood mobilization in streams. *Water Resources Research*, 46(10), 2010. doi: https://doi.org/10.1029/2009WR008772.

Máčka, Z., Krejčí, L., Loučková, B., and Peterková, L. A critical review of field techniques employed in the survey of large woody debris in river corridors: a Central European perspective. *Environmental Monitoring and Assessment*, 181(1-4):291–316, 2011. doi: https://doi.org/10.1007/s10661-010-1830-8.

Ortega-Terol, D., Moreno, M. A., Hernández-López, D., and Rodríguez-Gonzálvez, P. Survey and Classification of Large Woody Debris (LWD) in Streams Using Generated Low-Cost Geomatic Products. *Remote Sensing*, 6(12):11770–11790, 2014. doi: https://doi.org/10.3390/rs61211770.

Piégay, H., Arnaud, F., Belletti, B., Bertrand, M., Bizzi, S., Carbonneau, P., Dufour, S., Liébault, F., Ruiz-Villanueva, V., and Slater, L. Remotely Sensed Rivers in the Anthropocene: State of the Art and Prospects. *Earth Surface Processes and Landforms*, 45, 12

2019. doi: https://doi.org/10.1002/esp.4787.

Pásztory, Z. and Polgár, R. Photo Analytical Method for Solid Wood Content Determination of Wood Stacks. *Journal of Advanced Agricultural Technologies*, 3:54–57, 01 2016. doi: https://doi.org/10.18178/joaat.3.1.54-57.

Reymond, C. Instream wood detection using YOLOv4 object detection algorithm. Master project. Master's thesis, University of Lausanne, 2022, https://wp.unil.ch/dawn/files/2023/02/final_report_christophe_reymond.pdf.

Rouge, F. Instream wood detection using the YOLOv4 algorithm on aerial images of the Spöl River. Master project. Master's thesis, University of Lausanne, 2022, https://wp.unil.ch/dawn/files/2023/02/ML_project_FROUGE_2022.pdf.

Ruiz-Villanueva, V., Díez-Herrero, A., Bodoque, J., and Bladé Castellet, E. Large wood in rivers and its influence on flood hazard. *Cuadernos de Investigación Geográfica*, 40: 229, 06 2014. doi: https://doi.org/10.18172/cig.2523.

Ruiz-Villanueva, V., Piégay, H., Gurnell, A. M., Marston, R. A., and Stoffel, M. Recent advances quantifying the large wood dynamics in river basins: New methods and remaining challenges. *Reviews of Geophysics*, 54(3):611–652, 2016. doi: https://doi.org/10.1002/2015RG000514.

Sanhueza, D., Picco, L., Ruiz-Villanueva, V., Iroumé, A., Ulloa, H., and Barrientos, G. Quantification of fluvial wood using UAVs and structure from motion. *Geomorphology*, 345:106837, 2019. doi: https://doi.org/10.1016/j.geomorph.2019.106837.

Sanhueza, D., Picco, L., Paredes, A., and Iroumé, A. A Faster Approach to Quantify Large Wood Using UAVs. *Drones*, 6(8), 2022. doi: https://doi.org/10.3390/drones6080218.

Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015. doi: https://doi.org/10.1016/j.neunet.2014.09.003.

Scott, D. N. and Wohl, E. E. Natural and Anthropogenic Controls on Wood Loads in River Corridors of the Rocky, Cascade, and Olympic Mountains, USA. *Water Resources Research*, 54(10): 7893–7909, 2018. doi: https://doi.org/10.1029/2018WR022754.

Sendrowski, A. and Wohl, E. Remote sensing of large wood in high-resolution satellite imagery: Design of an automated classification work-flow for multiple wood deposit types. *Earth Surface Processes and Landforms*, 46(12):2333–2348, 2021. doi: https://doi.org/10.1002/esp.5179.

Sendrowski, A., Wohl, E., Hilton, R., Kramer, N., and Ascough, P. Wood-Based Carbon Storage in the Mackenzie River Delta: The World's Largest Mapped Riverine Wood Deposit. *Geophysical Research Letters*, 50(7):e2022GL100913, 2023. doi: https://doi.org/10.1029/2022GL100913. e2022GL100913 2022GL100913.

Sharma, C., Singh, P. S., and Shenoy, A. Performance Analysis of Object Detection Algorithms on YouTube Video Object Dataset. 2021, https://api.semanticscholar.org/CorpusID:237278049.

Spreitzer, G., Tunnicliffe, J., and Friedrich, H. Using Structure from Motion photogrammetry to assess large wood (LW) accumulations in the field. *Geomorphology*, 346:106851, 2019. doi: https://doi.org/10.1016/j.geomorph.2019.106851.

Tamminga, A., Hugenholtz, C., Eaton, B., and Lapointe, M. Hyperspatial Remote Sensing of Channel Reach Morphology and Hydraulic Fish Habitat Using an Unmanned Aerial Vehicle (UAV): A First Assessment in the Context of River Research and Management. *River Research and Applications*, 31(3):379–391, 2015. doi: https://doi.org/10.1002/rra.2743.

Tassielli, G., Notarnicola, B., Renzulli, P. A., De Molfetta, M., and Fosco, D. The Use of Unmanned Aerial Systems in Environmental Monitoring. In *Innovation, Quality and Sustainability for a Resilient Circular Economy*, edited by Lagioia, G., Paiano, A., Amicarelli, V., Gallucci, T., and Ingrao, C., pages 459–465, Cham, 2024. Springer International Publishing.

Ulloa, H., Iroumé, A., Mao, L., Andreoli, A., Diez, S., and Lara, L. E. Use of remote imagery to analyse changes in morphology and longitudinal large wood distribution in the blanco river after the 2008 chaitén volcanic eruption, southern chile. *Geografiska Annaler: Series A, Physical Geography*, 97(3):523–541, 2015. doi: https://doi.org/10.1111/geoa.12091.

Varghese, R. and M., S. YOLOv8: A Novel Object Detection Algorithm with Enhanced Performance and Robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6, 2024. doi: https://doi.org/10.1109/ADICS58448.2024.10533619.

Verdonschot, P. F. M. and Verdonschot, R. C. M. Ecological Functions and Management of Large Wood in Fluvial Systems. *Current Forestry Reports*, 10:39–55, 02 2023. doi: https://doi.org/10.1007/s40725-023-00209-x.

Wohl, E., Cenderelli, D. A., Dwire, K. A., Ryan-Burkett, S. E., Young, M. K., and Fausch, K. D. Large in-stream wood studies: a call for common metrics. *Earth Surface Processes and Landforms*, 35(5): 618–625, 2010. doi: https://doi.org/10.1002/esp.1966.

Wohl, E., Lininger, K. B., Fox, M., Baillie, B. R., and Erskine, W. D. Instream large wood loads across bioclimatic regions. *Forest Ecology and Management*, 404(May):370–380, 2017. doi: https://doi.org/10.1016/j.foreco.2017.09.013.

Wohl, E., Scott, D. N., and Lininger, K. B. Spatial Distribution of Channel and Floodplain Large Wood in Forested River Corridors of the Northern Rockies. *Water Resources Research*, 54(10): 7879–7892, 2018. doi: https://doi.org/10.1029/2018WR022750.

Wohl, E., Kramer, N., Ruiz-Villanueva, V., Scott, D. N., Comiti, F., Gurnell, A. M., Piegay, H., Lininger, K. B., Jaeger, K. L., Walters, D. M., and Fausch, K. D. The natural wood regime in rivers. *BioScience*, 69(4):259–273, 2019. doi: https://doi.org/10.1093/biosci/biz013.

Zapryanov, G., Ivanova, D., and Nikolova, I. Automatic White Balance Algorithms for Digital Still Cameras - a Comparative Study. *Information Technologies and Control*, 1:16–22, 01 2012.

Zhao, Z., Zheng, P., Xu, S.-t., and Wu, X. Object detection with deep learning: A review. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11):3212–3232, 2019. doi: https://doi.org/10.1109/TNNLS.2018.2876865.