International Journal of Combinatorial Optimization Problems and Informatics, 16(3), May-Aug 2025, 310-322. ISSN: 2007-1558.
https://doi.org/10.61467/2007.1558.2025.v16i3.714

www.editada.org

Implementing Deep Learning for Real-Time Fuel Tank Detection in UAV Surveillance Systems

Ali Bin Khalifa[1], Hamad Al-Nuaimi[1], Latifa Bint Ahmed[2], Salem Jassim[1], Affan Alkim[2],
[1] Computer Engineering Department, Qatar University, Qatar
[2] Department of Information Technology, Lusail University, Qatar
E-mails: alkim@lu.edu.qa

Abstract. This study investigates the performance of the YOLOv8 object detection algorithm for fuel tank detection using unmanned aerial vehicles (UAVs) under various conditions and perspectives. The model achieves a precision of 0.888, recall of 0.896, and mAP of 0.891, confirming its strong capabilities in detecting fuel tanks and supporting the sustainability of industrial and energy operations. With a processing time of 41 ms, YOLOv8 proves to be highly effective for real-time applications. This research highlights the importance of optimizing UAVs and deep learning models for reliable data collection in challenging environments and demonstrates their potential for use in fuel tank monitoring and infrastructure reliability tracking across industries. While accurate detection successes are noted, the study emphasizes the need for further optimization of the algorithm to address false positives and undetected objects in real-world applications. Future work will explore the adaptation of high-performance algorithms for broader object detection scenarios.

Keywords: YOLOv8 algorithm, Fuel tank detection, Real-time applications, Industrial automation, Model optimization

Article Info
Received January 28, 2025
Accepted March 21, 2025

1 Introduction

The constant increase in global energy demand and the expansion of international trade have heightened the importance of efficient storage and shipping processes for oil and its derivatives. This has made it essential for tanks to support the sustainability of industrial activities and energy infrastructure. Tanks are specifically designed to ensure the safe storage and transportation of crude oil and petroleum-derived products [1, 2]. Strategically located tanks within refinery areas streamline transportation and processing. These fuel tanks are manufactured using materials tailored to specific fuel types, such as oil, diesel, gasoline, or natural gas, and are designed to meet the requirements of their intended tasks. While many tanks, often constructed from steel, feature floating roofs, others are designed as pressurized systems. Safety considerations, particularly against fire and explosion risks, are a top priority in tank design for oil and gas production facilities [3]. Key design factors include resistance to pressure and impact under varying weather conditions, sealing to prevent environmental pollution and fuel waste, and adaptability to evolving global energy demands. To meet these criteria, the integration of advanced technologies, such as unmanned aerial vehicles (UAVs), is becoming increasingly important for monitoring and managing these critical structures, which play a vital role across various sectors.

Today, rapid advancements in remote sensing technology have brought about the diversification of methods used to generate high-resolution image outputs [4-8]. High-resolution remote detection techniques deliver superior image quality and detailed information, creating significant opportunities to advance object detection within the scope of remote sensing [9-12]. In the field of remote object detection, UAVs and satellites have emerged as leading methods, thanks to the imaging capabilities they have developed over time through progressive policies [13-17]. UAVs, with their advanced mobility, can be deployed swiftly and in succession, performing target detection while also capturing closer and higher-quality image outputs

[18-21]. Additionally, UAVs, when combined with the bird's-eye view provided by satellites, allow for target evaluation from diverse angles and perspectives. However, UAV performance is adversely impacted by harsh weather conditions, which can complicate and hinder target detection. In contrast, satellites remain unaffected by variable weather, enabling consistent and continuous data collection. This makes satellites particularly valuable for monitoring long-term changes with ease. Moreover, satellites can cover wide geographical areas, whereas UAVs typically operate at a local scale due to their limited range of activity. Despite these distinctions, UAVs possess unique advantages. Their advanced sensors and cameras allow them to navigate dangerous or hard-to-reach areas and efficiently carry out sensitive tasks within a short period. Unlike satellites, which rely on fixed equipment and require extensive planning, UAVs can be easily adapted with different sensors tailored to specific tasks, making them highly versatile and responsive to changing operational needs. This adaptability gives UAVs an edge over satellites in many scenarios, particularly under dynamic and variable conditions. Consequently, UAVs have surpassed satellites in certain applications due to their flexibility, efficiency, and ability to deliver customized solutions. With these advantages, UAVs are increasingly recognized as one of the most promising tools for target detection in remote sensing technology. By leveraging their strengths, UAVs stand out as a leading method, meeting the diverse demands of this evolving field.

Object detection involves tracking output information such as the location, orientation, and direction of predefined or undefined objects in an image [22-26]. The process of object detection is completed step by step, progressing through stages such as data input, preprocessing, feature extraction, feature selection, and finally classification, which we refer to as the data classification phase. Significant advancements in object detection and classification have been achieved alongside various developments in the field of deep learning [27–30]. Object detection algorithms are divided into two distinct branches: single-stage and two-stage detectors, depending on whether a single neural network or multiple steps are utilized [31, 32]. Single-stage algorithms, such as Single Shot Detector (SSD) and You Only Look Once (YOLO), perform both classification and location detection in a single step. This allows them to meet expectations for fast detection with reduced processing time [33, 34]. However, in scenarios requiring real-time, high-accuracy detection, single-stage algorithms often provide limited precision when detecting small objects. Single-stage detectors process an image through a convolutional neural network (CNN) only once and make direct predictions without requiring region proposals [35-38]. On the other hand, Region-Based Convolutional Neural Networks (R-CNN) and Faster R-CNN fall under the category of two-stage methods, dividing object detection into two steps [39, 40]. In the first step, region proposal networks (RPN) identify candidate object regions [41]. In the second step, these detected regions are classified to finalize the detection process. From this perspective, the choice between single-stage and two-stage object detectors depends on the specific requirements of the task they are designed to serve. Single-stage methods are often preferred in real-time detection scenarios where speed is a priority. In contrast, two-stage methods are advantageous in situations where accuracy takes precedence over speed, such as in critical target detection tasks. By evaluating the advantages offered by each approach, the appropriate method is selected based on the desired characteristics and application context.

UAVs have become a crucial tool for automated inspection and monitoring across various industrial domains, including infrastructure surveillance, environmental monitoring, and security applications. Their ability to capture high-resolution images in real time, navigate hard-to-reach areas, and operate autonomously makes them particularly well-suited for fuel tank monitoring. In parallel, advancements in deep learning have significantly improved object detection capabilities, with the YOLO algorithm emerging as one of the most efficient real-time detection models. YOLO follows a single-stage detection approach, enabling rapid and accurate object localization and classification in a single pass through a convolutional neural network. The latest iteration, YOLOv8, introduces enhanced architectural improvements, optimizing detection accuracy while maintaining high-speed performance. Given these strengths, integrating UAVs with YOLOv8 enables efficient, real-time fuel tank detection, offering a scalable and adaptable solution for industrial monitoring applications.

This study focuses on the detection of fuel tanks, which hold critical importance for global energy management, using image outputs captured by UAVs with the YOLOv8 algorithm. Additionally, it evaluates the effectiveness of this algorithm, which represents the 8th iteration of the YOLO family. YOLOv8 builds upon the high efficiency and rapid processing capabilities of its predecessors, further enhancing these features to offer significant advantages in time-sensitive applications. Leveraging state-of-the-art deep learning techniques, the YOLOv8 object detector aims to deliver superior performance in fuel tank detection, contributing to the effective monitoring and control of industrial and energy infrastructure. Given the strategic significance of fuel tanks in energy supply chains, ensuring their continuous monitoring and security is of paramount importance. Traditional inspection methods, which often rely on manual assessments or satellite imagery, can be time-consuming, costly, and subject to environmental limitations. UAVs equipped with advanced object detection algorithms, such as YOLOv8, provide a highly efficient alternative by offering rapid, high-resolution monitoring with real-time detection capabilities. This study specifically aims to demonstrate how UAV-based YOLOv8 detection can enhance the surveillance of fuel tanks, improving safety, operational efficiency, and predictive maintenance efforts. By integrating UAVs with deep learning-based object detection, this approach supports proactive monitoring strategies that can detect potential structural issues, fuel leaks, or security threats before they escalate into major hazards. The findings of this research will contribute to the advancement of automated monitoring solutions for critical energy infrastructure, ensuring reliability and sustainability in fuel storage management.

To ensure accurate and reliable detection, the YOLOv8 algorithm analyzes fuel tank images captured under diverse environmental and operational conditions. These images are obtained using UAV-mounted high-resolution cameras, which enable close-range monitoring and real-time data acquisition. The analysis considers factors such as varying illumination levels, weather conditions (e.g., cloudy, sunny, or low-light environments), and the presence of obstructions such as surrounding structures or vegetation. Additionally, the algorithm is evaluated on its ability to detect fuel tanks from different angles and altitudes, ensuring robust performance across multiple viewpoints. Given the industrial setting, challenges such as reflective surfaces, shadows, and overlapping objects are also taken into account during the detection process. By incorporating these factors, this study provides a comprehensive assessment of YOLOv8's effectiveness in real-world fuel tank monitoring scenarios.

2 Methods

The YOLO family, one of the single-stage object detectors, performs detection successfully without requiring multiple steps or region proposals. Consequently, its single-step detection approach significantly improves the speed of target detection. YOLO iterations utilize a single CNN to directly detect class labels and bounding boxes for objects in an image [42]. The YOLO approach was initially developed on the DarkNet framework, an open-source deep learning platform designed to enhance real-time object detection performance. Rather than relying on traditional object detection methods, YOLO employs a distinct approach. In the first step, the input image is divided into a predefined grid of cells, and each cell is scanned with bounding boxes. Predictions for each cell yield core outputs, including bounding boxes and classification information. In subsequent steps, bounding boxes and overlapping regions are processed with the Non-Maximum Suppression (NMS) algorithm. This algorithm selects the bounding box with the highest probability, suppressing less likely boxes and preventing the formation of an excessive number of bounding boxes. To achieve this, the NMS algorithm assigns a confidence score to each bounding box, indicating the likelihood that the detected object belongs to a specific class. This score, expressed as a value between 0 and 1, represents both the probability of the object belonging to a class and the algorithm's confidence in its prediction. Unlike earlier versions of YOLO, YOLOv8, which was introduced by Ultralytics in 2023, focuses on delivering superior performance for real-time detection tasks. Figure 1 illustrates the workflow diagram of UAV-based object detection using the YOLOv8 algorithm.

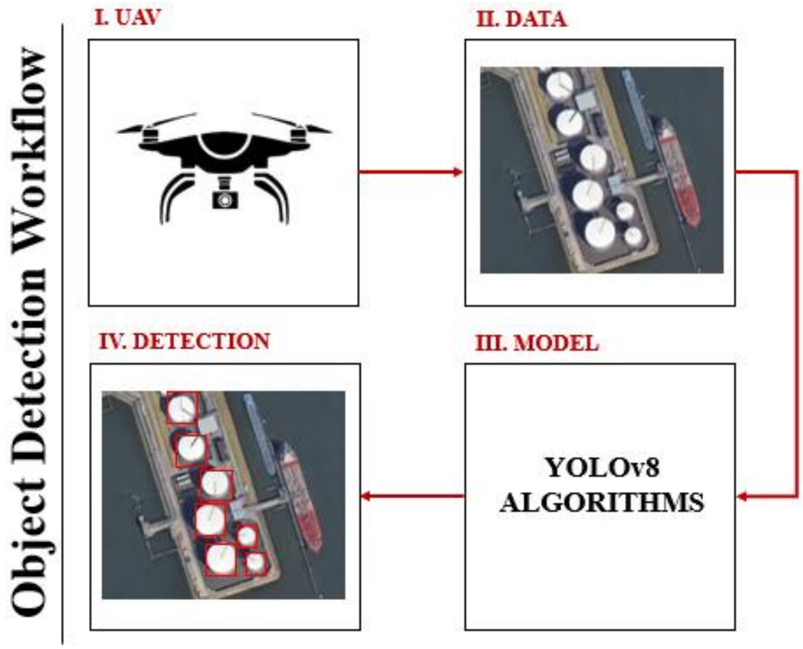

Figure 1. Flowchart illustrating the process for autonomous fuel tank detection.

2.1 YOLOv8 Algorithm

YOLOv8 represents a significant advancement in the YOLO family by building on the architecture of its predecessors, YOLOv5 and YOLOv7. Like the earlier versions, YOLOv8 maintains the single-stage detection mechanism [43]. This algorithm, which has gained attention for its anchor-free approach, can be divided into three main components: backbone, neck, and head. YOLOv8 introduces an enhanced version of the CSPDarknet series, known for its superior performance, and a new backbone architecture called CSPDarknet-AA, which takes model scalability to the next level. By reducing the number of parameters, CSPDarknet-AA achieves a robust balance between speed and performance. Using this technique, the object detector can effectively identify objects of varying sizes within an image. The C2f module, a merge-based replacement for the Cross Stage Partial (CSP) module, synthesizes outputs from different layers to improve detection accuracy. After the final convolution layer in the C2f block, the SPPF (Spatial Pyramid Pooling - Fast) block is applied. This algorithm replaces 6×6 convolutional layers with 3×3 convolutional layers, reducing parameter counts while achieving higher computational speed. Additionally, the SPPCSPC structure, which optimizes the spatial pyramid pooling (SPP) structure with a convolutional spatial pyramid (CSP), enhances the model's comprehension capabilities. SPPF, designed to adapt output dimensions, connects three max-pooling layers to minimize computational power requirements. In the neck section, a PAN-FPN (Path Aggregation Network with Feature Pyramid Network) structure, inspired by PANet, is utilized to preserve localization information of objects. This component is responsible for upsampling feature maps and integrating features extracted from previous layers. The use of PANet creates significant advantages, particularly in the detection of small objects. The upsample layer doubles the size of feature maps without altering the output channels, merging feature maps generated at different resolutions. Finally, the head section, which enhances speed and performance, focuses on predicting object classes and bounding boxes using the final outputs generated by the backbone and neck sections. The head structure, developed with the C2f module, improves the accuracy of predictions while enabling faster inference. In this section, three detect blocks process the objects passed through the C2f module to perform detection. Figure 2 details the architecture of YOLOv8, which is based on deep learning principles.

2.2 Faster R-CNN

The Faster R-CNN is a two-stage object detection framework that integrates the region proposal process directly into the convolutional neural network, making it both faster and more efficient than previous models. The first stage involves the Region Proposal Network (RPN), which generates a set of potential bounding box proposals from the input image. The RPN scans the image at multiple scales and aspect ratios, producing a large number of candidate regions that might contain objects. These proposals are evaluated based on their likelihood of containing an object, and those with the highest scores are selected for further processing. In the second stage, the selected proposals are passed through a CNN for object classification and precise bounding box regression. This classification step assigns labels to the objects within the proposals, while the bounding box regression refines the coordinates of each proposed box to improve the localization accuracy. One of the key innovations of Faster R-CNN is its end-to-end trainable architecture, where the RPN and the object detection network share convolutional layers. This shared architecture allows for faster processing and eliminates the need for separate proposal generation algorithms, which were required in earlier models like R-CNN and Fast R-CNN. Additionally, Faster R-CNN employs the use of anchor boxes, predefined bounding box templates, to further enhance the accuracy of the region proposal process. This approach makes the framework highly effective in detecting objects of various sizes and aspect ratios in a wide range of images. Faster R-CNN's two-stage approach, while more computationally intensive than single-stage methods like YOLO, provides exceptional accuracy and is particularly useful in applications where high precision is required, such as in fine-grained object detection or detecting small objects. The framework's success has made it a foundational model in the field of computer vision, and it continues to be a benchmark for evaluating newer object detection algorithms.

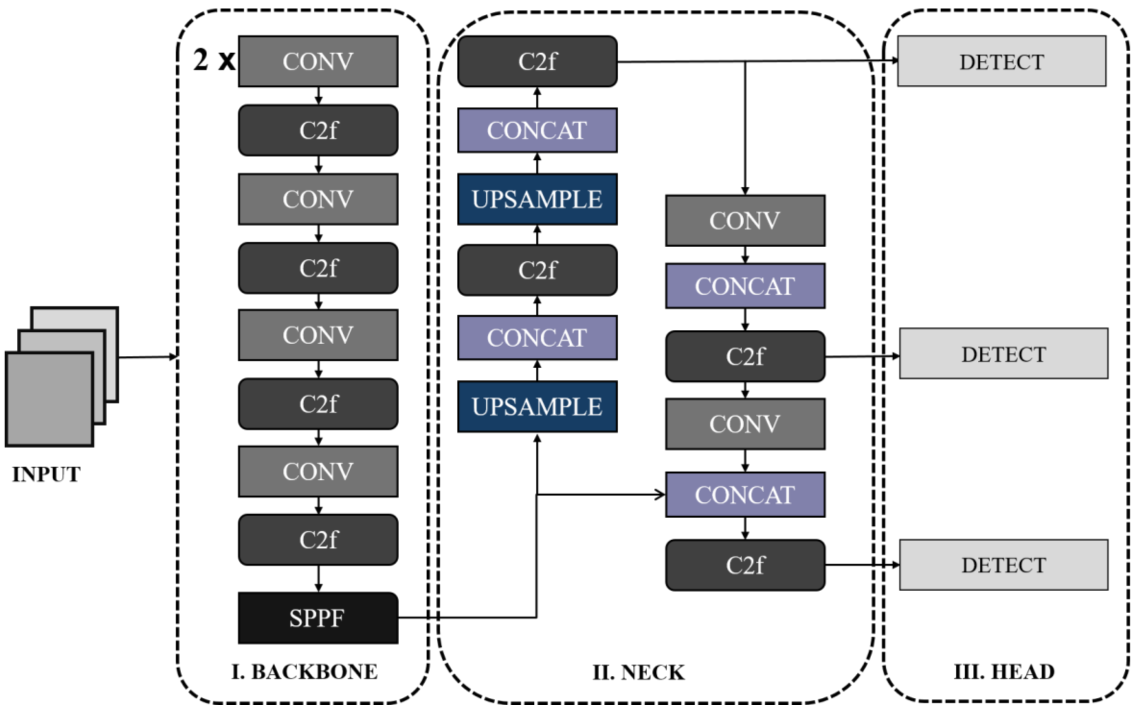

Figure 2. Detailed architecture of the YOLOv8 algorithm for object detection.

2.3 Performance Evaluation

The IoU (Intersection over Union) metric assigns a measure of accuracy to evaluate the performance of object detectors in localization. It calculates the overlap ratio between the predicted bounding box and the ground truth bounding box. High IoU values indicate better localization accuracy. IoU is a widely used metric for comparing the performance of various object detectors. This overlap ratio, which reflects the efficiency of the bounding box prediction and ranges between 0 and 1, is mathematically expressed as shown in Equation 1:

$$\mathrm{IoU} = \frac{\text{Intersection Area}}{\text{Union Area}} \quad (1)$$

However, the IoU metric has a limitation: it only functions optimally when bounding boxes overlap. This limitation prevents models from improving localization accuracy for non-overlapping bounding boxes. To address shortcomings in low-performance localization scenarios, the YOLOv8 algorithm employs the CIoU (Complete Intersection over Union) metric, which complements IoU.

The advanced CIoU metric includes additional features that optimize the center distance and aspect ratio of bounding boxes, alongside measuring the overlap. CIoU provides a more comprehensive assessment by not only measuring the intersection between the predicted and ground truth bounding boxes but also factoring in the distances between their centers and their dimensions. While CIoU offers a broader perspective, it comes with the trade-off of additional computational cost, particularly when working with large datasets and more complex calculations. The CIoU loss function is expressed as follows:

$$\mathrm{Loss}_{\mathrm{CIoU}} = 1 - \mathrm{IoU} + \frac{\rho^2(b, b^{gt})}{c^2} + \alpha v \tag{2}$$

In this equation, $\rho^2(b, b^{gt})$ represents the squared Euclidean distance between the center coordinates of the predicted bounding box and the ground truth bounding box. Reducing the center distance decreases the loss function and improves the model's positional accuracy. c is the diagonal distance of the smallest rectangle that encloses both the predicted and ground truth bounding boxes. An increase in c implies that errors have less impact on the model's overall performance.

The terms α and v, which optimize the dimensions of the bounding boxes, are defined as follows:

$$\alpha = \frac{v}{1 - \mathrm{IoU} + v} \tag{3}$$

$$v = \frac{4}{\pi^2}\left(\arctan + \frac{w^{gt}}{h^{gt}} - \arctan + \frac{w}{h}\right)^2 \tag{4}$$

In addition to these loss functions, Binary Cross-Entropy (BCE) is used to evaluate logarithmic losses for each class. BCE aims to achieve accurate class predictions and strengthen object classification. Furthermore, YOLOv8 incorporates the Distribution-Focused Loss (DFL) function to enhance localization accuracy in complex tasks.

2.4 Training the Algorithm

Before the tests, a UAV system was used to create a dataset for detecting fuel tanks. The dataset consists of images of fuel tanks captured from various locations. Ground conditions, camera angles, and altitudes were deliberately diversified during data collection to build a robust dataset. Particular attention was paid to capturing both vertical and angled images at various altitudes in real-world shooting scenarios. Additionally, the dataset includes both wide-angle shots and detailed, close-up images targeting specific objectives. By accounting for different conditions in the UAV-acquired images, the strengths of the network's performance were analyzed. The dataset comprises a total of 1,457 images containing 4,352 fuel tank annotations. The dataset was divided into training, validation, and test sets with proportions of 70%, 10%, and 20%, respectively.

To better illustrate the diversity and balance of the dataset, we present a detailed distribution of images based on key variables, as shown in Table 1. The dataset encompasses a range of conditions, with images captured under different lighting scenarios (33% sunny, 41% cloudy, 26% low-light), altitudes (spanning 10m to 120m), and angles (vertical, oblique, and side views). Additionally, object density per image varies, with 42% containing 1–2 fuel tanks, 36% with 3–5 tanks, and 22% featuring more than 5 tanks. We also accounted for obstacles such as pipelines, shadows, and surrounding infrastructure, ensuring a robust dataset for training and evaluating the YOLOv8 model. This comprehensive approach aims to enhance the model's generalization capacity and adaptability to real-world scenarios.

Table 1. Dataset distribution overview. The table illustrates the distribution of key variables in the dataset, including lighting conditions, altitudes, camera angles, object density per image, and the presence of obstacles, ensuring a diverse and representative set of training data for the YOLOv8 model.

Condition	Categories	Percentage (%)
Lightning Conditions	Sunny	33%
	Cloudy	41%
	Low-light	26%
Altitude (m)	10-50	45%
	51-90	35%
	91-120	20%
Camera Angles	Vertical	40%
	Oblique	35%
	Side view	25%
Object Density	1-2 fuel tanks per image	42%
	3-5 fuel tanks per image	36%
	More than 5 fuel tanks per image	22%
Presence of Obstacles	Pipelines, shadows, infrastructure	Included

2.5 UAV System Components

The object detection application for fuel tanks was systematically carried out through the integration of the YOLOv8 algorithm and a quad-rotor UAV. The UAV was configured and tailored to meet the requirements of the specified mission in the defined scenario. YOLOv8 was integrated with a Sony IMX industrial-grade camera system for high-resolution image processing, enabling precise and reliable detection. This optimized camera reduced critical challenges such as low-light conditions and occlusions that complicate object detection. A Pixhawk Cube Orange, an ArduPilot-based flight controller, was optimized to ensure controlled and stable flight. Thanks to its robust mechanisms, this flight controller provided reliable performance under various scenarios. The UAV's flight control system used the collaboration of an Inertial Measurement Unit (IMU) and Here 3 GPS sensor to process data, provide location information, and enhance navigation. The system employed the high-precision SkyWalker 80A ESC module to manage speed control based on incoming commands. An LTE module was integrated as the communication system, enabling real-time data transmission and instant notifications. Power was supplied using an XT60 Power Module, ensuring the UAV had sufficient energy to execute tasks in various scenarios, thereby increasing overall system efficiency. The components integrated into the UAV play a critical role in ensuring its functional performance. The hardware included an NVIDIA GeForce RTX 3050 GPU for processing and an Intel Xeon Silver 4116 CPU to handle computational tasks. Figure 3 illustrates the components of the unmanned aerial vehicle.

The hardware configuration for this study, consisting of an NVIDIA GeForce RTX 3050 GPU and an Intel Xeon Silver 4116 CPU, was chosen to balance computational power and portability for real-time fuel tank detection. While the RTX 3050 provides efficient parallel processing capabilities necessary for deep learning tasks, it consumes approximately 130W under full load. The Intel Xeon Silver 4116 CPU, with a thermal design power (TDP) of 85W, ensures reliable data processing and UAV control. Recognizing the importance of energy efficiency, future work will explore the integration of low-power alternatives, such as NVIDIA Jetson modules, to reduce energy consumption without compromising detection speed. Additionally, scalability considerations will involve testing the deployment of the YOLOv8 model on edge devices, assessing real-time performance and resource usage, and investigating dynamic power management techniques to enhance UAV endurance during prolonged operations. This approach aims to ensure a balance between high detection accuracy and practical energy efficiency for industrial applications.

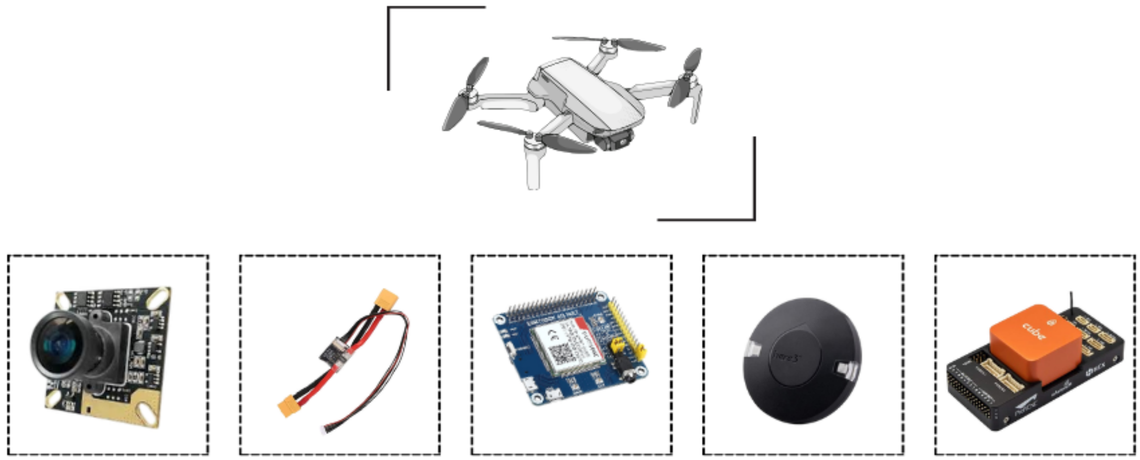

Figure 3. Overview of the key components in the UAV system.

3 Results

Through the experiments, fuel tank detection was performed using the YOLOv8 model in various scenarios, with the UAV's effective utilization, and the results obtained from test images were evaluated. Figure 4 shows some of the detection examples. The YOLOv8 algorithm analyzed in this study demonstrated superior performance by establishing a balance between speed and accuracy in detecting fuel tanks. However, limitations imposed by certain conditions led to some false detections. The accurate detection of fuel tanks in images captured by the UAV improved the algorithm's success rate, confirming its suitability for this specific application. The algorithm successfully detected object presence and provided location information across various scenes. However, in conditions involving intricate backgrounds and low-resolution images, the rate of false positives and false negatives increased. In images with complex backgrounds, such as those containing other industrial structures and similar colors within the same spectrum, it became more difficult for the YOLOv8 model to accurately detect fuel tanks. In some instances, detections were missed due to the similarity between fuel tanks and other objects nearby in the UAV-captured images.

Moreover, images of fuel tanks often included surrounding areas with shadows, pipelines, and other oil tanks, complicating detection. In certain scenarios, the algorithm failed to correctly detect existing objects. The challenges of resolution loss and blurring from images captured at different perspectives and angles caused additional masking effects, leading to false detections. These incorrect and missed detections undermine the reliability of the YOLOv8 algorithm. Especially in cases where object detection algorithms missed detections across different scenes, critical details were overlooked, posing a significant problem in applications requiring precise detection. Furthermore, detection of fuel tanks was more difficult when poor weather conditions prevailed during UAV filming. In such conditions, detection became more challenging in scenarios with shaded or cloudy weather, and detection of fuel tanks at long distances was complicated. Additionally, the presence of multiple similar fuel tanks in the same frame from a specific angle made object detection more complex, reducing the algorithm's performance. In particular, in images where tanks were very close to each other, misclassification of some objects occurred.

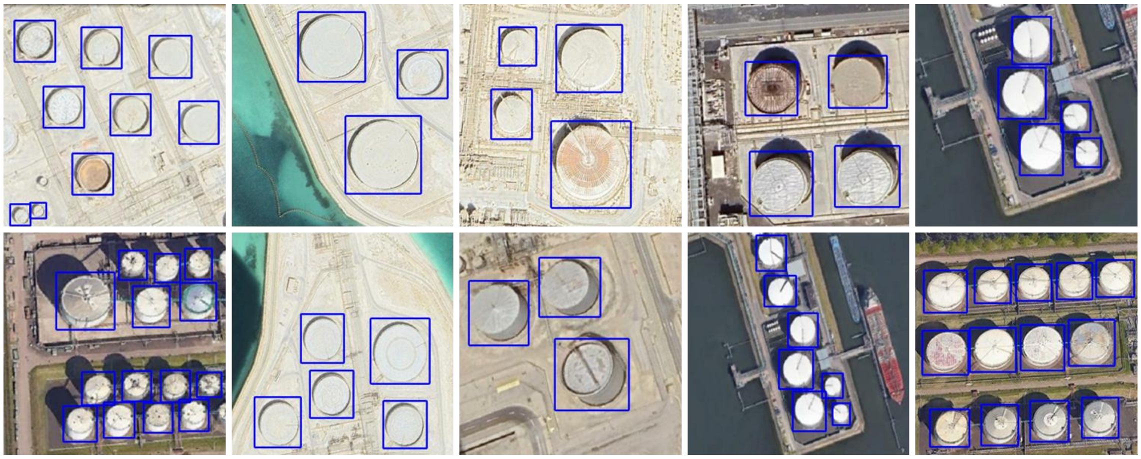

Figure 4. Representative examples of fuel tank detection results using the YOLOv8 algorithm.

Figure 5 shows the performance values of the YOLOv8 algorithm in fuel tank detection and its comparison with Faster R-CNN. The primary reason behind YOLOv8's superiority in object detection lies in its architectural innovations and the enhanced modules it incorporates. As a result of the tests, this model achieved a precision value of 0.888. This high precision is a result of its architecture, which is free from the complex structure of anchor boxes, and the advanced NMS support that prevents excessive predictions for the same object. Furthermore, the specialized head of the algorithm ensures better accuracy in classification and location detection, contributing to its superior precision. However, Faster R-CNN, a two-stage object detection process with strong feature extraction, outperforms YOLOv8 in this area. Unlike its predecessors, YOLOv8 uses PANet instead of FPN, positively influencing the Recall value, which measures the algorithm's ability to correctly identify positive examples. With the support of this module, YOLOv8 enhances the detection of smaller objects and achieves a higher Recall rate of 0.896. In comparison, Faster R-CNN, which involves anchor constraints, revealed YOLOv8’s superiority with a Recall value of 0.871.

The impact of detection errors, specifically false positives and false negatives, plays a critical role in industrial operations where fuel tank monitoring is essential for safety and efficiency. False positives, incorrectly identifying non-existent fuel tanks, could lead to unnecessary inspections, diverting resources and increasing operational costs. On the other hand, false negatives, failing to detect actual fuel tanks, pose a more severe risk by potentially overlooking structural issues or fuel leaks, compromising safety protocols and delaying crucial maintenance actions. These errors can disrupt real-time monitoring systems, emphasizing the need for a balanced model that minimizes both types of inaccuracies. Consequently, future iterations of this research will focus on enhancing YOLOv8’s error mitigation strategies, incorporating techniques such as confidence threshold tuning and ensemble modeling to improve detection reliability and ensure more informed decision-making processes.

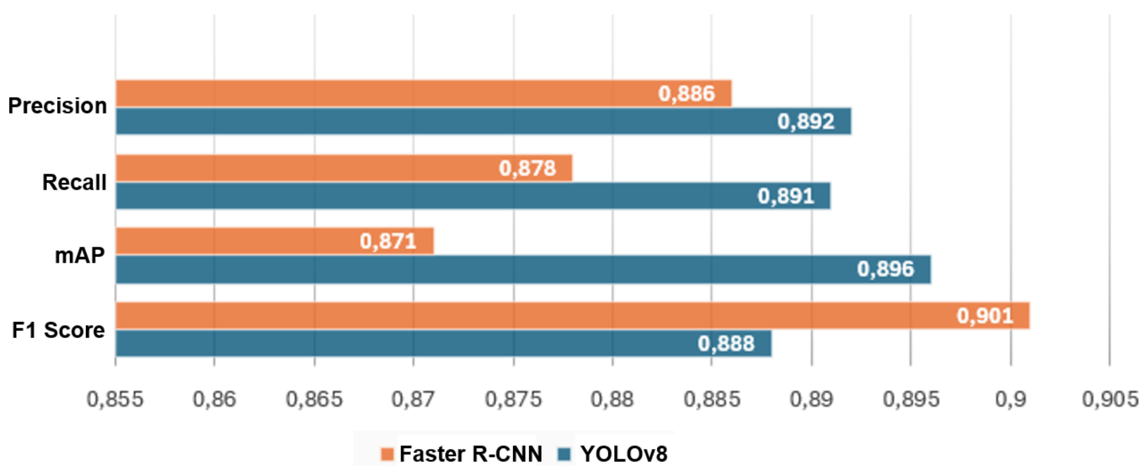

Figure 5. Performance comparison of YOLOv8 and Faster R-CNN algorithms for fuel tank detection.

One of the key metrics in performance evaluation, the mAP (mean Average Precision), is driven by the accuracy of object extraction at different scales, attributed to the SPPCSPC module. Additionally, the CIoU loss function, which elevates prediction precision, further improves YOLOv8's mAP value, resulting in a final mAP of 0.891. On the other hand, Faster R-CNN is likely to miss small objects or objects in complex conditions, resulting in an mAP value of 0.878. The F1 score, which bridges Recall and Precision, reached 0.892 for the YOLOv8 model, reflecting a more balanced and optimized training process. YOLOv8 delivers much faster performance compared to its previous iterations, with a response time of 41 ms. In contrast, Faster R-CNN, which executes the object recognition task in two steps, has a response time of 64 ms, making it less effective in real-time object detection scenarios. Given its superior balance of speed and accuracy, the YOLOv8 model is the preferred choice over Faster R-CNN, offering significant advantages.

4 Conclusion and Discussion

This study thoroughly examines the performance of the YOLOv8 object detection algorithm in detecting fuel tank images captured by UAVs from various angles and under different conditions. The agility of UAVs and their ability to capture high-quality images from different perspectives make them an effective tool for facilitating the detection of fuel tanks, which play a critical role in many industries. Together, UAVs and YOLOv8 have significantly advanced automation and precision in inspection processes within the energy and industrial sectors. The YOLOv8 model, a single-stage deep learning-based object detection algorithm, offers a superior architecture compared to previous YOLO iterations. It integrates speed and accuracy in the object detection process, showcasing its strengths in detection. With a high precision rate of 0.888, a recall value of 0.896, and an mAP value of 0.891, the model has demonstrated its strong capabilities in detecting fuel tanks, which support the sustainability of operations in the industrial and energy sectors. The object detector's fast processing time of 41 ms highlights its effectiveness in real-time applications that require rapid workflows. This study highlights the importance of optimizing UAVs and deep learning models, recognized for their reliable data collection capabilities even under challenging conditions, to efficiently monitor fuel tanks and ensure reliability across various sectors.

Future studies are expected to explore the adaptation of such high-performance algorithms to a wider range of object detection scenarios. Thus, future work will focus on enhancing the robustness of the YOLOv8 model for fuel tank detection by addressing current limitations such as false positives and undetected objects in real-world environments. This will involve fine-tuning the model with a more diverse and comprehensive dataset that includes various fuel tank types, environmental conditions, and varying degrees of occlusion. Additionally, the integration of data augmentation techniques, such as image rotation, scaling, and noise injection, will be explored to improve the algorithm's generalization ability. Another important avenue for future research will involve testing the model in real-world operational settings, such as industrial facilities and energy plants, to assess its performance under diverse environmental factors like varying lighting conditions, adverse weather, and crowded settings. Moreover, real-time deployment of the algorithm on UAVs will require further optimization for computational efficiency, minimizing both processing time and resource usage without compromising detection accuracy. Finally, an important direction will be exploring the adaptation of YOLOv8 for detecting other critical infrastructure elements beyond fuel tanks, thus broadening the applicability of UAV-based object detection in the monitoring and management of industrial assets.

To address the limitations observed in the current study, future optimization strategies will focus on fine-tuning the YOLOv8 model through advanced data augmentation techniques, such as random cropping, flipping, and CutMix, to enhance the model's ability to generalize across diverse scenarios. Additionally, transfer learning will be explored by leveraging pre-trained weights from large-scale object detection datasets, allowing the model to inherit robust feature representations. Another promising direction involves integrating attention mechanisms, such as the Efficient Channel Attention (ECA) module, to refine feature extraction and boost detection accuracy. We also plan to experiment with hyperparameter optimization using algorithms like Bayesian optimization to identify the most effective configurations for training. Finally, deploying model pruning and quantization methods will be considered to reduce computational overhead, ensuring the model's real-time applicability without compromising accuracy.

Moreover, future research will also explore the deployment of the YOLOv8 model on edge devices, such as NVIDIA Jetson Nano and Xavier modules, to enhance UAV autonomy and reduce reliance on high-powered hardware. These edge AI devices, designed for real-time inference with lower energy consumption, offer a practical solution for onboard processing. To achieve this, model compression techniques, including quantization and pruning, will be employed to reduce the size and computational requirements of YOLOv8 without compromising detection accuracy. Furthermore, we plan to evaluate the performance of lightweight YOLOv8 variants, such as YOLOv8n (nano), to identify optimal configurations for resource-limited environments. This approach aims to strike a balance between computational efficiency, detection speed, and UAV battery life, ultimately supporting real-time, autonomous fuel tank monitoring in industrial settings. While the study highlights the success of accurate detections, it also emphasizes that the occurrence of false detections or undetected objects suggests that further optimization of the algorithm is necessary for real-world applications. Furthermore, improving the algorithm's performance through training with datasets that feature diverse characteristics and testing under challenging environmental conditions is crucial.

Data Availability

The dataset used in this study is available upon request from the corresponding author. Due to licensing agreements, direct public access is restricted.

References

1. Rizk, M., Chehade, A. (2024). Efficient oil tank detection using deep learning: A novel dataset and deployment on edge devices. *IEEE Access*, 12, 170346-170378. https://doi.org/10.1109/ACCESS.2024.3495523
2. Bakirci, M. (2025). Advanced aerial monitoring and vehicle classification for intelligent transportation systems with YOLOv8 variants. *Journal of Network and Computer Applications*, 237, 104134. https://doi.org/10.1016/j.jnca.2025.104134
3. Sheng, H., Chen, G., Li, X., Men, J., Xu, Q., Zhou, L., Zhao, J. (2024). A novel unmanned aerial vehicle driven real-time situation awareness for fire accidents in chemical tank farms. *J. Loss Prev. Process Ind.*, 91(105357). https://doi.org/10.1016/j.jlp.2024.105357
4. Fan, L., Chen, X., Wan, Y., Dai, Y. (2022). Comparative analysis of remote sensing storage tank detection methods based on deep learning. *Remote Sensing*, 15(9), 2460, 2022. https://doi.org/10.3390/rs15092460
5. Xiao, Y., Tian, Z., Yu, Y. et al. (2020). A review of object detection based on deep learning. *Multimedia Tools and Applications*, 79, 23729–23791. https://doi.org/10.1007/s11042-020-08976-6
6. Bakirci, M. (2025). Vehicular mobility monitoring using remote sensing and deep learning on a UAV-based mobile computing platform. *Measurement*, 244, 116579. https://doi.org/10.1016/j.measurement.2024.116579
7. Liu, L., Ouyang, W., Wang, X. et al. (2020). Deep Learning for Generic Object Detection: A Survey. *International Journal of Computer Vision*, 128, 261–318. https://doi.org/10.1007/s11263-019-01247-4
8. Kaur, R., Singh, S. (2023). A comprehensive review of object detection with deep learning. *Digital Signal Processing*, 132(103812). https://doi.org/10.1016/j.dsp.2022.103812
9. Nie, G.T., Huang, H. (2021). A review of optical remote sensing image target detection algorithms. *J. Autom.*, 47, 1749–1768.
10. Bakirci, M., Bayraktar, I. (2024). The cutting-edge YOLO11 for advanced aircraft detection in synthetic aperture radar (SAR) imagery. In *2024 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS)* (pp. 1-6). IEEE Press. https://doi.org/10.1109/ISAS64331.2024.10845222
11. Ryan, A., Zennaro, M., Howell, A., Sengupta, R., Hedrick, J.K. (2004). An overview of emerging results in cooperative UAV control. In *2004 43rd IEEE Conference on Decision and Control* (pp. 602-607). IEEE Press.
12. Wang, G., Chen, Y., An, P., Hong, H., Hu, J., Huang, T. (2023). UAV-YOLOv8: A Small-Object-Detection Model Based on Improved YOLOv8 for UAV Aerial Photography Scenarios. Sensors, 23, 7190. https://doi.org/10.3390/s23167190
13. Mittal, P., Singh, R., Sharma, A. (2020). Deep learning-based object detection in low-altitude UAV datasets: A survey. *Image Vis. Comput.*, 104(104046).
14. Almaskati, D., Kermanshachi, S., Pamidimukkala, A., Loganathan, K., Yin, Z. (2024). A Review on Construction Safety: Hazards, Mitigation Strategies, and Impacted Sectors. *Buildings*, 14(526). https://doi.org/10.3390/buildings14020526

15. Mozaffari, M., Saad, W., Bennis, M., Nam, Y.H., Debbah, M. (2019). A Tutorial on UAVs for Wireless Networks: Applications, Challenges, and Open Problems. *IEEE Communications Surveys & Tutorials*, 21(3), 2334-2360. https://doi.org/10.1109/COMST.2019.2902862
16. Bakirci, M., Bayraktar, I. (2024). Assessment of YOLO11 for ship detection in SAR imagery under open ocean and coastal challenges. In *2024 21st International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE)* (pp. 1-6). IEEE Press. https://doi.org/10.1109/CCE62852.2024.10770926
17. Nex, F., Remondino, F. (2014). UAV for 3D mapping applications: a review. Appl Geomat, 6, 1–15. https://doi.org/10.1007/s12518-013-0120-x
18. Tsouros, D.C., Bibi, S., Sarigiannidis, P.G. (2019). A Review on UAV-Based Applications for Precision Agriculture. Information, 10, 349. https://doi.org/10.3390/info10110349
19. Bakirci, M., Bayraktar, I. (2024). Comparative performance of YOLOv9 and YOLOv10 for vehicle detection towards real-time traffic surveillance with UAVs. In *2024 21st International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE)* (pp. 1-6). IEEE Press. https://doi.org/10.1109/CCE62852.2024.10771048
20. Motlagh, N.H., Bagaa, M., Taleb, T. (2017). UAV-Based IoT Platform: A Crowd Surveillance Use Case. *IEEE Communications Magazine*, 55(2), 128-134.
21. Ju, C., Son, H.I. (2018). Multiple UAV Systems for Agricultural Applications: Control, Implementation, and Evaluation. Electronics, 7, 162. https://doi.org/10.3390/electronics7090162
22. Borji, A., Cheng, M.M., Hou, Q., Jiang, H., Li, J. (2019). Salient object detection: A survey. *Comput. Vis. Media*, 5, 117-150.
23. Bakirci, M., Dmytrovych, P., Bayraktar, I., Anatoliyovych, O. (2024). Multi-class vehicle detection and classification with YOLO11 on UAV-captured aerial imagery. In *2024 IEEE 7th International Conference on Actual Problems of Unmanned Aerial Vehicles Development (APUAVD)* (pp. 191-196). IEEE Press. https://doi.org/10.1109/APUAVD64488.2024.10765862
24. Lou, H., Duan, X., Guo, J., Liu, H., Gu, J., Bi, L., Chen, H. (2023). DC-YOLOv8: Small-Size Object Detection Algorithm Based on Camera Sensor. Electronics, 12, 2323. https://doi.org/10.3390/electronics12102323
25. Lee, J., Hwang, Ki. (2022). YOLO with adaptive frame control for real-time object detection applications. Multimed Tools Appl, 81, 36375–36396. https://doi.org/10.1007/s11042-021-11480-0
26. Bakirci, M., Cetin, M. (2023). Improving position-time trajectory accuracy in vehicle stop-and-go scenarios by using a mobile robot as a testbed. *Journal of Control Engineering and Applied Informatics*, 25(3), 35-44. https://doi.org/10.61416/ceai.v25i3.8365
27. Bakirci, M., Cetin, M. (2023). Reliability of MEMS accelerometers embedded in smart mobile devices for robotics applications. In: Garcia Marquez, F.P., Jamil, A., Eken, S., Hameed, A.A. (eds) Computational Intelligence, Data Analytics and Applications. ICCIDA 2022. *Lecture Notes in Networks and Systems*, vol 643, pp. 78-90, Springer, Cham. https://doi.org/10.1007/978-3-031-27099-4_7
28. Zhao, Z.-Q., Zheng, P., Xu, S.-T., Wu, X. (2019). Object Detection With Deep Learning: A Review. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11), 3212-3232.
29. Wu, X., Sahoo, D., Hoi, S.C.H. (2020). Recent advances in deep learning for object detection. *Neurocomputing*, 396, 39-64.
30. Bakirci, M. (2024). Enhancing vehicle detection in intelligent transportation systems via autonomous UAV platform and YOLOv8 integration. *Applied Soft Computing*, 164, 112015. https://doi.org/10.1016/j.asoc.2024.112015
31. Bakirci, M. (2024). Real-time vehicle detection using YOLOv8-nano for intelligent transportation systems. *Traitement du Signal*, 41(4), 1727-1740. https://doi.org/10.18280/ts.410407
32. Diwan, T., Anirudh, G., Tembhurne, J.V. (2023). Object detection using YOLO: Challenges, architectural successors, datasets and applications. *Multimed. Tools Appl.*, 82(6), 9243-9275.
33. Ma, G., Wu, M., Wu, Z., Yang, W. (2021). Single-shot multibox detector- and building information modeling-based quality inspection model for construction projects. *J. Build. Eng.*, 38(102216).
34. Huang, Z., Wang, J., Fu, X., Yu, T., Guo, Y., Wang, R. (2020). DC-SPP-YOLO: Dense connection and spatial pyramid pooling based YOLO for object detection. *Inf. Sci.*, 522, 241-258.
35. Hu, X., Liu, Y., Zhao, Z., Liu, J., Yang, X., Sun, C. et al. (2021). Real-time detection of uneaten feed pellets in underwater images for aquaculture using an improved YOLO-V4 network. *Comput. Electron. Agric.*, 185(106135).
36. Yin, Y., Li, H., Fu, W. (2020). Faster-YOLO: An accurate and faster object detection method. *Digit. Signal Process.*, 102(102756).
37. Bakirci, M., Bayraktar, I. (2024). Transforming aircraft detection through LEO satellite imagery and YOLOv9 for improved aviation safety. In *2024 26th International Conference on Digital Signal Processing and its Applications (DSPA)* (pp. 1-6). IEEE Press. https://doi.org/10.1109/DSPA60853.2024.10510106

38. Adade, R., Aibinu, A.M., Ekumah, B. et al. (2021). Unmanned Aerial Vehicle (UAV) applications in coastal zone management—a review. Environ Monit Assess, 193, 154. https://doi.org/10.1007/s10661-021-08949-8
39. Tian, Y., Yang, G., Wang, Z., Li, E., Liang, Z. (2020). Instance segmentation of apple flowers using the improved mask R–CNN model. *Biosyst. Eng.*, 193, 264-278.
40. Bakirci, M. (2025). Advanced ship detection and ocean monitoring with satellite imagery and deep learning for marine science applications. *Regional Studies in Marine Science*, 81, 103975. https://doi.org/10.1016/j.rsma.2024.103975
41. Ren, S., He, K., Girshick, R., Sun, J. (2016). Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(6), 1137-1149.
42. Redmon, J., Divvala, S., Girshick, R., Farhadi, A. (2015). You Only Look Once: Unified, real-time object detection. *arXiv*. [Online]. Available: https://arxiv.org/abs/1506.02640
43. Sohan, M., Ram, T.S., Reddy, R., Venkata, C. (2024). A review on YOLOv8 and its advancements. *Int. Conf. Data Intell. Cogn. Informatics*, 529-545, Springer, Singapore.