



OPEN Detection of anomalous activities around telecommunications infrastructure based on YOLOv8s

Enerst Edozie✉, Aliyu Nuhu Shuaibu, Ukagwu Kelechi John & Bashir Olaniyi Sadiq

This study explores the deployment of YOLOv8s for detecting anomalies in fiber optic cables mounted on poles, with a focus on climbing activities and environmental impediments. To address the lack of climbing-related annotations in current datasets, a custom dataset was generated, covering a variety of scenarios to enhance model adaptability. During training, various augmentation approaches were used, which greatly enhanced model performance and reduced overfitting. The proposed model was trained and tested over many epochs, with detection performance progressively improving: mAP@0.5 increased from 78.9% at 20 epochs to 87.5% at 50 epochs and 97.3% at 100 epochs, after which further increases plateaued. In comparison, the trained YOLOv8s-modified model outperformed the other models on all key metrics. It achieved a mAP@50 of 97.3% and a mAP@50:95 of 71.5%, outperforming YOLOv8-original at 89.6% and 59.0%, respectively. Additionally, it achieved higher precision (96.9%) and recall (86.6%), demonstrating superior detection accuracy, dependability, and robustness in detecting complex anomalies. These improvements due to model backbone optimizations and the usage of a well-balanced, scenario-rich dataset. This study shows that YOLOv8s is a highly accurate and efficient method for detecting anomalies in fiber optic infrastructure, making it appropriate for real-time deployment in operational field environments.

Keywords YOLOv8s, Anomaly detection, YOLOv5s, Fiber optic cables, Custom dataset

Abbreviations

AI	Artificial intelligence
CNNs	Convolutional neural networks
YOLO	You only look once
OTDR	Optical time domain reflectometer
DCAE	Deep convolutional autoencoder
BiGRU	Bidirectional gated recurrent unit
BiLSTM	Bidirectional long short-term memory
DCGAN	Deep convolutional generative adversarial network
DAS	Distributed acoustic sensing
BCE	Binary cross entropy
C2F	Cross stage partial with 2 and fusion
CBS	Conv + BatchNorm + SiLU
DFL	Distribution focal loss
SiLU	Sigmoid linear unit
SPPF	Spatial pyramid pooling fast
CIoU	Complete intersection over union

The fast construction of fiber optic networks has greatly improved global telecommunications infrastructure, allowing for high-speed data transmission with low latency¹. Aerial fiber optic cables installed on poles are frequently employed because of their low cost and ease of installation. However, these cables are susceptible to a variety of environmental and human-induced anomalies, such as sagging, unintentional cutting, vandalism, and illegal tampering². Such abnormalities can cause network outages, service degradation, and higher maintenance costs, posing significant problems to telecom operators. Traditional monitoring methods, such as human inspections and rule-based surveillance systems, are frequently inefficient, time-consuming, and reactive rather than proactive³. This emphasizes the importance of modern automated systems to assure the reliability and

Department of Electrical Engineering, Kampala International University, 20000 Ishaka, Uganda. ✉email: edozie.enerst.11275@studwc.kiu.ac.ug

security of fiber optic networks. Artificial intelligence (AI), specifically deep learning-based object identification models, has emerged as an effective method for detecting anomalies in infrastructure monitoring⁴. Convolutional Neural Networks (CNNs) have shown great accuracy in object categorization, but they frequently struggle with real-time processing demands and complicated backdrops⁵. You Only Look Once (YOLO) is a deep learning-based object recognition model that has gained popularity due to its capacity to recognize objects with high speed and precision in real-time applications^{6,7}. With recent developments in YOLO models, YOLOv8s now provides increased accuracy, faster inference time, and better feature extraction, making it a perfect choice for detecting anomalies in fiber optic cables mounted on poles^{8,9}. Several research have investigated deep learning-based surveillance systems for monitoring infrastructure security, including applications for electricity lines, railway tracks, and pipeline monitoring^{10,11}. However, limited study has concentrated on identifying anomalies in aerial fiber optic cables. Traditional approaches, such as Optical Time-Domain Reflectometry (OTDR) and thermal imaging, necessitate specialized equipment and frequently do not provide continuous real-time monitoring¹². Integrating AI-powered computer vision techniques into existing monitoring systems will improve detection capabilities and give a proactive approach to fiber optic network maintenance¹³.

This study presents a deep learning architecture that uses YOLOv8s-modified to detect anomalies in fiber optic cables. The model is trained on a dataset containing images of fiber optic cables in various states, including normal, sagging, detached, and manipulated. The system is designed for deployment on edge computing devices such as NVIDIA Jetson to ensure real-time monitoring with low latency. The proposed method uses AI-driven anomaly detection to reduce manual inspection efforts, lower operating expenses, and increase the overall resilience of fiber optic networks. Furthermore, this study compares the performance of proposed model to other object detection models, examining key parameters like as precision, recall, and inference time. The findings of this study contribute to the advancement of AI-powered telecoms infrastructure monitoring by providing a scalable and cost-effective approach for ensuring the reliability of aerial fiber optic networks. Therefore, increasing dependence on fiber optic networks needs the use of intelligent monitoring technologies to efficiently manage operational difficulties. This research emphasizes the importance of deep learning in telecommunications infrastructure security and proposes a complete methodology for detecting anomalies with YOLOv8s-modified. The proposed approach has the potential to revolutionize fiber optic cable monitoring by increasing network reliability and decreasing downtime, paving the path for future AI-driven breakthroughs in the telecoms sector. The contributions of this paper are as follows:

- This paper implemented extensive data augmentation techniques to simulate diverse weather and lighting conditions, enhancing the model's generalization to real-world deployment challenges.
- This research presents a YOLOv8s-modified based model specifically designed for anomaly detection in aerial fiber optic cables, offering enhanced speed, precision, and real-time detection capability.
- This study introduces a custom-labeled dataset of fiber cable anomalies, including climbing activities, poles, and person and animal, captured from field environments to improve model training and validation.

The rest of the paper is structured as follows: The remainder of this article is organized as follows. The Section “[Related works](#)” addresses current advances and existing approaches to deep learning-based anomaly detection for fiber optic cables. Section “[Methodology](#)” describes the YOLOv8s-based detection model and system architecture. The Section “[Experimental environment and dataset setup](#)” describes the dataset used, the experimental setup, and the model training techniques. Section “[Results](#)” displays and analyzes performance metrics such as accuracy and inference time. Section “[Conclusion](#)” summarizes the research findings and proposes prospective areas for further investigation.

Related works

Several studies have looked into the use of deep learning in infrastructure anomaly detection, particularly in telecommunications and electric power networks. In order to analyze noisy OTDR traces for precise fault identification and localization in fiber optic cables, Abdelli et al.¹⁴ used a denoising convolutional autoencoder (DCAE) with Bidirectional Long Short-Term Memory (BiLSTM). Even in situations with a poor signal-to-noise ratio, the model's accuracy was 96.7%. A CNN-based technique for identifying reflected events in optical fiber networks utilizing noisy OTDR signals was presented by Abdelli et al.¹⁵. When compared to conventional techniques, their model demonstrated better detection performance and fewer false alarms. A multi-task deep learning model employing LSTM was developed by Abdelli et al.¹⁶ to identify, locate, and estimate reflectance of fiber optic defects. In noisy settings, the model continued to perform well, particularly when identifying reflected flaws.

Natalino et al.¹⁷ developed a framework for detecting spectrum anomalies using unsupervised deep learning and computer vision. Their technique detected errors in optical transmission spectra with 100% accuracy without using labeled training data. Wang and Zhang¹⁸ used CNN model for detecting anomalies in fiber optic cables, which outperformed prior versions in real-time. However, their research concentrated mostly on lab-based testing, leaving real-world validation restricted.

Abdelli et al.¹ proposed a machine learning-driven monitoring system for optical networks that uses autoencoders and attention-based Bidirectional Gated Recurrent Unit (BiGRU) models to detect and classify anomalies. Their work focuses on lightweight inference to enable scalable deployment in real-time monitoring. Rizzo et al.¹⁹ proposed a deep learning architecture based on Faster R-CNN to recognize known and unknown events in OTDR recordings. The model was highly accurate in detecting many types of anomalies, including previously unknown flaws. Huot et al.²⁰ created a CNN-based model for detecting microseismic events in Distributed Acoustic Sensing (DAS) data obtained from fiber-optic networks. Their method has a detection accuracy of 98.6%, demonstrating the promise of deep learning in seismic anomaly monitoring.

Zhang et al.²¹ proposed a hybrid deep learning framework that combines Deep Convolutional Generative Adversarial Networks (DCGAN) and Convolutional Neural Networks (CNN) to detect and localize faults in fiber optic cables. Their model exhibited remarkable accuracy (98.5%) in controlled test scenarios, indicating potential for real-world applications. Ghamisi et al.²² proposed an unsupervised deep learning strategy for detecting anomalies in Automated Fiber Placement systems. The approach eliminates the need for labeled data and shown usefulness in detecting fiber flaws during composite material manufacture. Zhang et al.²³ performed a neural network-based study to anticipate defects in fiber optic cables used in power distribution communication networks. Their model displayed forecast accuracy and dependability, which helped reduce network downtime in smart grid networks.

This area of study has been expanded further by recent 2025 studies. With a focus on the model's lightweight structure and quick inference, Sheng et al.²⁴ presented an enhanced YOLOv8s-based model for real-time classroom behavior monitoring. This contribution may be applicable to infrastructure surveillance scenarios. A contextual information-based anomaly detection framework for multi-scene aerial videos was presented by G et al.²⁵ to handle scene variation issues that are pertinent to wide-area surveillance. Road anomaly identification in remote sensing was accomplished by Khatoun et al.²⁶ by combining fuzzy logic and TinyML, advancing edge-enabled intelligent monitoring. To detect walkway irregularities in intelligent transportation systems, Alotaibi et al.²⁷ used deep learning in conjunction with optimization approaches. This idea can be extended to telecom pedestrian safety monitoring. Veeram et al.²⁸ developed a robust anomaly detection model based on temporal graph attention and transformer-augmented RNNs, demonstrating enhanced spatiotemporal feature learning essential for dynamic event detection in telecom infrastructures. Finally, Liu et al.²⁹ introduced a deep learning-based method for identifying abnormal behaviors and items in public buses, suggesting feasible cross-domain applications in anomaly recognition.

The YOLO (You Only Look Once) models have evolved quickly, influencing real-time object identification tasks. Beginning with Redmon's et al. YOLOv1³⁰, which introduced a one-stage detection technique, the model faced issues in localization accuracy and small object recognition. YOLOv2⁷ and YOLOv3³¹ improved performance with multi-scale predictions, anchor boxes, and batch normalization. YOLOv4³² and YOLOv5³³ significantly improved detection speed and usability, with YOLOv5 gaining popularity due to its PyTorch-based implementation and excellent performance on customized datasets. The adaptability of YOLOv5 has been proved in a variety of infrastructure-related applications. Wang et al.³⁴ used UAVs and YOLOv5 to create a rebar-counting model for reinforced concrete columns, and Wang³⁵ investigated the effect of picture augmentation on window identification tasks. These examples demonstrate the applicability of YOLO models in real-world engineering scenarios.

In addition to these developments, YOLOv6³⁶ and YOLOv7³⁷ added more performance enhancements. With an anchor-free design, a redesigned backbone network, improved loss functions, and ONNX/TensorRT export compatibility, Ultralytics' 2023 release of YOLOv8³⁸ represented a significant architectural change that improved detection accuracy and deployment efficiency, especially for edge computing. Transformer-enhanced models, such as YOLOv10³⁹, have surfaced more lately. For instance, Eum et al.⁴⁰ used attention-based transformer modules to show how well YOLOv10 works for recognizing heavy machinery in crowded construction sites. Although YOLOv9⁴¹ and YOLOv10³⁹ have emerged and provide outstanding performance on challenging tasks, YOLOv8 balances detection accuracy, model size, and inference speed—all of which are important considerations for deployment in resource-constrained or rural landscapes. The smallest and quickest variant in the YOLOv8 family, YOLOv8s, is therefore used in this investigation. It is designed to identify small, mobile objects, like people and animals, in real time close to fiber optic infrastructure.

This study extends previous deep learning-based anomaly detection measures by using the efficiency and robustness of YOLOv8s. It provides a scalable, real-time solution for improving surveillance and operational monitoring of pole-mounted telecommunications equipment by incorporating data-driven optimization methodologies and reacting to changing environmental circumstances. Furthermore, by combining YOLOv8s with edge computing, the study addresses key limitations of prior research, such as limited real-time efficiency, environmental sensitivity, and large-scale deployability, thus offering a practical path toward intelligent infrastructure monitoring in low-resource environments.

In Table 1, a range of anomaly detection approaches across telecommunications and related domains are compared in terms of techniques, contributions, limitations, and improvements. The evaluated works range from CNNs, LSTMs, and autoencoder-based models to modern YOLO frameworks, demonstrating progress toward efficient, accurate, and real-time solutions, with the proposed YOLOv8s-modified being the most recent step forward.

Methodology

This study adopts the YOLOv8s deep learning model to detect and categorize abnormal activities such as human or animal presence near pole-mounted optical fiber infrastructure. The nature of the data, which consists of unstructured, high-resolution images and video frames obtained using static or mobile surveillance devices in outdoor areas, justifies the use of deep learning. Traditional machine learning models frequently require hand-engineered features and perform poorly under changing illumination conditions, background noise, or when objects appear in different orientations or scales. Deep learning, particularly CNN-based architectures such as YOLOv8, automatically extracts hierarchical characteristics, allowing it to generalize effectively across various anomaly environments. YOLOv8s, an anchor-free model with enhanced detection heads and augmented loss functions, has achieved world-class performance in real-time object detection tasks. Its lightweight design and high inference speed make it suitable for use on embedded field surveillance systems. Previous empirical investigations (e.g., Wang²⁴, Eum et al.²⁵; Sheng et al.²⁵; Girisha et al.⁴⁰) have validated the practical usefulness of YOLO models in high-stakes visual monitoring scenarios, which influenced this methodological decision. Thus,

Author	Technical details	Contribution	Limitations	Suggested improvements
Abdelli et al. ¹⁴	DCAE + BiLSTM on OTDR traces	96.7% fault localization accuracy under low SNR	Limited to cable fault detection	Extend to broader telecom anomaly scenarios
Abdelli et al. ¹⁵	CNN on OTDR signals	Improved reflected-event detection, fewer false alarms	Focused on specific event type	Generalize to diverse anomaly categories
Abdelli et al. ¹⁶	Multi-task LSTM	Simultaneous detection, localization & reflectance estimation	Controlled lab setting	Field validation under varied conditions
Abdelli et al. ¹	Autoencoder + attention BiGRU	Lightweight real-time anomaly detection	Optical-only focus	Apply to physical infrastructure surveillance
Rizzo et al. ¹⁹	Faster R-CNN	Accurate detection of known + unknown events	Computationally heavy	Optimize for real-time edge deployment
Huot et al. ²⁰	CNN on DAS data	98.6% accuracy in microseismic anomaly detection	Non-telecom domain	Adapt to telecom event monitoring
Zhang et al. ²¹	DCGAN + CNN hybrid	98.5% accuracy in fault detection & localization	Controlled scenario	Validate in outdoor/field settings
Yang et al. ⁴²	Improved YOLOv5s for wildlife detection	Demonstrated lightweight, adaptable detection algorithm in unstructured environments	Lower recall (63.9%) and weaker accuracy on small/complex objects	Enhance sensitivity and optimize for infrastructure-focused anomalies
Sakiba et al. ⁴³	YOLOv7 + ConvLSTM for real-time crime detection	Achieved higher precision (86.4%) with improved real-time detection capability	Recall and mAP limited under complex scenes	Refine feature learning to improve recall and anomaly coverage
Zhang et al. ²³	Neural networks	Reliable defect forecasting	Dataset limited	Larger dataset, broader infra contexts
Veesam et al. ²⁸	Temporal graph attention + Transformer-RNN	Captures complex spatiotemporal anomalies	Model complexity for edge	Develop lightweight transformer hybrids
Ours (2025)	Modified YOLOv8s with lightweight backbone (48–768 channels), reduced C2f. repetitions, SPPF(768), and LSKAttention	mAP@0.5 = 97.3%. Faster inference, reduced computational cost, improved adaptability to varied object scales & complex scenes. High robustness for edge devices and real-world deployments	–	Extensible to multimodal sensing and wider smart-city infra-applications

Table 1. The comparison of previous works on deep learning-based anomaly detection in telecoms infrastructure and related applications.

YOLOv8s offers a well-suited and experimentally supported strategy to detecting anomalies surrounding key telecommunications infrastructure.

YOLOv8s model for anomaly detection

Overview of YOLOv8s

YOLOv8 (You Only Look Once version 8) is a cutting-edge one-stage object detection method that transforms the object detection task into a regression issue, considerably increasing speed and accuracy over prior versions³⁰. YOLOv8 outperforms other object detection frameworks in terms of detection speed and precision, making it ideal for real-time applications, particularly those that need microscopic object detection^{44–47}. To meet a variety of deployment requirements, the YOLOv8 architecture provides several model variants: YOLOv8n (nano), YOLOv8s (small), YOLOv8m (medium), YOLOv8l (large), and YOLOv8x (extra-large). YOLOv8s is a lightweight, real-time object detection model developed by Ultralytics, designed for speed and efficiency without significantly compromising accuracy. It is part of the YOLOv8 family and uses a streamlined architecture featuring a modular design: a backbone for feature extraction (typically based on C2f. blocks), a neck for feature fusion (like PAN or FPN structures), and a head for final object classification and localization. YOLOv8s supports a variety of tasks including detection, segmentation, and pose estimation, and is optimized for deployment on edge devices due to its small size and fast inference speed. The following subsection provides an in-depth discussion of the architectural components of both the original YOLOv8s and the modified YOLOv8s.

Architectural components of YOLOv8s

YOLOv8s has several advanced modules that dramatically improve detection efficiency and accuracy. The CBS module (Convolution, Batch Normalization, SiLU Activation) is at the heart of the system's ability to extract hierarchical image features as shown in Fig. 1a. The convolutional layers capture spatial information, batch normalization keeps training stable by normalizing feature distributions, and the SiLU activation function improves gradient flow for enhanced feature learning⁴⁸. This combination enables YOLOv8s to detect things in complex surroundings, including fiber optic poles, climbing activity, and other anomalies. Furthermore, the SPPF module (Spatial Pyramid Pooling Fusion) enhances detection accuracy by collecting and combining high-level characteristics from various receptive fields⁴⁹. SPPF uses multi-scale pooling to ensure that objects of varied sizes from little details like climbing activities to massive structures like utility poles are detected with high precision. Another significant improvement to YOLOv8s is the C2F module (C3-inspired Lightweight Module with ELAN Concepts), which improves computational efficiency while maintaining detection performance³⁷. This module decreases computational effort by reusing feature maps, reducing gradient dispersion, and freeing up deeper layers for sophisticated feature extraction. As a result, YOLOv8s strikes the ideal balance between speed and accuracy, making it extremely useful for real-time fiber optic anomaly detection. Furthermore, the use of a decoupled detecting head in YOLOv8s differs significantly from prior YOLO architectures⁵⁰. This method isolates classification tasks (predicting object classes) from bounding box regression (localizing objects), eliminating conflicts and improving precision, particularly when detecting small and occluded objects in challenging environments.

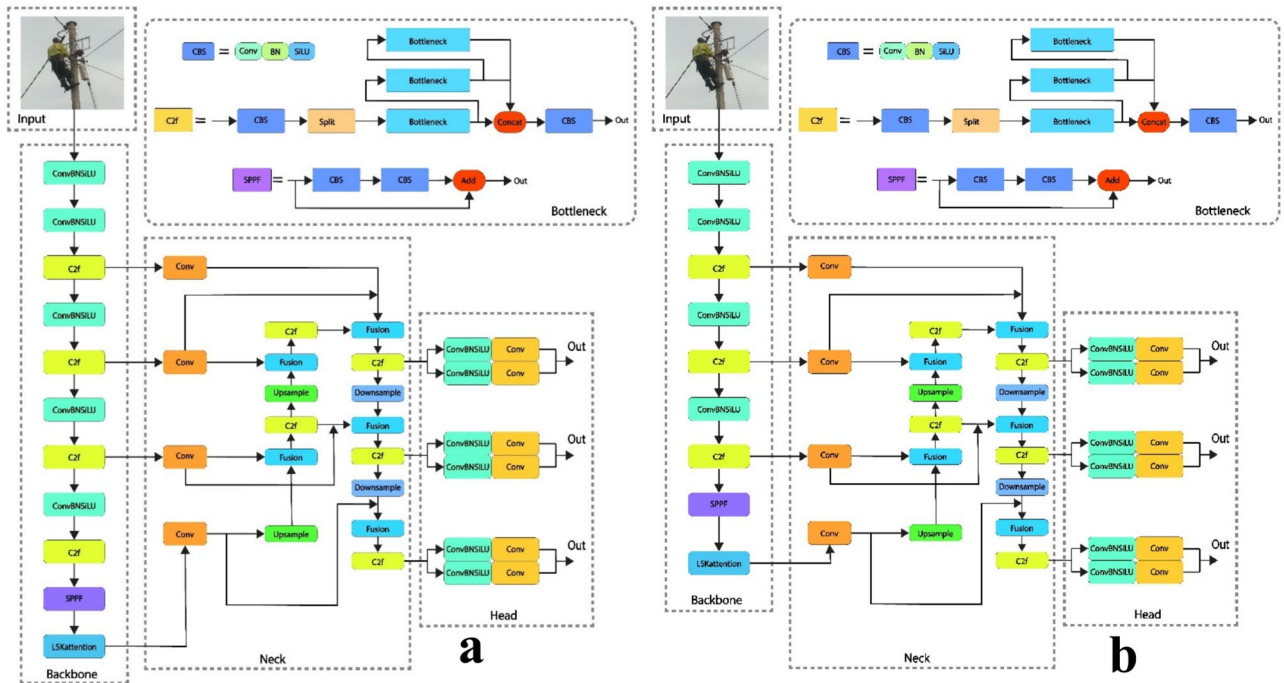


Fig. 1. The structure of (a) YOLOv8s-baseline model and (b) YOLOv8s-modified.

This modified YOLOv8s backbone is designed for the purpose to increase inference speed, reduce computational cost and improve adaptability when training on custom-datasets particularly those ones with varied object scales and complex visual structures. It starts with two lightweight convolutional layers using 48 and 96 channels, which are reduced from the original 64 and 128 to minimize early-stage memory and processing overhead. The architecture continues with progressively increasing channel sizes at each spatial stage (192, 384, and 768 for P3, P4, and P5, respectively), which allows for efficient hierarchical feature extraction while keeping the model lightweight. The number of C2f (Cross-Stage Partial Fusion) repetitions is reduced to 3 in the backbone stage as shown in Fig. 1b, cutting down parameter count and computational load without compromising the richness of features. These architectural choices strike a strong balance between accuracy and efficiency, making the model highly suitable for real-time applications on edge devices or GPU-constrained environments. The key improvement in this backbone is the introduction of a lower SPPF module with 768 channels, down from 1024 in the original YOLOv8, which takes contextual information from different receptive fields at a low additional cost. Importantly, this is followed by an LSKAttention block also reduced to 768 channels, which provides a large-kernel attention mechanism that selectively focuses on the most interesting spatial regions while reducing irrelevant background noise. This attention modification greatly improves the model’s generalizability, particularly when trained on custom or imbalanced datasets that lack the constancy of large-scale benchmarks. However, this backbone not only runs faster and lighter than the default YOLOv8s, but it also has higher detection accuracy and robustness in domain-specific or real-world deployments. The next section presents a detailed discussion of the loss functions used in YOLOv8s.

Loss functions in YOLOv8s

YOLOv8s uses a combination of advanced loss functions to optimize training stability and detection accuracy. These loss functions ensure that the model learns to make exact classifications, generate correct bounding boxes, and maximize confidence ratings for discovered objects⁵¹.

BCE loss for classification

The YOLOv8s model uses Binary Cross-Entropy (BCE) Loss for classification to maximize the probability estimation of each observed object class. BCE Loss is commonly employed in deep learning for binary and multi-class classification tasks because of its ability to adequately handle probability distributions^{9,52}.

BCE loss function definition The BCE loss function is defined in Eq. (1):

$$L_{BCE} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \tag{1}$$

where: y_i is the true class label (0 or 1). \hat{y}_i is the predicted probability for class 1. N is the total number of predictions.

BCE calculates the logarithmic difference between the true label and the anticipated probability, which penalizes inaccurate predictions more heavily⁵³. This guarantees that the model assigns high confidence scores to right classifications and discourages uncertain predictions. In YOLOv8, BCE Loss is used in object classification tasks to ensure that each spotted object is allocated the best accurate class probability. Because YOLOv8s uses a decoupled detection head, classification and regression tasks are optimized individually, resulting in higher precision. The model learns to distinguish between many object classes (for example, poles, humans, animals, and climbing activities) by refining class probability distributions via BCE Loss.

CIOU loss and DFL for regression

- In YOLOv8s, CIOU Loss (Complete Intersection over Union Loss) optimizes bounding box regression by taking into account overlap, distance, and aspect ratio disparities between predicted and ground-truth bounding boxes⁵⁴. This improves on normal IoU loss. The CIOU ensures that the model correctly localizes objects with little misalignment.
- DFL (Distribution-Focal Loss) improves bounding box localization by learning a probability distribution of alternative placements⁵⁵. In addition to CIOU, DFL loss is used to increase the precision of bounding box coordinates. Instead of anticipating fixed numerical values for bounding box positions, DFL depicts geographical coordinates using a probability distribution. This enables the model to modify predictions using probability density, improving localization accuracy in highly changeable situations⁵⁶. The DFL function is expressed in Eq. (2):

$$DFL(s_i, s_{i+1}) = -((y_{i+1} - y) \log(s_i) + (y - y_i) \log(s_{i+1})) \quad (2)$$

where: s_i represents the network's sigmoid output. y is the label. y_{i+1}, y_i are interval orders.

Anchor-free object detection in YOLOv8s

Unlike classic anchor-based YOLO models, YOLOv8 takes an anchor-free method, which improves generalization while lowering computational costs⁵⁷. The model employs a Task-Aligned Assigner for dynamic matching, calculating anchor-level alignment as shown in Eq. (3):

$$t = s^\alpha \times u^\beta \quad (3)$$

where s is the classification score, u is the IoU value, and α and β are weight hyperparameters. Choosing the best anchors maximizes training efficiency while minimizing false positives.

Object detection process in YOLOv8s

The YOLOv8s model employs a structured object detection pipeline to provide rapid and accurate detection of anomalies in fiber optic cables mounted on poles. This pipeline is divided into numerous essential stages, starting with image preparation and ending with the final prediction output. Each stage has been extensively adjusted to improve detection speed and accuracy, making it ideal for real-time monitoring applications.

Input processing

Before feeding images into the YOLOv8s model, they are preprocessed to standardize input dimensions and improve detection consistency. Images are scaled to 640×640 pixels to ensure consistent input sizes across datasets. This phase is critical for maintaining computational efficiency, lowering processing load, preventing fluctuations in object scale, assuring constant feature extraction, and facilitating batch training for deep learning models with fixed input dimensions. Also, YOLOv8s normalizes pixel values to a range of 0 to 1, ensuring stable gradients throughout training and increasing convergence speed.

Feature extraction using the backbone module

Once the input image has been processed, it is routed through the Backbone module, which is in charge of extracting multi-scale feature maps. To capture important spatial and semantic information, the Backbone uses convolutional layers, batch normalization, and activation functions. This module detects key aspects of poles, persons, animals, and climbing activities, even in low-light conditions or when occlusions are present. The Backbone extracts feature maps at various resolutions, letting the model detect small and large objects. Spatial Pyramid Pooling Fusion (SPPF) is used to preserve high-level semantic features while increasing classification accuracy.

Detection head processing

The Detection Head module processes the retrieved feature maps to provide object bounding boxes, class probabilities, and confidence scores. This stage has a decoupled head structure that separates classification and regression operations, resulting in better bounding box predictions.

The detection procedure involves predicting object existence in grid cells.

- Refined bounding box coordinates to match actual object placements.
- Assigning class labels using probability distributions.
- Filtering predictions with confidence ratings and non-maximum suppression (NMS) to eliminate redundant detections.

The model automatically adjusts bounding box sizes based on object properties, ensuring precise localization of anomalies near fiber optic cables.

Prediction output and final detection

The final prediction output is represented as a $B \times B \times C$ feature map, where:

- $B \times B$ defines the spatial resolution of the detection grid.
- C is the total number of prediction parameters, as shown in Eq. (4):

$$C = n \times (5 + N) \quad (4)$$

where: n = Number of bounding boxes per grid cell (typically 3), 5 = Four bounding box coordinates (x , y , width, height) + confidence score, and N = Number of object classes (e.g., poles, humans, animals, climbing activities).

Each grid cell predicts multiple bounding boxes and assigns a confidence score that indicates the probability of an object's presence. Non-Maximum Suppression (NMS) improves the results by eliminating overlapping detections and picking the best accurate bounding box for each object.

Experimental environment and dataset setup

Dataset description

The data used in this study was collected in Bushenyi-Ishaka, an urban–rural area with a variety of environmental conditions. The dataset focuses on abnormalities surrounding fiber optic cables mounted on poles, and it includes images of poles, humans, animals, and climbing activities. The goal is to create a comprehensive detection model that can identify unauthorized access, potential damage hazards, and other security concerns in a variety of environments.

Dataset production

To provide a diverse and comprehensive dataset, random images were collected at various times of day and weather conditions, including sunny, cloudy, and rainy conditions. The data was collected using high-resolution surveillance cameras strategically placed throughout fiber optic routes. Data was collected in a variety of settings, including metropolitan streets, rural areas, and solitary pole installations, to improve the model's generalizability. The combination of these diverse scenarios ensures that the proposed model learns to discriminate abnormalities effectively, hence increasing its robustness and dependability. Figure 2 shows a representative selection of image samples from the obtained dataset, demonstrating various conditions and anomaly kinds observed during the data collection procedure.

Image labeling

The obtained images were manually labeled with LabelImg⁵⁸, an open-source graphical image annotation tool commonly used to label datasets in object detection applications. This procedure was critical to ensure that the improved YOLOv8s model could reliably recognize and classify target objects such as poles, persons, animals, and climbing activities. Labeling entails creating exact bounding boxes around each object of interest in the image, while ensuring that each labeled instance provides explicit positional data to the model. The bounding boxes were altered to securely surround objects while eliminating unwanted background inclusion, enhancing the model's feature extraction capabilities.

Class label assignment and dataset consistency

Once the bounding boxes had been created, each object was allocated an appropriate class label. A consistent labeling strategy was used to avoid differences that could have a negative impact on model training. Humans engaged in climbing, for example, were clearly identified in order to distinguish between typical human presence and potential security hazards. Furthermore, several quality control checks were carried out during the annotating process. These tasks involved manually examining labeled images, correcting any misclassified objects, and ensuring that overlapping or obscured objects were correctly categorized. This refinement approach improved dataset consistency and reduced annotation errors that could harm model performance. The completed annotated dataset was saved in YOLO format, which is a lightweight annotation style that represents each labeled object with a text file providing normalized bounding box coordinates and class labels. This format offers direct compatibility with the proposed model, allowing for smooth integration throughout the training process. The dataset was optimized using the YOLO annotation structure for fast loading, low computational overhead, and increased detection accuracy. These measures improved the model's ability to generalize across different environmental scenarios, making it extremely successful for detecting anomalies in fiber optic infrastructure monitoring.

Dataset splitting

To train the YOLOv8s model effectively, the dataset was split into three subsets:

- *Training set 70%* Used for learning patterns and feature extraction.
- *Validation set 20%* Used to fine-tune hyperparameters and prevent overfitting.
- *Test set 10%* Used to evaluate model performance on unseen data.



Fig. 2. The image samples from the dataset.

Dataset training

Training configuration

The improved YOLOv8s model was trained and evaluated on a machine powered by an Intel Core i7-6700HQ CPU running at 2.60 GHz, providing a solid basis for general computing and model training activities. The input image resolution was set to 640×640 pixels, a standard for YOLOv8 models. This ensures that the model can analyze and extract characteristics from objects of varied sizes. The system was built with 24 GB of RAM to ensure effective data handling and computational efficiency, particularly when working with huge image datasets and deep learning models. For storage, the system used a 500 GB SSD, which considerably improves data read/write performance, allowing for smooth dataset processing and model checkpoint saving. The deep learning framework utilized was PyTorch, a popular and adaptable library for deep learning applications. The development environment was set up with Python 3.9, which is compatible with a variety of deep learning frameworks and tools. The model training used an NVIDIA Quadro 1000 M GPU, which, while not the most recent in deep learning acceleration, has dedicated CUDA cores for accelerating computation-intensive operations. To expedite the workflow and manage dependencies efficiently, the current version of Anaconda was utilized as the package manager, allowing for smooth environment setup and library installation.

Evaluation metrics

The performance of the YOLOv8s-modified model is measured using commonly established object detection measures such as Precision, Recall, Average Precision (AP), Mean Average Precision (mAP), and the F1-score^{59,60}. These metrics evaluate the model's accuracy, dependability, and robustness in detecting key objects such as poles, persons, animals, and climbing activities near fiber optic infrastructure in a variety of environmental circumstances.

Precision is measured as the fraction of successfully detected positive detections among all anticipated positives, indicating the model's ability to reduce false alarms. The mathematical definition is given in Eq. (5):

$$P = \frac{TP}{TP + FP} \quad (5)$$

Recall quantifies the model's capability to recognize true positive cases, demonstrating how well it decreases missed anomalies. It's computed in Eq. (6):

$$R = \frac{TP}{TP + FN} \quad (6)$$

AP describes the precision-recall trade-off for a single class, usually determined at a conventional Intersection over Union (IoU) threshold of 0.5. Equation (7) provides the formula for calculating AP for a single category.

$$AP = \int_0^1 P(R) dR \quad (7)$$

mAP is the average of AP scores across all target classes, providing a comprehensive measure of detection accuracy and localization performance as shown in Eq. (8).

$$mAP = \frac{1}{N} \sum_{i=1}^N AP_i \quad (8)$$

The F1-score, as the harmonic mean of precision and recall, provides a balanced single metric that is especially valuable in anomaly detection scenarios where both false positives and false negatives have large costs. It is expressed in Eq. (9):

$$F1 = \frac{2 \times P \times R}{P + R} \quad (9)$$

These measures are particularly important in this study because they directly indicate the model's ability to consistently detect anomalous activities such as illegal human presence or animal interference while limiting false warnings that could result in wasteful operational responses. The tight assessment criteria (e.g., mAP@0.5:0.95) further ensure that detections are spatially accurate, which is crucial for precisely localizing anomalies on pole-mounted fiber optic cables. By focusing on these criteria, the study emphasises not just detection accuracy but also the practical application of YOLOv8s-modified in real-world telecom infrastructure monitoring scenarios.

Results

Model performance analysis

The trained YOLOv8s-modified model demonstrated high detection capabilities, notably in detecting anomalies related to fiber optic cables mounted on poles. The model attained a mAP@0.5 of more than 85%, indicating that it can reliably recognize and classify items such as poles, humans, animals, and climbing activities while minimizing false positives and false negatives. This great precision implies that the model can consistently distinguish between normal and unusual activities, making it ideal for telecom infrastructure monitoring. However, the mAP@0.5:0.95, which assesses performance across several intersection-over-union (IoU) criteria, demonstrated a progressive deterioration. This tendency is expected since tougher IoU thresholds necessitate higher localization precision, which makes object detection challenging. Figure 3a illustrates the precision-recall curve, demonstrating that the model effectively maintains a decent balance between accurately identified anomalies and erroneous detections. When comparing YOLOv8s-modified to other state-of-the-art models shown significant increases in detection accuracy, resilience, and computational efficiency. Figure 3b shows the precision-confidence curve and which indicates that YOLOv8s-modified maintains a greater degree of precision across different confidence levels, implying that its high-confidence predictions are more dependable and contain fewer false positives. Furthermore, Fig. 3c presents the F1-confidence curve, which assesses the model's balance of precision and recall. This increase can be ascribed to YOLOv8s' improved design, superior feature extraction, and optimized computing performance, which enable it to handle complicated settings such as crowded backdrops and variable illumination conditions commonly found in telecom networks. Figure 3d illustrates another critical feature of YOLOv8s' performance: the recall-confidence curve.

The curve indicates that improved YOLOv8s maintains a high recall value even at moderate confidence thresholds, showing that it consistently detects the majority of anomalies without significantly increasing false alarms. This is particularly useful for fiber optic monitoring applications, where missing an anomaly such as unauthorized climbing or cable tampering could result in costly infrastructure damage or service disruptions. The proposed model provides higher generalization, enhanced object localization, and fewer false alarms, giving it a more resilient and efficient solution for real-time telecom anomaly detection. Given the time restrictions of monitoring fiber optic cables in pole-mounted infrastructure, YOLOv8s' higher accuracy and inference speed make it the preferable choice for high-volume, real-time anomaly detection tasks in telecom networks.

Detection results

The detection results shown in Fig. 4, indicate that the YOLOv8s model has the capability of detecting anomalies in fiber optic cable systems with high accuracy and dependability. Even in complicated landscapes with varying lighting conditions and partial occlusions, the model effectively detected poles, humans, animals, and climbing activities. YOLOv8s-modified outperformed earlier models in terms of object localization and classification, resulting in less false positives and negatives. The clear boundary boxes and high-confidence detections shows the model's ability to discriminate key anomalies from irrelevant background objects. Also, the precision-recall and confidence-based curves from prior performance evaluations are consistent with these results, confirming that the model maintains a balance of high precision and recall. These findings establish YOLOv8s-modified as a dependable solution for real-time monitoring, allowing telecom operators to promptly identify and respond to possible threats to fiber optic infrastructure.

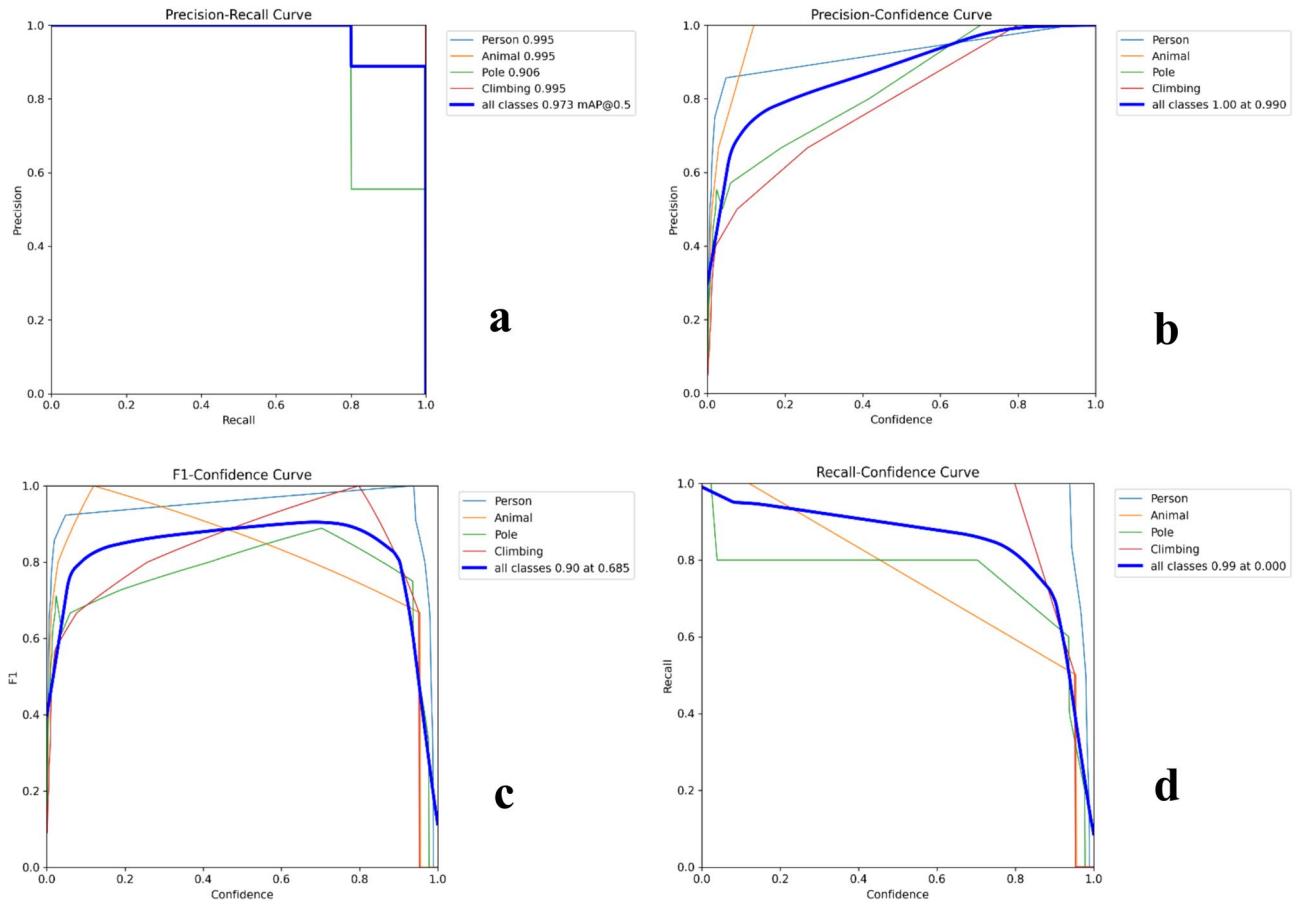


Fig. 3. (a) the precision-recall curve, (b) the precision-confidence curve, (c) the F1-confidence curve, and (d) the recall-confidence curve.




Fig. 4. Some samples of detection results.

Impact of augmentation techniques

The augmentation technique utilized in this study are widely known in the computer vision literature, but their strategic application to the novel domain of fiber optic infrastructure monitoring in outdoor telecom environments was important for improving model performance and dependability. Rather than providing

Metric	With augmentation (%)	Without augmentation (%)	Change (%)
Precision	96.9	93.4	+ 3.75
Recall	86.6	83.9	+ 3.22
mAP @ 50	97.3	87.1	+ 11.71
mAP@50:95	71.5	66.8	+ 7.04

Table 2. Comparison of YOLOv8s-modified performance metrics with and without data augmentation techniques.




Fig. 5. PR curves comparing YOLOv8s-improved performance trained with (a) augmentation, and (b) without augmentation.

methodological innovation, these strategies were purposefully selected and customized to address domain-specific environmental problems faced during real-world surveillance. Random flipping was employed to simulate orientation variance, allowing the improved YOLOv8s model to recognize humans and animals approaching fiber optic poles from various directions. This is especially true in unstructured outdoor contexts where object positioning is not stable. Gaussian Noise Addition improved the model's resilience to sensor noise and visual artifacts caused by low-light circumstances or compressed transmission feeds, which are typical in real-time field deployments. Brightness modifications adjusted for variations in natural lighting caused by time of day, shadows and weather changes. By training the algorithm on images with varying brightness levels, it enhanced detection consistency across circumstances ranging from direct sunlight to cloudy or dusk scenarios. Rotation Augmentation enabled the model to manage object orientation variability, such as animals mounting poles at unexpected angles or maintenance people assuming unique postures. This increased the model's capacity to generalize to less predictable field interactions. These enhancements significantly enhanced training data diversity, lowering overfitting and increasing generalization. While not innovative in itself, these strategies were successfully used to train an improved YOLOv8s model using a special telecom anomaly detection dataset. This is consistent with Han's et al.⁶¹ findings, which underscore the importance of environment-specific augmentation. The essential innovation is the intentional integration of standard methodologies into a domain-adapted pipeline, which allows for robust, real-time anomaly identification in demanding outdoor fiber optic monitoring circumstances.

The comparison of performance metrics with and without data augmentation techniques

The model's performance with and without data augmentation is compared in the Table 2, which shows that applying augmentation consistently improves all key metrics. Precision increased from 93.4 to 96.9%, while recall climbed from 83.9 to 86.6%, demonstrating higher accuracy and sensitivity in object detection. Additionally, the mAP@50 increased dramatically from 87.1 to 97.3%, while mAP@50:95, a tougher and more comprehensive measure, improved from 66.8 to 71.5%. These results show that data augmentation improves the model's generalization capacity by exposing it to a broader number of training instances, resulting in more robust and reliable detections in a variety of operational environments.

The Precision-Recall (PR) curves shown in Fig. 5 for the YOLOv8s model trained with and without data augmentation demonstrate how augmentation affects detection performance. In subplot (a), the model trained with augmentation has a better total class score of 97.3%, compared to 87.1% in subplot (b). The PR curve with augmentation demonstrates more consistent and elevated precision across all recall levels, indicating that the model is more confident and accurate in detecting objects in various scenarios. This improvement emphasizes the importance of augmentation in improving the model's robustness, generalization, and ability to maintain high detection quality even under difficult or varied input conditions.

Model	Precision (%)	Recall (%)	mAP@50 (%)	mAP@50:95 (%)	Inference time (ms)
YOLOv5	76.4	67.5	66.4	43.2	75.9
YOLOv7	85.4	73.6	80.6	48.8	72.9
YOLOv8s	86.9	79.8	89.6	59.0	71.9
YOLOv8s (modified)	96.9	86.6	97.3	71.5	66.8

Table 3. The comparison of Modified YOLOv8s with other models' performance based on key evaluation metrics.




Fig. 6. The confusion matrix of (a) YOLOv8s-modified model and (b) YOLOv8s model.

Performance comparison of YOLOv8s-modified and other models

This section compares the performance of improved YOLOv8s, YOLOv8 (baseline), YOLOv7, and YOLOv5, with a focus on their ability to detect anomalies associated with fiber optic cable infrastructure specifically, human climbing activities on pole-mounted systems. Conventional object detection datasets frequently lack annotated examples of such specialized events. As a result, a bespoke dataset was created for this investigation, which included high-resolution photos and annotated bounding boxes for unauthorized human presence, climbing, and other suspicious activity near optical fiber poles. To ensure experimental consistency and fairness, each model was trained independently but under the same settings. The improved YOLOv8s has a lightweight backbone with smaller channel sizes, which increases speed and lowers computational cost. Improved FPN and anchor-free heads increase localization accuracy, particularly in obstructed scenes. It's designed for real-time anomaly detection on edge devices that have limited processing power. The comparative study results are summarized in Table 3. YOLOv8s-modified performed better than all other models in every key statistic. It had the highest precision (96.9%), demonstrating the capacity to reduce false positives, and the highest recall (86.6%), suggesting strong detection capability without missing crucial events. Furthermore, it achieved a mAP@50 of 97.3% and a mAP@50:95 of 71.5%, proving its stability across various detection thresholds and confirming the high localization accuracy.

The improved YOLOv8s also achieved the fastest inference time of 66.8 ms, making it ideal for real-time monitoring applications that require immediate danger identification. In comparison, YOLOv5 demonstrated significant flaws, with a lower mAP@50 (66.4%) and recall (67.5%), occasionally failing to detect climbing behaviors, which could result in missing anomalies during live monitoring. Additionally, YOLOv8s-modified performed more smoothly and consistently on confidence-based metrics, such as the precision-confidence, recall-confidence, and F1-confidence curves as shown in Fig. 3. These curves showed its consistent detection thresholds and minimal variability, resulting in fewer misclassifications in uncertain environments. YOLOv8s-modified has greater performance due to its advanced transformer-inspired architecture, adaptive feature fusion algorithms, and increased object tracking consistency between frames. Therefore, YOLOv8s-modified emerged as the best effective model for detecting anomalies near pole-mounted fiber optic infrastructures. It combines high detection accuracy, rapid inference speed, and robustness against severe environmental conditions making it an optimum choice for implementation in automated telecom monitoring and maintenance systems.

Comparison of YOLOv8s-modified and YOLOv8s models

The confusion matrix comparison of modified YOLOv8s and original YOLOv8s shows significant improvements in classification accuracy following model augmentation as presented in Fig. 6. The improved YOLOv8s has

Training epochs	mAP@50 (%)	Observations
20	78.9	Early-stage training provides a reasonable baseline but has higher false positives and missed detections
50	87.5	Significant improvement in accuracy; better feature extraction and reduced false detections
100	97.3	Performance stabilizes, confirming optimal convergence. Further training yields diminishing returns

Table 4. Effect of training epochs on model accuracy.

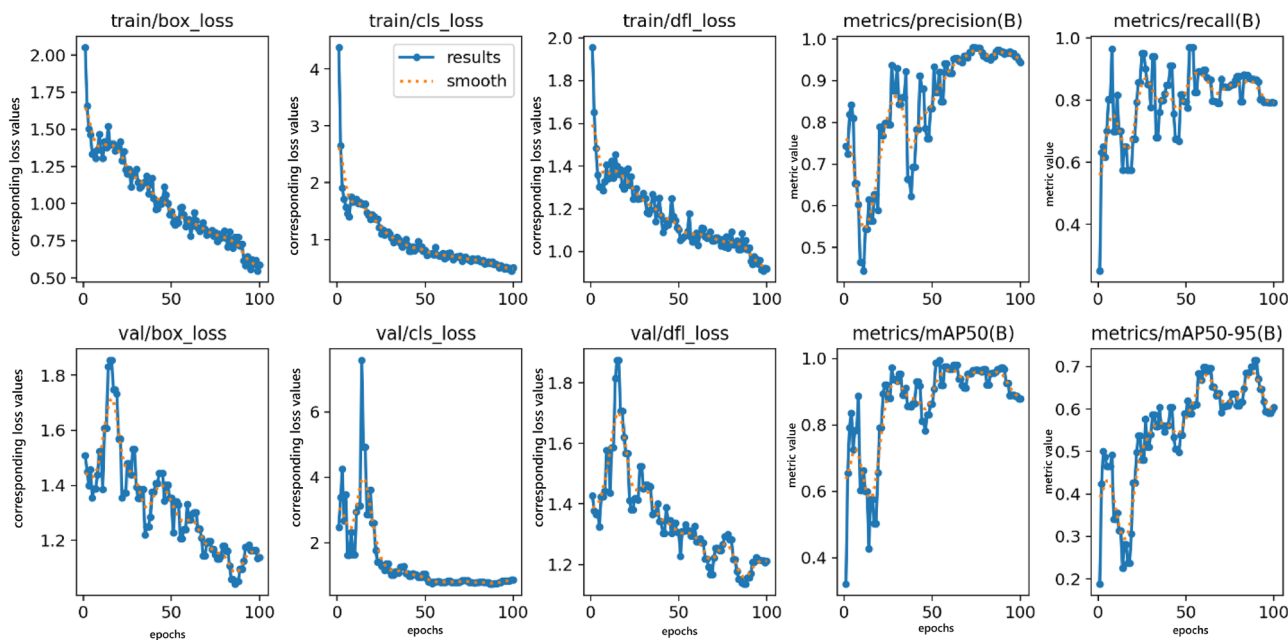


Fig. 7. Training and validation metrics across 100 epochs for the improved yolov8s.

less false positives and false negatives, with more predictions correctly falling along the diagonal of the matrix, indicating higher precision and recall across all classes. In contrast, the original YOLOv8s matrix shows more class confusion, especially in overlapping or visually similar activity categories. This increase indicates how adjustments, such as augmentation, attention modules, or architecture tuning, can improve detection reliability and resilience in complicated surveillance tasks.

Effect of training epochs on model accuracy

The effect of training epochs on model accuracy was investigated by comparing performance at various stages, as shown in Table 4. At 20 epochs, the model's mAP@0.5 was 78.9%, indicating early learning but with some false positives and missed detections. By 50 epochs, accuracy had increased significantly to 87.5%, indicating better feature extraction and fewer errors. After 100 epochs, the model stabilized at mAP@0.5 of 97.3%, showing optimal convergence. While more epochs improved accuracy, the gains tapered off beyond this point, implying that additional training could lead to overfitting without significant performance gains. These findings emphasize the significance of balancing training duration and computational efficiency, with 50–100 epochs being ideal for deep learning-based object detection models such as YOLOv8s.

Model training and validation loss summary

The improved YOLOv8s model's training process was evaluated with key loss metrics and performance indicators. The train/box_loss metric assesses the model's ability to predict accurate bounding box coordinates; a lower value indicates improved localization. Similarly, train/cls_loss captures classification mistakes, and train/df_l_loss measures the quality of bounding box regression refinement as illustrated in Fig. 7. On the performance front, metrics/precision(B) and metrics/recall(B) are crucial for assessing the model's capacity to detect true abnormalities correctly. Precision determines how many predicted anomalies were correct, whereas recall reveals how many real abnormalities were correctly detected. metrics/mAP50(B) and metrics/mAP50-95(B) summarize detection accuracy at various IoU levels, with mAP@50 indicating leniency and mAP@50–95 indicating stringent examination. The model improved consistently across various metrics during training, with convergence occurring after multiple epochs. The high mAP values validate the model's robust generalization and anomaly detection capabilities.

Dataset	Precision (%)	Recall (%)	mAP@50 (%)	mAP@50:95 (%)	Inference time (ms)
Roboflow Universe ⁶²	93.2	84.6	91.7	68.3	71.8
Kaggle ⁶³	93.8	85.1	92.4	74.1	71.5
Open images Dataset ⁶⁴	85.9	86.7	85.3	70.2	74.6

Table 5. The modified YOLOv8s performance on anomaly detection datasets.

Experiment ID	Model variant	Precision (%)	Recall (%)	mAP@50 (%)	mAP@50:95 (%)	Inference time (ms)	Notes
A	YOLOv8s (original)	86.9	79.8	89.6	59.0	71.9	Baseline
B	Early-channel reduction (48, 96)	87.5	80.4	90.2	59.8	70.3	Lighter first layers
C	Hierarchical channels (P3 = 192, P4 = 384, P5 = 768)	88.6	81.2	91.5	61.0	70.5	Better multi-scale features
D	Reduced C2f repetitions (3)	88.0	80.6	90.9	60.2	67.8	Lower parameter count
E	SPPF reduced (768 channels)	88.7	81.0	91.4	60.6	70.2	Contextual pooling at lower cost
F	LSKAttention added (768 channels)	92.5	83.8	94.6	65.2	72.5	Spatial focus on targets
H	YOLOv8s (modified, all changes)	96.9	86.6	97.3	71.5	66.8	Final model

Table 6. Ablation experiment results of improved YOLOv8s.

Evaluation on external anomaly datasets

The proposed model was tested for generalizability and robustness using three public datasets: Roboflow Universe, Kaggle, and the Open Images Dataset, which included a wide range of image types and instances. The Open Images Dataset, which lacked the YOLOv8 format, was converted before training. The results in Table 5 show persistent high performance: Kaggle achieved the highest precision (93.8%) and mAP scores, while Open Images had the highest recall (86.7%), suggesting flexibility to complex imagery. Roboflow Universe achieved balanced results, including high precision (93.2%) and mAP@50 (91.7%). Inference times remained within 71–75 ms across datasets, allowing for real-time application. Overall, the model displayed resilience and accuracy across multiple datasets, proving its potential for robust anomaly identification in a variety of surveillance scenarios.

Ablation study

In order to evaluate the contribution of each architectural component to the performance of the modified YOLOv8s, an ablation study was conducted. The results are presented in Table 6. The original YOLOv8s (Experiment A) achieves a high accuracy of 89.6% mAP@50 and an inference time of 71.9 ms, which indicates a good trade-off between accuracy and computational efficiency for real-time object detection tasks. Introducing early-channel reduction (Experiment B) reduces the number of parameters in the first convolutional layers, resulting in a modest 0.6% improvement in mAP@50 while also lowering inference time by 1.6 ms. This demonstrates that limiting early-stage overhead enhances efficiency without compromising accuracy. Applying hierarchical channel scaling in deeper backbone stages (Experiment C) yields a further 1.3% increase in mAP@50 and improves fine-grained multi-scale feature representation, especially beneficial for detecting small and irregularly shaped objects. Importantly, inference time remains stable, showing this modification enhances accuracy at minimal additional cost. Reducing the number of C2f repetitions (Experiment D) achieves a noticeable latency reduction (−4.1 ms) and decreases parameter count, while maintaining accuracy close to baseline. This indicates that redundant backbone depth can be removed without significant degradation in detection quality.

Replacing the original SPPF with a reduced-channel SPPF module (Experiment E) slightly boosts accuracy (+0.8% mAP@50) while lowering computational cost, confirming the efficiency of streamlined contextual aggregation. The addition of the LSKAttention block (Experiment F) provides the single largest accuracy improvement (+4.0% mAP@50 compared to baseline). By reducing background noise and emphasizing informative spatial regions, the network's large-kernel attention improves generalization on intricate custom datasets. The trade-off is a somewhat longer inference time (~+0.6 ms). Finally, merging all improvements into the proposed YOLOv8s (modified) (Experiment H) results in the greatest performance across all metrics, attaining 97.3% mAP@50 and 71.5% mAP@50:95 while reducing average inference time to 66.8 ms. Compared to the original YOLOv8, this is a significant +7.7% improvement in mAP@50 and −5.1 ms latency reduction, showing that the improvements strike a superior balance between accuracy and efficiency.

Comparative study

The comparative evaluation highlights the performance differences between traditional and recent deep learning models for anomaly detection. Faster R-CNN¹⁹ achieves a precision of 82.3%, recall of 75.4%, and mAP@50 of 84.1%, but its mAP@50:95 is relatively low at 54.2%. With an inference time of 93.5 ms as shown in Table 7, it

Model	Precision (%)	Recall (%)	mAP@50 (%)	mAP@50:95 (%)	Inference time (ms)	Strengths	Limitations
Faster R-CNN ¹⁹	82.3	75.4	84.1	54.2	93.5	Strong detection of diverse anomalies	Computationally heavy, poor real-time applicability
YOLOv5s ⁴²	82.2	63.9	72.6	–	–	Lightweight and widely used	Lower accuracy on small/complex anomalies
YOLOv7 ⁴³	86.4	71.3	75.9	–	–	High precision, improved real-time detection	Recall and mAP limited on challenging scenes
YOLOv8s (original)	86.9	79.8	89.6	59.0	71.9	Balanced accuracy and speed	Accuracy limited on small, complex anomalies
YOLOv8s (modified) Ours	96.9	86.6	97.3	71.5	66.8	Highest accuracy, robust generalization, faster inference, optimized for edge deployment	–

Table 7. Comparative performance of anomaly detection models.

is the slowest model in the comparison. While it is strong at detecting diverse anomalies due to its region-based architecture, the computational overhead makes it unsuitable for real-time monitoring tasks.

YOLOv5s⁴² provides a lightweight and widely adopted solution with a precision of 82.2% and mAP@50 of 72.6%. However, the recall rate is substantially lower (63.9%), indicating that it frequently fails to recognize minor or more complicated anomalies. This shortcoming restricts its utility in infrastructure surveillance applications where sensitivity to uncommon anomalies is essential. YOLOv7⁴³ outperforms YOLOv5 with a precision of 86.4% and recall of 71.3%. Its mAP@50 is 75.9% and mAP@50:95 is 54.4%, indicating improved detection consistency than YOLOv5, but it still struggles with more difficult anomalous conditions. Despite its precision advantage, the lack of reported inference time limits assessment of its real-time practicality.

The original YOLOv8s strikes a stronger balance between accuracy and efficiency. It records 86.9% precision, 79.8% recall, 89.6% mAP@50, and 59.0% mAP@50:95, with a faster inference time of 71.9 ms compared to Faster R-CNN. This makes it more appropriate for real-time applications, although its speed remains limited when dealing with small-scale or visually complicated abnormalities.

In contrast, the updated YOLOv8s proposed in this work had a clear advantage across all criteria. It reduces the inference time to 66.8 ms while achieving 96.9% precision, 86.6% recall, 97.3% mAP@50, and 71.5% mAP@50:95. This performance boost is due to architectural improvements such as lightweight channel reduction, hierarchical scaling, fewer C2f repetitions, streamlined SPPF, and the introduction of LSKAttention. These enhancements not only increase feature extraction and contextual awareness, but they also reduce processing costs, allowing for more robust detection in real time. Overall, the updated YOLOv8s model strikes the best balance of accuracy, generalization, and efficiency, making it the most ideal model for identifying abnormalities in telecommunications infrastructure.

Conclusion

This study successfully demonstrated the use of YOLOv8s for real-time anomaly detection in fiber optic cables mounted on poles, with the emphasis on detecting climbing activities and other objects such as, person, animal and pole. Considering the lack of climbing-related annotations in available datasets, a custom dataset was created to incorporate a variety of context-specific events. This significantly improved the model's learning capabilities, allowing it to detect human presence, climbing motions, and key ambient parameters with high precision. During training, various data augmentation procedures were used, such as random flipping, Gaussian noise, brightness modifications, and rotation. These strategies helped improve the model's generalization and robustness, particularly under changing environmental conditions. After 100 training epochs, the model attained a peak mAP@0.5 of 97.3%, indicating that it struck the right balance between learning depth and overfitting mitigation. A training progression analysis revealed mAP gains ranging from 78.9% at 20 epochs to 87.5% at 50 epochs, showing that extended training resulted in more refined feature extraction and classification. When compared to other state-of-the-art models, the improved YOLOv8s outperformed all main evaluation metrics. Precision raised significantly to 96.9%, recall to 86.6%, mAP@50 increased to 97.3%, and mAP@50:95 improved to 71.5%. These results show significant improvements over the baseline YOLOv5 (precision: 76.4%, mAP@50: 66.4%) and even the original YOLOv8s (precision: 86.9%, mAP@50: 89.6%). The reduction in inference time to 66.8 ms emphasizes its effectiveness. These increases are mostly due to model backbone optimizations, which increased the model's capacity to detect small, overlapping, and complicated anomalous occurrences with more confidence and speed.

However, several constraints should be considered. While the custom dataset is extensive, it is very small in size and geographic breadth, which may limit the model's applicability to diverse environments or infrastructure types. The study concentrated on pole-mounted fiber optic cables in rural and peri-urban settings, therefore performance in dense urban contexts or with alternative camera systems requires more confirmation. Furthermore, the computational resources utilized for training and inference were moderate. Scaling this technique to massive network deployments with real-time restrictions may require further optimization. This study has important practical implications for the monitoring of telecommunications infrastructure. The demonstrated accuracy and resilience of YOLOv8s justify their use in intelligent, automated surveillance systems capable of proactively identifying security concerns, unlawful access, and environmental dangers affecting fiber optic networks. This can lower maintenance costs, avoid service disruptions, and improve overall network

stability. Furthermore, YOLOv8s' lightweight design makes them easy to integrate with edge computing devices, allowing for near real-time decision making and decreasing data transmission loads. Therefore, this study demonstrated YOLOv8s-modified as an extremely effective approach in detecting fiber optic cable anomalies, enabling increased accuracy and dependability while also laying the groundwork for real-world deployment in intelligent telecom monitoring systems. Future research could look into optimization to improve inference speed and scalability for large-scale network applications.

Future work

While YOLOv8s demonstrated high accuracy, several areas for future improvement remain:

- Optimizing model efficiency by exploring lightweight architectures to enhance inference speed on edge devices.
- Expanding dataset diversity by incorporating images from various geographic locations, lighting conditions, and environmental settings to improve generalization.
- Integrating temporal analysis by leveraging video-based anomaly detection techniques to track dynamic changes and detect patterns over time.

Data availability

The datasets utilized and analyzed in this study are available from the corresponding author upon reasonable request.

Received: 19 April 2025; Accepted: 30 September 2025

Published online: 05 November 2025

References

1. Abdelli, K. et al. Machine-learning-based anomaly detection in optical fiber monitoring. *J. Opt. Commun. Netw.* **14**(5), 365–375 (2022).
2. J. Farooq, M. Muaz, K. Khan Jadoon, N. Aafaq, and M. K. A. Khan (2024) An improved YOLOv8 for foreign object debris detection with optimized architecture for small objects. *Multimedia Tools and Applications*, 83(21): 60921–60947.
3. Lv, Z., Li, Y., Feng, H. & Lv, H. Deep learning for security in digital twins of cooperative intelligent transportation systems. *IEEE Trans. Intell. Transp. Syst.* **23**(9), 16666–16675 (2021).
4. Liu, W. et al. Ssd: Single shot multibox detector. In *Computer vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14* 21–37 (Springer, Cham, 2016).
5. Alzubaidi, L. et al. Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *J. big Data* **8**, 1–74 (2021).
6. Xia, L., Yang, D., Zhang, J., Yang, H. & Chen, J. Enhanced semantic information transfer of multi-domain samples: An adversarial edge detection method using few high-resolution remote sensing images. *Sensors* **22**(15), 5678 (2022).
7. J. Redmon and A. Farhadi, YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7263–7271 (2017).
8. Wang, J. et al. Road defect detection based on improved YOLOv8s model. *Sci. Rep.* **14**(1), 16758 (2024).
9. K. Aishwarya, R. Suthan, C. C. Deboral, and S. KV, A Deep Learning Strategy for Abnormal Object Detection by YOLOv8 Model. In *2024 International Conference on Power, Energy, Control and Transmission Systems (ICPECTS)*, 2024: IEEE, pp. 1–5.
10. Duong, H.-T., Le, V.-T. & Hoang, V. T. Deep learning-based anomaly detection in video surveillance: A survey. *Sensors* **23**(11), 5024 (2023).
11. P. Wu, C. Pan, Y. Yan, G. Pang, P. Wang, and Y. Zhang, Deep learning for video anomaly detection: A review. *arXiv preprint arXiv:2409.05383*, 2024.
12. Lalam, N. et al. Achieving precise multiparameter measurements with distributed optical fiber sensor using wavelength diversity and deep neural networks. *Commun. Eng.* **3**(1), 121 (2024).
13. Chen, Z. et al. Machine learning-enabled IoT security: Open issues and challenges under advanced persistent threats. *ACM Comput. Surv.* **55**(5), 1–37 (2022).
14. Abdelli, K., Griefser, H., Trotschug, C. & Pachnicke, S. Optical fiber fault detection and localization in a noisy OTDR trace based on denoising convolutional autoencoder and bidirectional long short-term memory. *J. Lightwave Technol.* **40**(8), 2254–2264 (2021).
15. K. Abdelli, H. Griesser, and S. Pachnicke, Convolutional neural networks for reflective event detection and characterization in fiber optical links given noisy OTDR signals. In *Photonic Networks; 22th ITG Symposium*; VDE, pp. 1–5, (2021).
16. Abdelli, K., Griefser, H., Ehrle, P., Trotschug, C. & Pachnicke, S. Reflective fiber fault detection and characterization using long short-term memory. *J. Opt. Commun. Netw.* **13**(10), E32–E41 (2021).
17. Natalino, C., Udalcovs, A., Wosinska, L., Ozolins, O. & Furdek, M. Spectrum anomaly detection for optical network monitoring using deep unsupervised learning. *IEEE Commun. Lett.* **25**(5), 1583–1586 (2021).
18. Wang, D. & Zhang, M. Artificial intelligence in optical communications: from machine learning to deep learning. *Front. Commun. Netw.* **2**, 656786 (2021).
19. Rizzo, A. M. et al. Known and unknown event detection in OTDR traces by deep learning networks. *Neural Comput. Appl.* **34**(22), 19655–19673 (2022).
20. Huot, F. et al. Detection and characterization of microseismic events from fiber-optic DAS data using deep learning. *Seismolog. Soc. Am.* **93**(5), 2543–2553 (2022).
21. Zhang, L., Gao, W. & Yan, L. Deep learning-based fault diagnosis and localization method for fiber optic cables in communication networks. *Appl. Math. Nonlinear Sci.* **9**, 1–14 (2023).
22. Ghamisi, A. et al. Anomaly detection in automated fibre placement: Learning with data limitations. *Front. Manuf. Technol.* **4**, 1277152 (2024).
23. L. Zhang, L. Yan, W. Shen, F. Li, J. Wu, and W. Liang, “Neural network-based fiber optic cable fault prediction study for power distribution communication network. *Appl. Math. Nonlinear Sci.*, **9**(1), (2024).
24. Sheng, X., Li, S. & Chan, S. Real-time classroom student behavior detection based on improved YOLOv8s. *Sci. Rep.* **15**(1), 14470 (2025).
25. Girisha, S., Verma, U., Pai, M. M. & Pai, R. M. Contextual information based anomaly detection for multi-scene aerial videos. *Sci. Rep.* **15**, 25805 (2025).
26. Khatoon, A., Wang, W., Wang, M., Li, L. & Ullah, A. TinyML-enabled fuzzy logic for enhanced road anomaly detection in remote sensing. *Sci. Rep.* **15**(1), 20659 (2025).

27. Alotaibi, S. R. et al. Harnessing optimization with deep learning approach on intelligent transportation system for anomaly detection in pedestrian walkways. *Sci. Rep.* **15**(1), 17358 (2025).
28. Veeram, S. B. et al. Design of an integrated model with temporal graph attention and transformer-augmented RNNs for enhanced anomaly detection. *Sci. Rep.* **15**(1), 2692 (2025).
29. Liu, S. et al. A deep learning based detection algorithm for anomalous behavior and anomalous item on buses. *Sci. Rep.* **15**(1), 2163 (2025).
30. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788 (2016).
31. J. Redmon and A. Farhadi, Yolov3: An incremental improvement, *arXiv preprint arXiv:1804.02767* (2018).
32. A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, YOLOv4: Optimal speed and accuracy of object detection. *arXiv e-prints*, p. arXiv:2004.10934 (2020).
33. G. Jocher et al., ultralytics/yolov5: v3. 0. *Zenodo*, 2020.
34. Wang, S., Kim, M., Hae, H., Cao, M. & Kim, J. The development of a rebar-counting model for reinforced concrete columns: Using an unmanned aerial vehicle and deep-learning approach. *J. Constr. Eng. Manag.* **149**(11), 04023111 (2023).
35. Wang, S. Evaluation of impact of image augmentation techniques on two tasks: Window detection and window states detection. *Results in Engineering* **24**, 103571 (2024).
36. C. Li et al., YOLOv6: A single-stage object detection framework for industrial applications. *arXiv preprint arXiv:2209.02976*, 2022.
37. C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7464–7475 (2023).
38. G. Jocher, Chaurasia, A., and Qiu, J., YOLO by Ultralytics (Version 8.0.0). In *Computer software*. ed. (2023).
39. Wang, A. et al. Yolov10: Real-time end-to-end object detection. *Adv. Neural. Inf. Process. Syst.* **37**, 107984–108011 (2024).
40. Eum, I., Kim, J., Wang, S. & Kim, J. Heavy equipment detection on construction sites using you only look once (yolo-version 10) with transformer architectures. *Appl. Sci.* **15**(5), 2320 (2025).
41. Wang, C.-Y., Yeh, I.-H. & Mark Liao, H.-Y. Yolov9: Learning what you want to learn using programmable gradient information. In *European conference on computer vision* 1–21 (Springer, 2024).
42. Yang, W. et al. A forest wildlife detection algorithm based on improved YOLOv5s. *Animals* **13**(19), 3134 (2023).
43. C. Sakiba, S. M. Tarannum, F. Nur, F. F. Arpan, and A. A. Anzum, *Real-time crime detection using convolutional LSTM and YOLOv7*. Brac University (2023).
44. Wang, X., Gao, H., Jia, Z. & Li, Z. BL-YOLOv8: An improved road defect detection model based on YOLOv8. *Sensors* **23**(20), 8361 (2023).
45. Yang, D. et al. A streamlined approach for intelligent ship object detection using EL-YOLO algorithm. *Sci. Rep.* **14**(1), 15254 (2024).
46. Liu, J. et al. Optimization of a multi-environmental detection model for tomato growth point buds based on multi-strategy improved YOLOv8. *Sci. Rep.* **15**(1), 25726 (2025).
47. Zhang, H., Liang, M. & Wang, Y. YOLO-BS: a traffic sign detection algorithm based on YOLOv8. *Sci. Rep.* **15**(1), 7558 (2025).
48. Elfving, S., Uchibe, E. & Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Netw.* **107**, 3–11 (2018).
49. He, K., Zhang, X., Ren, S. & Sun, J. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(9), 1904–1916 (2015).
50. Z. Ge, S. Liu, F. Wang, Z. Li, and J. Sun, Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430* (2021).
51. Wang, A. et al. NVW-YOLOv8s: An improved YOLOv8s network for real-time detection and segmentation of tomato fruits at different ripeness stages. *Comput. Electron. Agric.* **219**, 108833 (2024).
52. I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning* (no. 2). MIT press Cambridge, 2016.
53. Knoblauch, A. Adapting loss functions to learning progress improves accuracy of classification in neural networks. In *International Symposium on Methodologies for Intelligent Systems* 272–282 (Springer, 2022).
54. Zheng, Z. et al. Distance-IoU loss: Faster and better learning for bounding box regression. *Proc. AAAI Conf. Artif. Intell.* **34**(07), 12993–13000 (2020).
55. Su, Q. et al. Drone object detection incorporating multi-head mixed self-attention and dynamic regression mapping loss function. *J. Real-Time Image Proc.* **22**(2), 56 (2025).
56. Xiao, B., Nguyen, M. & Yan, W. Q. Fruit ripeness identification using YOLOv8 model. *Multimed. Tools Appl.* **83**(9), 28039–28056 (2024).
57. J. Liu, C. Zhang, and J. Li, Use anchor-free based object detectors. In *Proceedings of the TEPEN International Workshop on Fault Diagnostic and Prognostic: TEPEN2024-IWFDP-Volume 1*, vol 1. p. 348 (Springer, Cham, 2024).
58. D. Tzatalin, *LabelImg Free Software: MIT License*. ed. MIT: Cambridge, MA, USA, 2015.
59. Lin, T.-Y. et al. Microsoft coco: Common objects in context. In *European conference on computer vision* 740–755 (Springer, 2014).
60. Everingham, M., Van Gool, L., Williams, C. K., Winn, J. & Zisserman, A. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vision* **88**(2), 303–338 (2010).
61. Han, J., Kim, J., Kim, S. & Wang, S. Effectiveness of image augmentation techniques on detection of building characteristics from street view images using deep learning. *J. Constr. Eng. Manag.* **150**(10), 04024129 (2024).
62. R. Universe, *Open Source Computer Vision Community*. ed. (2025).
63. K. datasets, *Public Datasets*. ed. (2025).
64. Kuznetsova, A. et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *Int. J. Comput. Vision* **128**(7), 1956–1981 (2020).

Acknowledgements

Not applicable.

Author contributions

Each author listed has made substantial contributions to the development and composition of this manuscript. EE conceived the initial idea, while ANS, UKJ, and BOS provided supervision throughout the research process. EE, BOS, ANS and UKJ were involved in revising and refining the final manuscript. All authors have reviewed the manuscript and provided their approval for its publication.

Funding

The authors have no relevant financial or non-financial interests to disclose.

Declarations

Competing interests

The authors declared that they have no competing interests.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Additional information

Correspondence and requests for materials should be addressed to E.E.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025

Terms and Conditions

Springer Nature journal content, brought to you courtesy of Springer Nature Customer Service Center GmbH (“Springer Nature”).

Springer Nature supports a reasonable amount of sharing of research papers by authors, subscribers and authorised users (“Users”), for small-scale personal, non-commercial use provided that all copyright, trade and service marks and other proprietary notices are maintained. By accessing, sharing, receiving or otherwise using the Springer Nature journal content you agree to these terms of use (“Terms”). For these purposes, Springer Nature considers academic use (by researchers and students) to be non-commercial.

These Terms are supplementary and will apply in addition to any applicable website terms and conditions, a relevant site licence or a personal subscription. These Terms will prevail over any conflict or ambiguity with regards to the relevant terms, a site licence or a personal subscription (to the extent of the conflict or ambiguity only). For Creative Commons-licensed articles, the terms of the Creative Commons license used will apply.

We collect and use personal data to provide access to the Springer Nature journal content. We may also use these personal data internally within ResearchGate and Springer Nature and as agreed share it, in an anonymised way, for purposes of tracking, analysis and reporting. We will not otherwise disclose your personal data outside the ResearchGate or the Springer Nature group of companies unless we have your permission as detailed in the Privacy Policy.

While Users may use the Springer Nature journal content for small scale, personal non-commercial use, it is important to note that Users may not:

1. use such content for the purpose of providing other users with access on a regular or large scale basis or as a means to circumvent access control;
2. use such content where to do so would be considered a criminal or statutory offence in any jurisdiction, or gives rise to civil liability, or is otherwise unlawful;
3. falsely or misleadingly imply or suggest endorsement, approval, sponsorship, or association unless explicitly agreed to by Springer Nature in writing;
4. use bots or other automated methods to access the content or redirect messages
5. override any security feature or exclusionary protocol; or
6. share the content in order to create substitute for Springer Nature products or services or a systematic database of Springer Nature journal content.

In line with the restriction against commercial use, Springer Nature does not permit the creation of a product or service that creates revenue, royalties, rent or income from our content or its inclusion as part of a paid for service or for other commercial gain. Springer Nature journal content cannot be used for inter-library loans and librarians may not upload Springer Nature journal content on a large scale into their, or any other, institutional repository.

These terms of use are reviewed regularly and may be amended at any time. Springer Nature is not obligated to publish any information or content on this website and may remove it or features or functionality at our sole discretion, at any time with or without notice. Springer Nature may revoke this licence to you at any time and remove access to any copies of the Springer Nature journal content which have been saved.

To the fullest extent permitted by law, Springer Nature makes no warranties, representations or guarantees to Users, either express or implied with respect to the Springer nature journal content and all parties disclaim and waive any implied warranties or warranties imposed by law, including merchantability or fitness for any particular purpose.

Please note that these rights do not automatically extend to content, data or other material published by Springer Nature that may be licensed from third parties.

If you would like to use or distribute our Springer Nature journal content to a wider audience or on a regular basis or in any other manner not expressly permitted by these Terms, please contact Springer Nature at

onlineservice@springernature.com