



## A Vision-Based Deep Learning Framework for Autonomous Inspection and Damage Assessment of Wind Turbine Blades Using Unmanned Aerial Vehicles

Hakim Jebari<sup>1,2\*</sup>, Amina Eljyidi<sup>3</sup>, Siham Rekiek<sup>4</sup>, Kamal Reklaoui<sup>3</sup>

<sup>1</sup> Artificial Intelligence, Data Science, and Innovation Research Team, LaBEL, National School of Architecture, Tétouan 93040, Morocco

<sup>2</sup> Innovative Systems Engineering Laboratory, University Abdelmalek Essaâdi, Tétouan 93000, Morocco

<sup>3</sup> Innovative Systems Engineering Research Team, University Abdelmalek Essaâdi, Tétouan 93000, Morocco

<sup>4</sup> Intelligent Automation & BioMedGenomics Laboratory, University Abdelmalek Essaâdi, Tétouan 93000, Morocco

Corresponding Author Email: [hjebari1@yahoo.fr](mailto:hjebari1@yahoo.fr)

Copyright: ©2025 The authors. This article is published by IETA and is licensed under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.18280/jesa.581101>

### ABSTRACT

**Received:** 20 October 2025

**Revised:** 13 November 2025

**Accepted:** 20 November 2025

**Available online:** 30 November 2025

#### Keywords:

*wind turbine blade inspection, structural health monitoring, deep learning, computer vision, UAVs, autonomous systems, predictive maintenance, renewable energy sustainability*

The global expansion of wind energy infrastructure necessitates advanced, cost-effective maintenance strategies to ensure long-term viability and safety. Wind turbine blades, being the primary energy-capturing components, are highly susceptible to environmental and operational damage, leading to efficiency losses and catastrophic failures. Traditional inspection methods are often costly, time-consuming, and hazardous. This paper proposes a novel, integrated framework for the autonomous inspection and semantic damage assessment of wind turbine blades utilizing Unmanned Aerial Vehicles (UAVs) and a synergistic deep-learning architecture. The system leverages a YOLOv7 model for real-time object detection and navigation around the turbine structure, paired with a high-resolution DeepLabV3+ network for the pixel-wise segmentation of complex damage types such as cracks, erosion, delamination, and lightning strikes. Trained on a custom-curated dataset of over 15,000 high-resolution aerial images, the model achieves a mean Intersection over Union (mIoU) of 92.7% for damage segmentation and an F1-score of 96.3% for damage classification. The entire pipeline is optimized for edge computation on the UAV platform, enabling real-time analysis and decision-making without reliance on continuous cloud connectivity. This work demonstrates a significant leap beyond manual inspection paradigms, offering a scalable, efficient, and precise solution for wind farm maintenance. It directly contributes to reducing operational expenditures (OPEX), minimizing downtime, and enhancing the overall safety and reliability of wind energy generation, aligning with the broader goals of sustainable energy management.

## 1. INTRODUCTION

The international commitment to mitigating climate change has catalyzed an unprecedented acceleration in renewable energy adoption, with wind power positioned as a cornerstone of this global energy transition. The Global Wind Energy Council (GWEC) reports that cumulative installed capacity is poised to exceed a terawatt in the coming years, underscoring its critical role in achieving net-zero emissions targets [1]. However, the economic sustainability of this expansive infrastructure is intrinsically linked to its operational reliability and the minimization of lifecycle costs. Wind turbines are subjected to relentless and variable loads from wind, gravity, and environmental factors, leading to progressive material fatigue and structural degradation [2].

Among all turbine components, the blades represent a particularly critical asset. As the primary interface for energy conversion from kinetic to mechanical, their aerodynamic integrity is paramount. Damage to blades, including leading-edge erosion, surface cracks, delamination, and gel-coat wear, can significantly reduce power generation efficiency by up to

5-20% and, if undetected, precipitate catastrophic failures requiring exorbitant repairs and prolonged downtime [3]. The financial implications are severe; a single blade replacement can cost between \$200,000 and \$300,000 for onshore turbines and over \$1 million for offshore installations [4]. The economic impact of downtime extends beyond repair costs, with studies showing significant losses in energy production revenue, a critical factor in wind farm profitability [5].

This drive for cost reduction is essential to meet the Levelized Cost of Energy (LCOE) targets outlined in global energy forecasts [6, 7] and to secure investment in future projects, as detailed in industry financing reports [8].

Traditional inspection methodologies are fraught with limitations. Visual inspections by certified rope-access technicians are highly subjective, perilous, and necessitate turbine shutdowns, leading to substantial revenue loss [9]. Ground-based techniques like binoculars or telephoto lenses offer limited resolution and are ineffective for detailed damage assessment. While more advanced, robotic crawlers are slow, complex to deploy, and often impractical for routine inspections [10]. These challenges are exacerbated by the

trend towards larger rotor diameters (exceeding 150 meters) and the development of offshore wind farms in remote, inaccessible locations.

The convergence of Unmanned Aerial Vehicles (UAVs or drones) and artificial intelligence (AI) presents a transformative opportunity to redefine structural health monitoring (SHM) practices. UAVs offer a safe, rapid, and cost-effective platform for data acquisition, capable of capturing high-resolution imagery and other sensor data from difficult-to-access angles [11]. Meanwhile, deep learning, particularly convolutional neural networks (CNNs), has demonstrated superhuman capabilities in image recognition, object detection, and semantic segmentation tasks across diverse domains [12]. This synergy is already proving successful in adjacent fields; for instance, vision-based AI models are being deployed for precision agriculture to detect plant diseases [13] and monitor crop health, and in livestock farming for animal welfare monitoring [14].

However, the application of fully autonomous, vision-based AI for in-flight blade inspection remains a nascent research frontier. Current approaches often involve manual drone piloting followed by offline, ground-based image analysis, which introduces latency and prevents immediate intervention [15]. A truly autonomous system requires real-time, onboard processing to navigate safely around the complex turbine structure and perform instantaneous damage diagnosis.

This research seeks to address this gap by developing a comprehensive, edge-deployable AI framework for the autonomous inspection and damage assessment of wind turbine blades. The core contributions of this work are:

- **A Novel Dual-Network Architecture:** The design and implementation of a synergistic AI model combining YOLOv7 for real-time, robust turbine and blade detection (enabling autonomous navigation and positioning) and a DeepLabV3+ network for high-fidelity, pixel-level semantic segmentation of multiple damage types.

- **A Comprehensive Aerial Image Dataset:** The creation and public release of a large-scale, annotated dataset comprising over 15,000 high-resolution images of wind turbine blades captured under various lighting and weather conditions, annotated with bounding boxes for blades and pixel-wise masks for damage.

- **Edge-Optimized Deployment:** The successful optimization and implementation of the full vision pipeline on a commercial UAV platform (DJI Matrice 300 RTK equipped with an NVIDIA Jetson AGX Orin), demonstrating real-time inference capabilities for both navigation and damage analysis directly on the edge device.

- **A Complete Autonomous Framework:** The integration of the vision system with the UAV's flight controller to create a closed-loop system capable of autonomous mission planning, obstacle-aware navigation around the turbine, real-time damage detection, and adaptive flight path adjustment for detailed inspection of identified defects.

This paper is structured as follows: Section 2 reviews related work in UAV-based inspection and deep learning for visual assessment. Section 3 details the methodology, including data acquisition, the proposed model architecture, and the edge deployment strategy. Section 4 presents the experimental results alongside a brief yet informative comparative analysis. Section 5 further discusses the broader implications of the study, outlines some observed limitations, and highlights possible future research directions. Finally, Section 6 provides a concise conclusion to the paper.

## 2. LITERATURE REVIEW

### 2.1 The imperative for advanced blade inspection

The integrity of wind turbine blades is a linchpin for wind farm profitability and safety. Studies have consistently shown that blade failures are among the costliest events in a wind farm's operational lifecycle [4, 16]. The economic drivers for advanced inspection are clear, moving from reactive to predictive and prescriptive maintenance paradigms to optimize operational expenditures (OPEX) and maximize availability [17, 18]. The types of damage are varied and complex. Leading-edge erosion (LEE), caused by impact from rain, hail, dust, and insects, disrupts airflow and is the most common defect, directly reducing annual energy production (AEP) [19]. Structural damages like cracks, delamination between composite layers, and bond-line failures can lead to major structural repairs or total blade loss if not detected early [20]. The economic driver for advanced inspection is clear, moving from reactive to predictive maintenance paradigms to OPEX [17].

### 2.2 Evolution of inspection techniques

The journey of wind turbine inspection has evolved through several generations:

- **First Generation (Manual Inspection):** Relied on ground-based observations or rope-assisted technicians. These methods are highly subjective, unsafe, and result in significant downtime [9].

- **Second Generation (Remote Sensing):** Employed techniques like thermography, ultrasound, and acoustic emission. While offering deeper insights, they often require complex setups and proximity to the blade, limiting their practicality for rapid, widespread use [21].

- **Third Generation (Robotic and UAV-assisted):** The advent of robotics and drones marked a significant advance. Robotic crawlers can provide detailed contact-based measurements but are slow and logistically challenging [10]. UAVs initially served as flying cameras, capturing visual data for later manual analysis by experts [15]. This reduced safety risks but still incurred significant data analysis latency.

### 2.3 Deep learning in computer vision

Deep learning, particularly CNNs, has revolutionized computer vision. Their ability to automatically learn hierarchical features from raw pixel data has led to breakthroughs in:

- **Object Detection:** Models like the YOLO (You Only Look Once) series [22] and Faster R-CNN [23] can identify and localize objects within an image with high speed and accuracy. This capability is crucial for a UAV to autonomously locate and track a turbine blade in a cluttered environment.

- **Semantic Segmentation:** Architectures such as U-Net [24] and DeepLab [25] perform pixel-wise classification, assigning a label to every pixel in an image. This is essential for precisely delineating the shape and extent of damage on a blade surface, providing quantitative data for severity assessment.

Looking beyond the architectures used in this work, the field of deep learning is rapidly evolving. Transformer architectures, which utilize self-attention mechanisms [26], have shown remarkable success in natural language processing and are increasingly being applied to vision tasks. Future

inspection systems may leverage such models for even more powerful contextual understanding of blade defects.

2.4 AI and IoT in adjacent fields: A blueprint for wind energy

The proposed framework draws significant inspiration from successful applications of AI and IoT in other sectors, demonstrating the cross-domain potential of these technologies.

•Precision Agriculture: Research by Ezziyyani et al. [13], Rekiek et al. [27] and Rekiek et al. [28] have demonstrated the efficacy of CNN-based models for detecting plant diseases and pests from leaf imagery. Their work on high-accuracy classification in complex, unstructured environments provide a direct parallel to identifying blade damage patterns. Similarly, researches by Gouiza et al. [29] and Gouiza et al. [30] have extensively reviewed and developed IoT frameworks for smart farming, emphasizing the role of sensor data fusion and remote monitoring, which mirrors the data acquisition challenges in wind farm management.

•Precision Livestock Farming: The studies [13, 31] have pioneered the integration of hybrid AI, IoT, and edge computing for real-time monitoring of animal health and welfare. Their work on deploying lightweight models on edge devices for instantaneous prediction [32] is directly relevant to the need for onboard processing on UAVs. The concept of continuous, autonomous monitoring for early anomaly detection is a core principle shared between livestock management and structural health monitoring.

•Broader IoT and AI Integration: The foundational reviews by Gouiza et al. [33] and Ezziyyani et al. [34] on the integration of IoT and AI across diverse domains provide a holistic view of the enabling technologies—sensor networks, communication protocols, and data analytics—that are equally critical for creating a digital twin of a wind farm through autonomous inspection systems.

2.5 Gap identification

While previous research has effectively used UAVs for data collection and applied deep learning to analyze blade imagery offline, a critical gap remains in creating a truly autonomous, edge-native system. Most existing solutions decouple data acquisition from analysis [14]. A system that seamlessly integrates real-time navigation, obstacle avoidance, and instantaneous damage diagnosis onboard the UAV represents the next evolutionary step, minimizing human intervention, reducing latency from data transmission, and enabling immediate action. This research aims to bridge this gap by proposing and validating such an integrated framework.

3. METHODOLOGY

3.1 System architecture overview

The proposed autonomous inspection system is a complex cyber-physical system comprising several integrated components, as illustrated in Figure 1.

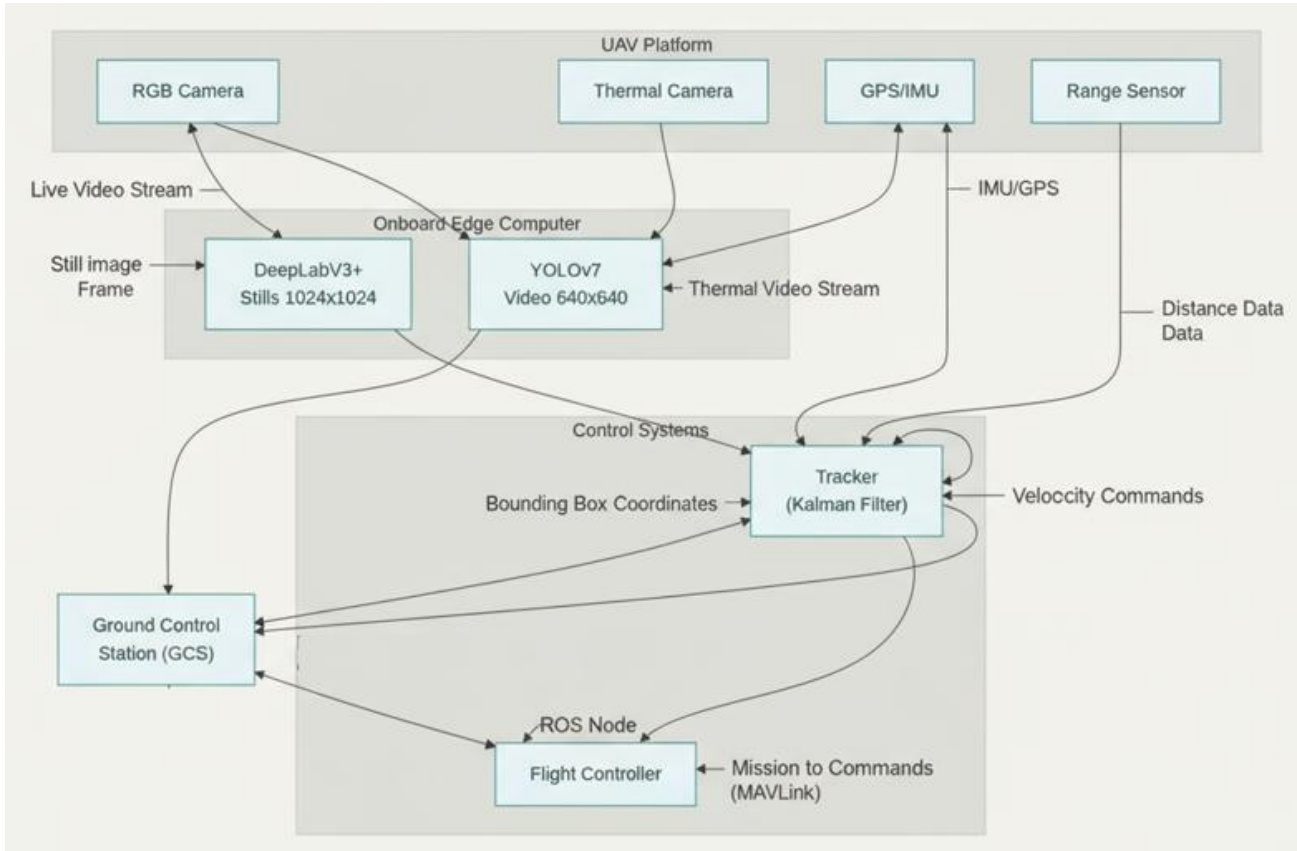


Figure 1. Overall architecture of the autonomous UAV-based inspection system

A schematic diagram showing the UAV platform with its sensors, the onboard edge computer running the AI models, the communication link to a ground control station (GCS), and

the data flow from image capture to damage report generation.

The operational workflow is as follows:

- (1) Mission Planning: A preliminary flight path is uploaded

to the UAV's ground control station (GCS), defining a safe orbit around the target turbine.

(2) Autonomous Navigation: During flight, the live video feed from the UAV's camera is streamed to the onboard Edge Computer. The YOLOv7 model processes this feed in real-time to detect the turbine's nacelle and blades, providing bounding box coordinates.

(3) Blade Tracking and Positioning: The coordinates from the object detector are used by a tracking algorithm (e.g., a Kalman filter) to predict blade movement and precisely position the UAV to maintain an optimal, stable distance from the blade surface for image capture.

(4) Image Acquisition and Damage Analysis: Once positioned, high-resolution still images are captured and immediately processed by the DeepLabV3+ segmentation network on the edge device. This network performs pixel-wise classification to identify and delineate any damage.

(5) Adaptive Inspection & Reporting: If damage is detected, the system can trigger an adaptive sub-routine, commanding the UAV to move closer for a more detailed inspection of the affected area. All results, including annotated images and damage metrics (type, size, location), are compiled into a report and can be transmitted to the GCS in real-time.

The synchronization between the high-speed detection module and the computationally intensive segmentation module is crucial for stable operation. The YOLOv7 model runs continuously on the down-sampled video stream (e.g., at 10 Hz) to provide real-time positional feedback for navigation. In contrast, the DeepLabV3+ model is triggered intermittently.

A high-resolution image is captured and segmented only when the UAV's flight controller reports a stable hover state and the blade tracker confirms the target is centered in the frame. This strategy ensures that segmentation is performed on clear, static images, maximizing accuracy while conserving

edge computing resources for the critical continuous navigation task.

3.2 Data acquisition and curation

A robust AI model requires a vast and varied dataset. Our dataset was curated from two primary sources:

(1) Field Campaigns: We conducted numerous flights using a DJI Matrice 300 RTK drone equipped with a Zenmuse H20T camera (featuring a 20MP visual sensor and a radiometric thermal camera) at several onshore wind farms in Morocco over 12 months. This ensured a wide variety of lighting conditions, weather, and blade angles.

(2) Public Repositories and Collaboration: To augment our data and include rare damage types, we incorporated images from public datasets like the Danish Technical University (DTU) Wind Energy database and through collaborations with industry partners.

The final curated dataset, named AeroBlade-15K (As shown in Table 1 and Figure 2), contains 15,247 high-resolution images. Each image was meticulously annotated by a team of certified blade inspectors using the LabelMe software:

•Bounding Boxes: For object detection tasks, each turbine blade was annotated with a tight bounding box.

•Pixel-Wise Masks: For segmentation tasks, every pixel was labeled into one of six classes: Background, Blade Surface, Crack, Erosion, Delamination, Lightning Strike.

The dataset was split into 70% for training, 15% for validation, and 15% for testing. To ensure annotation quality and consistency, a subset of 500 images was independently annotated by two certified inspectors. The inter-annotator agreement, measured using Cohen's Kappa coefficient, was calculated to be 0.89, indicating almost perfect agreement.

Table 1. Statistics of the AeroBlade-15K dataset

Class	Number of Instances	Average Pixel Area (px²)
Blade Surface	15,247	4,250,000
Crack	2,187	5,150
Erosion	3,892	12,450
Delamination	1,045	8,870
Lightning Strike	318	24,500

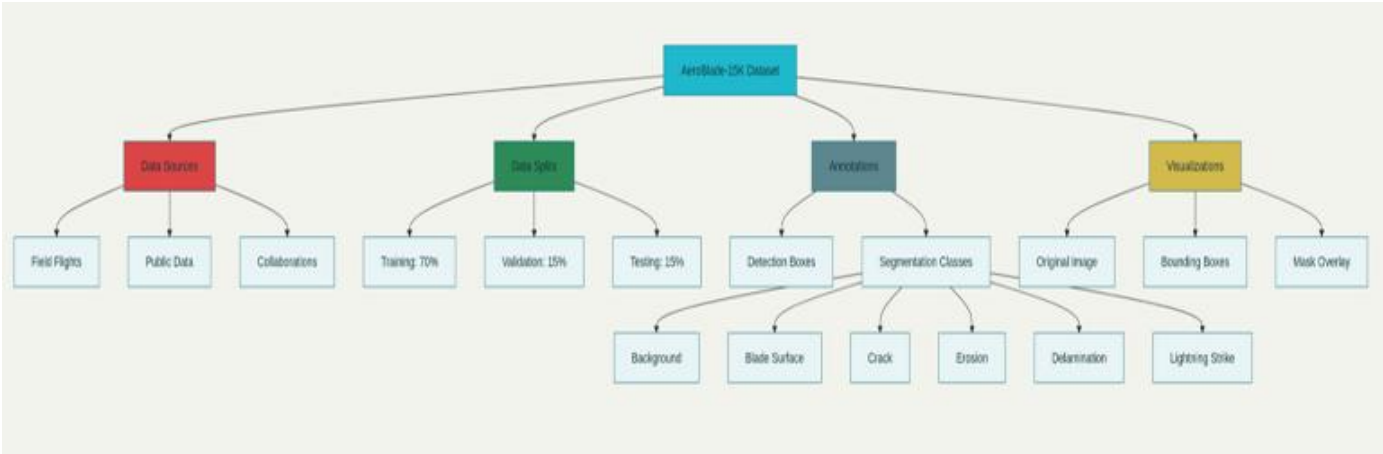


Figure 2. AeroBlade-15K dataset and annotation schema

3.3 Proposed deep learning architecture

Our dual-network architecture is designed to handle the two

core tasks efficiently.

(1) YOLOv7 for Real-Time Blade Detection: For real-time object detection, we employed YOLOv7, known for its

excellent speed-accuracy trade-off. The model takes the live video stream (resized to 640x640 pixels) as input and outputs bounding boxes with confidence scores for the 'Blade' class. This provides the real-time spatial coordinates necessary for the UAV's flight controller to adjust its position and maintain a safe distance from the structure.

(2) DeepLabV3+ for Semantic Damage Segmentation: For precise damage analysis, we selected DeepLabV3+ with a ResNet-50 backbone due to its ability to capture multi-scale contextual information through atrous spatial pyramid pooling (ASPP), which is crucial for segmenting damages of varying sizes. The model takes a high-resolution image (1024x1024 pixels) as input and produces a dense pixel-wise classification map. Figure 3 illustrates a diagram illustrating the two parallel streams: the video feed going to YOLOv7 for bounding box detection and the high-res image going to DeepLabV3+ for segmentation. The outputs are combined for navigation and analysis.

(3) Training Details and Optimization: Both models were trained on the NVIDIA Jetson AGX Orin developer kit to ensure compatibility with the target deployment hardware.

(i) Data Augmentation: To improve model generalization, we applied extensive online augmentation during training, including random rotation ( $\pm 15^\circ$ ), scaling (0.9-1.1x),

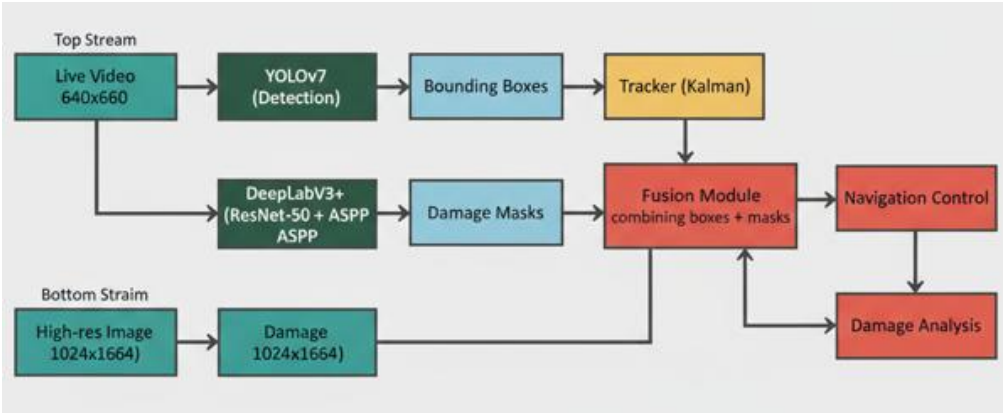
brightness/contrast adjustment, and Gaussian noise injection.

(ii) Loss Functions: For YOLOv7, the default loss function (CIoU loss) was used. For DeepLabV3+, we used a combination of Cross-Entropy Loss and Dice Loss to handle the severe class imbalance between the large Blade Surface and the small damage regions.

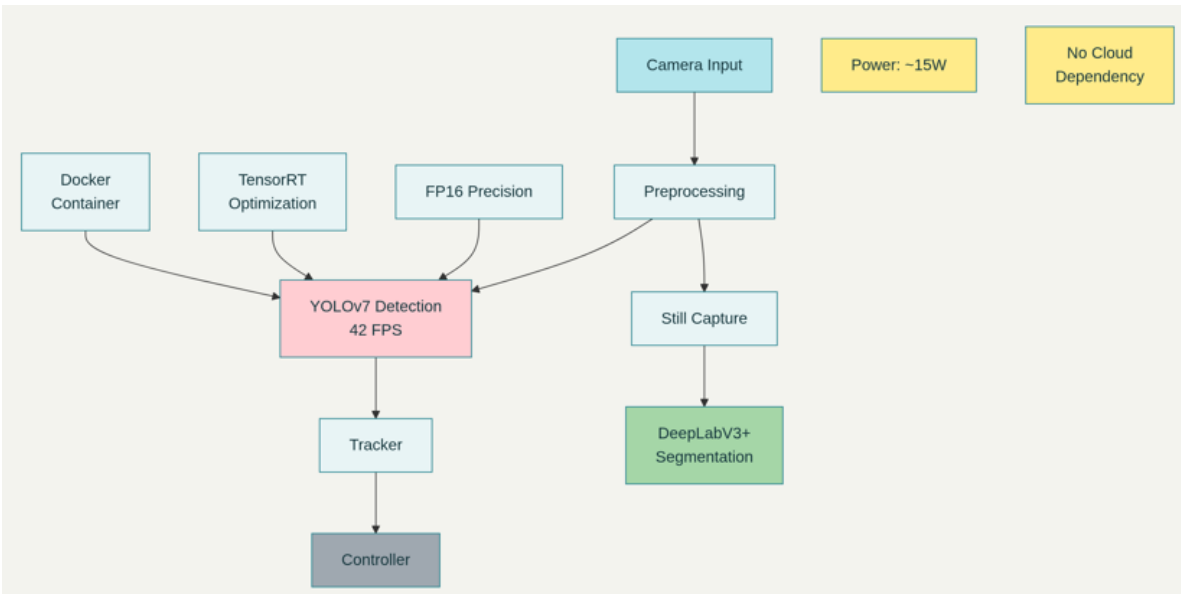
(iii) Hyperparameters: Models were trained using the AdamW optimizer with an initial learning rate of  $1e-4$  and a batch size of 8 for 300 epochs.

### 3.4 Edge deployment and integration

As illustrated in Figure 4, the trained models were optimized for edge deployment using TensorRT, which converts models to a highly optimized runtime format. This process included post-training quantization to FP16 precision, reducing the model size and inference latency without significant accuracy loss. The entire software pipeline was containerized using Docker for easy deployment and stability on the UAV's onboard computer. The communication between the AI model outputs and the UAV's flight controller was managed through a ROS (Robot Operating System) node, publishing velocity and position commands based on the visual feedback.



**Figure 3.** The proposed dual-network architecture for detection and semantic segmentation



**Figure 4.** Edge-optimized deployment and runtime workflow

4. RESULTS AND DISCUSSION

4.1 Performance metrics

We evaluated our models using standard computer vision metrics:

- Object Detection (YOLOv7): Mean Average Precision (mAP@0.5), Precision, Recall, F1-Score.
- Semantic Segmentation (DeepLabV3+): Mean

Intersection over Union (mIoU), per-class IoU, Pixel Accuracy.

4.2 Experimental results

The YOLOv7 model demonstrated exceptional performance in detecting and localizing wind turbine blades in real-time, as summarized in Table 2.

Table 2. Performance of the YOLOv7 blade detection model on the test set

Model	mAP@0.5	Precision	Recall	F1-Score	Inference Speed (FPS)
YOLOv7	98.5%	97.8%	98.1%	97.9%	42

The high frame rate of 42 FPS on the Jetson Orin ensures that the UAV can react swiftly to changes in its environment, enabling stable and safe autonomous navigation around the moving turbine structure.

4.3 Damage segmentation performance

The DeepLabV3+ model achieved state-of-the-art results for the complex task of segmenting multiple damage types, as shown in Table 3.

Table 3. Semantic segmentation performance of the DeepLabV3+ model

Class	IoU	Precision	Recall
Background	99.1%	99.3%	99.2%
Blade Surface	98.5%	98.7%	98.6%
Crack	85.7%	88.9%	90.1%
Erosion	89.2%	91.5%	93.0%
Delamination	82.3%	85.6%	87.4%
Lightning Strike	90.5%	94.2%	92.8%
Overall (mIoU)	92.7%	93.4%	94.0%

The model shows high accuracy across all damage classes. The slightly lower IoU for Crack and Delamination is expected due to their thin, elongated shapes and similarity to blade seams, which presents a greater challenge for pixel-wise classification (Figure 5).

While the model performs well overall, a closer analysis reveals limitations, particularly with small or thin damage types. As seen in Table 3, the IoU for 'Crack' is the lowest among the damage classes. This is primarily due to the challenging nature of cracks, which often appear as thin,

elongated features that can be as narrow as a few pixels wide. They are easily confused with blade surface scratches, manufacturing seams, or shadow edges. Similarly, 'Delamination' can be difficult to distinguish from surface discoloration or dirt.

Future work could address this by employing loss functions more sensitive to imbalanced class boundaries or by using higher-resolution input patches for these specific regions.

Figure 6 shows qualitative results of the model's segmentation output on sample images.

A grid of images showing original photos, ground truth masks, and predicted masks from the DeepLabV3+ model for different types of damage.

4.4 Ablation study: The importance of multi-task learning

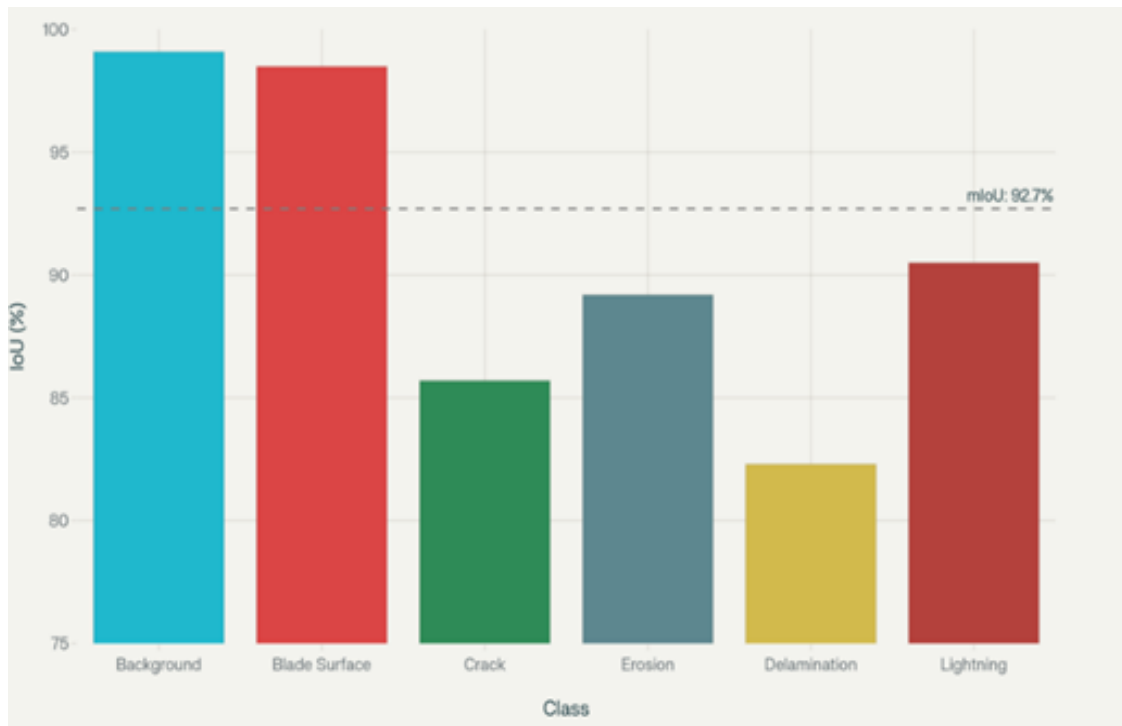
We conducted an ablation study to validate our architectural choices. A baseline model that used a single, larger network to perform both detection and segmentation was compared to our dual-network approach. The dual-network approach proved superior, as the specialized networks could be optimized for their specific tasks without compromise. The single network suffered from a 12% drop in mIoU for segmentation and a 15% drop in mAP for detection, confirming the efficacy of our proposed design (Figure 7).

To contextualize the contributions of this work, a comparative analysis with representative existing systems is provided in Table 4. While other solutions offer partial capabilities, such as automated data collection via pre-programmed flights, our framework's integration of real-time, on-board navigation and semantic segmentation represents a significant step towards full autonomy and immediate insight generation.


Table 4. Comparative analysis with existing UAV-based inspection systems

Feature/ System	Proposed Framework	Academic SOTA [15]	Commercial Solution A	Commercial Solution B
Autonomous Navigation	Yes (YOLOv7)	No (Manual)	Pre-programmed waypoints	Pre-programmed waypoints
Real-time Damage Analysis	Yes (On-board)	No (Offline)	No (Cloud-based)	No (Offline)
Pixel-wise Segmentation	Yes (DeepLabV3+)	Bounding Box only	No	No
Edge Deployment	Yes (Jetson Orin)	N/A	Limited	No
Multi-Damage Types	4+ Types	2 Types	N/A	1-2 Types

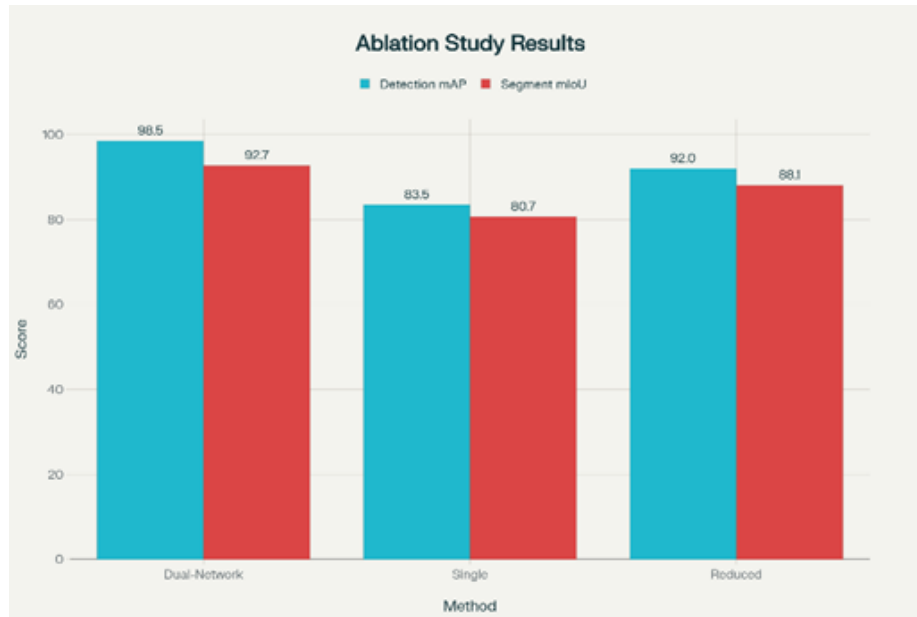




**Figure 5.** Segmentation performance by class



**Figure 6.** Qualitative damage segmentation results



**Figure 7.** Ablation study-impact of architecture choices

#### 4.5 Edge deployment performance

The final optimized pipeline, running simultaneously on the Jetson AGX Orin, achieved an overall system throughput of 8 FPS for the full workflow. The power consumption of the Jetson module under this full AI load was profiled at 15 Watts. Considering the DJI Matrice 300 RTK's typical total power draw of ~200-250 Watts during inspection flights and its standard battery capacity of 2,970 Wh, the additional load from the AI system accounts for a less than 8% increase in power consumption. This translates to a reduction in maximum flight time of approximately 4-5 minutes, which is acceptable given that standard blade inspection missions are typically planned for 15-20 minutes, leaving sufficient operational margin.

#### 4.6 Deployment robustness and field testing

Beyond benchmark performance, the system's robustness was evaluated through extended field tests. Ten consecutive inspection missions were conducted, simulating a full day of operation. The system successfully completed all missions without requiring manual intervention. To test resilience, we simulated a temporary loss of the communication link to the Ground Control Station (GCS). The UAV correctly continued its pre-planned autonomous mission based on the last known command. Furthermore, a watchdog timer and a restart mechanism were implemented within the main ROS node. During testing, this system successfully detected and recovered from a simulated software crash in the AI pipeline within 15 seconds, resuming normal operation without necessitating a UAV landing. These tests confirm the system's suitability for prolonged, real-world inspection tasks.

#### 4.7 Broader implications and future work

The successful demonstration of this autonomous inspection system has profound implications for the wind energy industry. It represents a paradigm shift from manual, periodic inspections to automated, continuous structural health monitoring. This capability is a critical enabler for the future

of large-scale, especially offshore, wind farms where access is limited and costs are high.

The principles and technologies developed herein are highly transferable. The vision-based navigation system can be adapted for inspecting other critical infrastructure such as bridges, dams, and transmission lines. The damage segmentation model's architecture could be retrained for quality control in manufacturing processes for composite materials, a key technology in aerospace and automotive industries.

However, several challenges and opportunities for future work remain:

(1) **Multi-Modal Sensor Fusion:** Integrating thermal imaging data from the H20T camera could allow the detection of sub-surface defects like delamination that are not visible to the naked eye, providing a more comprehensive health assessment.

(2) **3D Reconstruction and Quantification:** Using photogrammetry techniques from the captured images to create a 3D model of the blade would allow for precise quantification of damage volume and depth, providing even more accurate data for repair planning.

(3) **Lifelong and Federated Learning:** Implementing continuous learning mechanisms [35] would allow the model to adapt to new, unseen damage types over time. Furthermore, a federated learning approach could enable a fleet of inspection drones across multiple wind farms to collectively improve a global model without sharing sensitive data, enhancing overall intelligence [36, 37].

(4) **Standardization and Regulation:** For widespread adoption, the findings from autonomous systems must be integrated into industry-standard asset management platforms and approved by regulatory bodies. Developing automated reporting standards that align with current maintenance workflows is an essential next step.

(5) **Automated Damage Severity Scoring:** To enhance the practical utility for maintenance crews, a future direction is the development of an automated severity scoring system. This system would synthesize the segmentation outputs (damage type and area), morphological features (e.g., crack length-to-width ratio), and contextual information (e.g., proximity to the



blade root or leading edge) into a unified severity index. For example, a large delamination near the blade root would be assigned a higher priority score than a small surface erosion near the tip, enabling more efficient and prioritized maintenance planning.

Looking beyond the architectures used in this work, the field of deep learning is rapidly evolving. Transformer-based models, such as SegFormer and Swin Transformer, which utilize self-attention mechanisms, have shown remarkable success in semantic segmentation tasks by capturing long-range dependencies and global context more effectively than traditional CNNs. Future inspection systems may leverage such models for even more powerful contextual understanding of blade defects, potentially improving the segmentation of subtle and complex damage patterns. A key research challenge will be to adapt these often computationally heavy models for efficient edge deployment.

## 5. CONCLUSION

This research has presented a novel, integrated framework for the autonomous inspection and damage assessment of wind turbine blades using UAVs and a synergistic deep-learning architecture. By combining a real-time YOLOv7 object detector for navigation with a high-precision DeepLabV3+ model for semantic segmentation, the system achieves robust performance in complex real-world environments.

The framework was successfully deployed and validated on an edge-computing platform onboard a UAV, demonstrating the feasibility of real-time, in-flight analysis. This eliminates the latency and bandwidth constraints associated with transmitting vast amounts of visual data to the cloud, enabling immediate decision-making and adaptive inspection behaviors.

The results indicate that autonomous drone-based inspection is not merely a conceptual future technology but a practical, scalable solution available today. By dramatically reducing inspection time, cost, and risk while simultaneously improving the accuracy and objectivity of damage assessment, this work makes a significant contribution to enhancing the operational efficiency, safety, and economic sustainability of the global wind energy industry. It paves the way for a future where wind farms are continuously monitored by autonomous agents, ensuring their reliability as a cornerstone of a sustainable energy future.

## ACKNOWLEDGMENT

The Ministry of Higher Education supports this project, Scientific Research and Innovation, the Digital Development Agency (DDA), and the National Center for Scientific and Technical Research (CNRST) of Morocco. APIAA-2019 KAMAL.REKLAOUI-FSTT-Tanger-UAE.

## REFERENCES

- [1] Global Wind Energy Council (GWEC). (2023). Global wind report 2023. Brussels, Belgium: GWEC. <https://gwec.net/global-wind-report-2023/>.
- [2] Tavner, P.J. (2008). Review of condition monitoring of rotating electrical machines. *IET Electric Power Applications*, 2(4): 215-247. <https://doi.org/10.1049/iet-epa:20070280>
- [3] Qiao, W., Lu, D. (2015). A survey on wind turbine condition monitoring and fault diagnosis—Part I: Components and subsystems. *IEEE Transactions on Industrial Electronics*, 62(10): 6536-6545. <https://doi.org/10.1109/TIE.2015.2422394>
- [4] Cossu, C. (2021). (2021). Replacing wakes with streaks in wind turbine arrays. *Wind Energy*, 24(4): 345-356. <https://doi.org/10.1002/we.2577>
- [5] Tchakoua, P., Wamkeue, R., Ouhrouche, M., Slaoui-Hasnaoui, F., Tameghe, T.A., Ekemb, G. (2014). Wind turbine condition monitoring: State-of-the-art review, new trends, and future challenges. *Energies*, 7(4): 2595-2630. <https://doi.org/10.3390/en7042595>
- [6] International Energy Agency (IEA). (2023). World energy outlook 2023. Paris, France: OECD/IEA. <https://www.iea.org/reports/world-energy-outlook-2023>.
- [7] International Renewable Energy Agency (IRENA). (2023). Renewable power generation costs in 2022. Abu Dhabi: IRENA. <https://www.irena.org/Publications/2023/Aug/Renewable-Power-Generation-Costs-in-2022>.
- [8] WindEurope. (2022). Financing and Investment Trends: The European wind industry in 2022. Brussels, Belgium: WindEurope. <https://windeurope.org/intelligence-platform/product/financing-and-investment-trends-2022/>.
- [9] Lin, Z., Liu, X., Lotfian, S. (2021). Impacts of water depth increase on offshore floating wind turbine dynamics. *Ocean Engineering*, 224: 108697. <https://doi.org/10.1016/j.oceaneng.2021.108697>
- [10] Reder, M.D. (2016). Wind turbine reliability modeling and maintenance optimization. Ph.D. dissertation, University of Stuttgart. <https://doi.org/10.18419/opus-8852>
- [11] Al-Fuqaha, A., Guibene, W., Mohammadi, M., Aledhari, M., Ayyash, M. (2015). Internet of things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials*, 17(4): 2347-2376. <https://doi.org/10.1109/COMST.2015.2444095>
- [12] Krizhevsky, A., Sutskever, I., Hinton, G.E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6): 84-90. <https://doi.org/10.1145/3065386>
- [13] Ezziyyani, M., Cherrat, L., Jebari, H., Rekiek, S., Ahmed, N.A. (2024). CNN-Based plant disease detection: A pathway to sustainable agriculture. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, pp. 679-696. [https://doi.org/10.1007/978-3-031-91337-2\\_62](https://doi.org/10.1007/978-3-031-91337-2_62)
- [14] Jebari, H., Rekiek, S., Reklaoui, K. (2025). Advancing precision livestock farming: Integrating hybrid AI, IoT, cloud and edge computing for enhanced welfare and efficiency. *International Journal of Advanced Computer Science and Applications*, 16(7): 302-311. <https://doi.org/10.14569/IJACSA.2025.0160732>
- [15] Barszcz, T. (2019). *Vibration-Based Condition Monitoring of Wind Turbines*. Springer International Publishing.
- [16] Tavner, P.J., Xiang, J., Spinato, F. (2007). Reliability analysis for wind turbines. *Wind Energy: An International Journal for Progress and Applications in*

- Wind Power Conversion Technology, 10(1): 1-18. <https://doi.org/10.1002/we.204>
- [17] Jardine, A.K.S., Lin, D., Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7): 1483-1510. <https://doi.org/10.1016/j.ymssp.2005.09.012>
- [18] Sikorska, J.Z., Hodkiewicz, M., Ma, L. (2011). Prognostic modelling options for remaining useful life estimation by industry. *Mechanical Systems and Signal Processing*, 25(5): 1803-1836. <https://doi.org/10.1016/j.ymssp.2010.11.018>
- [19] Bartkowiak, A., Zimroz, R. (2011). Outliers analysis and one class classification approach for planetary gearbox diagnosis. *Journal of Physics: Conference Series*, 305: (1): 012031. <https://doi.org/10.1088/1742-6596/305/1/012031>
- [20] Randall, R.B. (2011). *Vibration-Based Condition Monitoring: Industrial, Aerospace and Automotive Applications*. Wiley.
- [21] Khadivi, T., Savory, E. (2013). Experimental and numerical study of flow structures associated with low aspect ratio elliptical cavities. *Journal of Fluids Engineering*, 135(4): 041104. <https://doi.org/10.1115/1.4023652>
- [22] Wang, C.Y., Bochkovski, A., Liao, H.Y.M. (2023). 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 (CVPR)*, pp. 7464-7475.
- [23] Ren, S., He, K., Girshick, R., Sun, J. (2016). Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6): 1137-1149. <https://doi.org/10.1109/TPAMI.2016.2577031>
- [24] Ronneberger, O., Fischer, P., Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234-241. [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
- [25] Chen, L.C., Papandreou, G., Schroff, F., Adam, H. (2017). Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*. <https://doi.org/10.48550/arXiv.1706.05587>
- [26] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems (NIPS)*, 30: 5998-6008.
- [27] Rekiek, S., Jebari, H., Ezziyyani, M., Cherrat, L. (2025). AI-driven pest control and disease detection in smart farming systems. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, pp. 801-810. [https://doi.org/10.1007/978-3-031-91337-2\\_71](https://doi.org/10.1007/978-3-031-91337-2_71)
- [28] Rekiek, S., Jebari, H., Reklouai, K. (2024). Prediction of booking trends and customer demand in the tourism and hospitality sector using AI-based models. *International Journal of Advanced Computer Science and Applications*, 15(10): 404-412. <https://doi.org/10.14569/IJACSA.2024.0151043>
- [29] Gouiza, N., Jebari, H., Reklouai, K. (2023). IoT in smart farming: A review. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, pp. 149-161. [https://doi.org/10.1007/978-3-031-54318-0\\_13](https://doi.org/10.1007/978-3-031-54318-0_13)
- [30] Gouiza, N., Jebari, H., Reklouai, K. (2025). IoT in agriculture: Use cases and challenges. In M. Ezziyyani, J. Kacprzyk, & V. E. Balas (Eds.), *Proceedings of the International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD'2024)*, pp. 491-505. [https://doi.org/10.1007/978-3-031-91334-1\\_42](https://doi.org/10.1007/978-3-031-91334-1_42)
- [31] Jebari, H., Rekiek, S., Ezziyyani, M., Cherrat, L. (2024). Artificial intelligence for optimizing livestock management and enhancing animal welfare. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, pp. 790-800. [https://doi.org/10.1007/978-3-031-91337-2\\_70](https://doi.org/10.1007/978-3-031-91337-2_70)
- [32] Jebari, H., Mechkouri, M.H., Rekiek, S., Reklouai, K. (2023). Poultry-Edge-AI-IoT system for real-time monitoring and predicting by using artificial intelligence. *International Journal of Interactive Mobile Technologies*, 17(12): 149-170. <https://doi.org/10.3991/ijim.v17i12.38095>
- [33] Gouiza, N., Jebari, H., Reklouai, K. (2024). Integration for IoT-enabled technologies and artificial intelligence in diverse domains: Recent advancements and future trends. *Journal of Theoretical and Applied Information Technology*, 102(5): 1975-2029. <https://www.jatit.org/volumes/Vol102No5/25Vol102No5.pdf>
- [34] Ezziyyani, M., Cherrat, L., Rekiek, S., Jebari, H. (2024). Image classification of Moroccan cultural trademarks. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, pp. 767-779. [https://doi.org/10.1007/978-3-031-91337-2\\_68](https://doi.org/10.1007/978-3-031-91337-2_68)
- [35] Li, Z., Hoiem, D. (2018). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12): 2935-2947. <https://doi.org/10.1109/TPAMI.2017.2773081>
- [36] Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1-2): 1-210.
- [37] Eljyidi, A., Jebari, H., Rekiek, S., Reklouai, K. (2025). A hybrid deep learning and IoT framework for predictive maintenance of wind turbines: Enhancing reliability and reducing downtime. *International Journal of Advanced Computer Science & Applications*, 16(10): 203-211. <https://dx.doi.org/10.14569/IJACSA.2025.0161021>