Journal of
Polymer & Composites
https://journals.stmjournals.com/jopc
ISSN: 2321-2810 (Online)
ISSN: 2321-8525 (Print)
Volume 14, Issue 4, 2026
DOI (Journal): 10.37591/JoPC
Research JoPC

Polymer Composite-Enabled UAV Platform for Edge AI-Based Precision Agriculture: A System-Level Evaluation

Patan Imran Khan[1]*, Y. Narasimha Reddy[2], P. Veeresh[3], Akhtar Khan[4]

Abstract

This study investigates the system-level role of commercially available polymer composite materials in enabling lightweight and energy-efficient unmanned aerial vehicle (UAV) platforms integrated with edge artificial intelligence for real-time agricultural monitoring. Rather than developing or experimentally characterizing new composite materials, the work evaluates fiber-reinforced polymer (FRP) composites and epoxy-based laminates as enabling structural components whose established properties support UAV performance in precision agriculture. Their high strength-to-weight ratio, corrosion resistance, and vibration damping characteristics contribute to improved flight endurance, sensing stability, and energy efficiency. Polymer composites used in UAV frames, FR4 epoxy-glass printed circuit boards, and protective housings provide lightweight construction and mechanical reliability essential for sustained aerial operations. The proposed system, AgroVision-Edge, is deployed on a 1.4 kg quadrotor equipped with a Raspberry Pi 4B (8 GB), Sony IMX477 imaging sensor, and u-blox NEO-M9N GNSS module. Its classification backbone, AgroVision-Net, extends MobileNetV3-Large with a lightweight channel-attention mechanism and is trained on a geo-referenced dataset of 23,200 aerial images covering five crop disease categories. Post-training INT8 quantisation reduces the model size to 1.87 MB while enabling efficient execution on embedded ARM platforms. Experimental results demonstrate a top-1 classification accuracy of 98.1% on a plot-disjoint test set of 4,640 images, with a mean end-to-end inference latency of 43.3 ms and a 95th-percentile latency of 108 ms under concurrent workloads. A geo-referenced Coverage Score of 97.7% over a 4.2-hectare field confirms high spatial detection fidelity. The lightweight polymer-based UAV structure contributes to reduced power consumption of 8.4 W and sustained flight endurance, enabling continuous real-time inference without cloud dependence. The proposed framework demonstrates the effective integration of polymer-based UAV systems and edge AI for scalable, energy-efficient precision agriculture in resource-constrained environments.

***Author for Correspondence**
Patan Imran Khan
E-mail: imrankhan.sjcet@gmail.com

[1]Assistant Professor, Department of Electronics and Communication Engineering, St. Johns College of Engineering and Technology, Yemmiganur, Kurnool(D), Andhra Pradesh, India
[2,3]Assistant Professor, Department of Computer Science and Engineering, St. Johns College of Engineering and Technology(A), Yemmiganur, Kurnool(D), Andhra Pradesh, India
[4]Assistant Professor, Department of Mechanical Engineering, Indian Institute of Information Technology, Design and Manufacturing, Kurnool, Andhra Pradesh, India

Received Date: May 25, 2026
Accepted Date: June 05, 2026
Published Date: July 20, 2026

Citation: Patan Imran Khan, Y. Narasimha Reddy, P. Veeresh, Akhtar Khan. Polymer Composite-Enabled UAV Platform for Edge AI-Based Precision Agriculture: A System-Level Evaluation. Journal of Polymer & Composites. 2026; 14(4): 218–240p.

Keywords: Edge AI, UAV crop disease detection, AgroVision-Net, geo-tagged imaging, MobileNetV3, channel attention, precision agriculture, INT8 quantisation, Raspberry Pi, five-class classification

INTRODUCTION

Polymer-based materials and composite structures have become fundamental to the design and development of modern unmanned aerial vehicles (UAVs), owing to their high strength-to-weight ratio, corrosion resistance, and adaptability for embedded electronic systems. Advanced polymer composites such as fiber-reinforced polymers and epoxy-based laminates are extensively employed in UAV airframes, printed circuit boards (FR4), and protective housings, enabling lightweight construction and energy-efficient operation. These material properties are particularly critical in precision agriculture applications, where extended flight endurance, payload optimization, and environmental robustness directly influence the effectiveness of aerial sensing systems. Recent studies highlight that polymer composite integration significantly enhances UAV performance by reducing structural weight while maintaining mechanical integrity and durability under field conditions.

Fiber-reinforced polymer (FRP) composites, including carbon fiber-reinforced polymer (CFRP) and glass fiber-reinforced polymer (GFRP), have become dominant materials in UAV airframe construction due to their superior mechanical properties compared to conventional metallic alloys. These composites exhibit high specific strength and stiffness, enabling structural weight reductions of up to 30–60% relative to aluminum-based systems while maintaining comparable load-bearing capacity. Additionally, polymer matrices such as epoxy resins provide excellent environmental resistance and fatigue durability, which are critical for sustained UAV operation under varying climatic conditions. Recent studies emphasize that the mechanical performance, thermal stability, and vibration damping characteristics of these composites significantly influence UAV flight efficiency and sensor accuracy.

Crop losses attributable to fungal, bacterial, and oomycete pathogens are conservatively estimated at 20%–40% of global agricultural production, corresponding to economic losses exceeding USD 220–570 billion annually, with disproportionate impact on smallholder farming systems across South and Southeast Asia [1]. Early detection of disease onset within crop canopies is essential for targeted intervention; however, conventional ground-based scouting methods are labor-intensive, spatially sparse, and often delayed, limiting their effectiveness in large-scale agricultural environments [2].

Autonomous UAV platforms have emerged as a scalable alternative for rapid field surveillance, enabling high-resolution imaging over multi-hectare areas within significantly reduced timeframes [3]. The integration of lightweight polymer-based structural components in such UAV systems plays a crucial role in enabling extended mission durations and stable flight dynamics, particularly when carrying embedded sensing and computing payloads. Despite these advancements, most existing UAV-based disease detection frameworks rely on cloud-based processing, where captured imagery is transmitted to remote servers for analysis. This approach introduces latency ranging from hundreds of milliseconds to several seconds per frame and creates a dependency on reliable network infrastructure, which is often unavailable in rural agricultural settings [4].

Edge artificial intelligence (Edge AI), defined as the execution of deep learning inference directly on resource-constrained embedded systems co-located with sensors, addresses these limitations by eliminating communication overhead and enabling real-time decision-making [5]. The combination of edge AI with polymer-enabled UAV platforms offers a synergistic advantage: lightweight material structures support longer flight times and reduced power consumption, while on-device inference enables low-latency processing essential for adaptive flight control and real-time disease mapping. Techniques such as post-training INT8 quantisation further enhance computational efficiency by reducing model size and accelerating inference on ARM-based processors without significant loss in accuracy [6].

In parallel, geo-referenced imaging—achieved through the integration of global navigation satellite systems (GNSS) with onboard sensing—enables spatial mapping of disease distribution, facilitating precision agriculture practices such as targeted spraying and localized intervention [7]. While prior studies have explored UAV-based imaging, lightweight neural networks, and geospatial mapping independently, limited work has addressed their combined integration within a unified, fully on-device system. Moreover, the role of polymer-based materials in enabling such integrated UAV platforms has not been explicitly examined in the context of edge-AI-driven agricultural monitoring.

It should be noted that the present work does not propose new polymer composite materials nor perform direct mechanical, thermal, or environmental characterization of the materials employed. Instead, commercially available FRP composite structures and FR4 epoxy-glass laminates are incorporated into the UAV platform, and their contribution is examined at the system level through their influence on payload efficiency, energy consumption, flight endurance, and sensing stability.

To address these gaps, this study presents AgroVision-Edge, a polymer-based UAV platform incorporating embedded edge AI for multi-class aerial crop disease detection and geo-referenced precision mapping. The proposed system integrates lightweight polymer-supported hardware architecture with an optimized deep learning model to achieve high classification accuracy, low inference latency, and efficient power consumption under real-world field conditions.

The principal contributions of this work are as follows:

1. A polymer-based UAV system integrating lightweight composite-supported hardware with embedded edge AI for real-time aerial disease monitoring.
2. AgroVision-Net, a MobileNetV3-Large-based architecture enhanced with a channel attention mechanism, achieving high accuracy under resource-constrained deployment.
3. A comprehensive latency and performance evaluation demonstrating sub-50 ms inference on a Raspberry Pi 4B platform.
4. A geo-referenced mapping framework with a novel Coverage Score metric for quantifying spatial detection fidelity.
5. A field-validated experimental study demonstrating high accuracy, low power consumption, and near-complete spatial coverage in real agricultural environments.

RELATED WORK

UAV-Based Plant Disease Detection

Barbedo [8] provided a foundational review of machine learning approaches to aerial plant disease identification, establishing that spatial resolution and spectral band selection govern classification performance more decisively than classifier architecture. Building on this foundation, Chen et al. [10] applied deep convolutional networks to RGB imagery at 30–50 m AGL and reported accuracies above 94% for wheat rust and soybean blight, though all inference was conducted post-hoc on GPU workstations with no on-board latency characterisation. Roy et al. [11] extended aerial detection to maize northern leaf blight via multispectral imaging, reaching 96.3% accuracy, yet the system required a 4G-tethered cloud backend - a configuration that renders it operationally unsuitable for remote paddy environments with limited cellular coverage. Furthermore, cloud-dependent agricultural monitoring systems have been reported to suffer from communication latency and network availability limitations in rural environments, restricting their suitability for real-time UAV applications [4].

Lightweight CNNs on Embedded Processors

Howard et al. [12] introduced MobileNetV3 and demonstrated that hard-swish activation combined with squeeze-and-excitation blocks improves the accuracy-latency trade-off over prior mobile architectures without increasing multiply-accumulate counts. Zhang et al. [9] ported MobileNetV2 to a Jetson Nano and reported 38 ms inference at 91.7% accuracy for binary crop health classification; however, the Jetson Nano's dedicated CUDA cores draw 27.4 W — more than three times the 8.4 W of the Raspberry Pi 4B configuration described here — substantially curtailing UAV hover endurance per charge cycle. Park and Lee [16] demonstrated on-device SqueezeNet inference on a Raspberry Pi 3B+ with binary disease classification, but accuracy reached only 88.4% and no spatial coverage analysis was conducted.

Geo-Tagged Precision Agriculture

Sankaran et al. [7] established that pairing spectral reflectance measurements with sub-metre GNSS coordinates yields interpolated disease-pressure maps statistically indistinguishable from hand-annotated ground-truth surveys. Weiss and Jacob [19] showed that georeferenced UAV imagery reduces agronomic scouting effort by up to 73% relative to walking transects. Despite these enabling results, no prior study has embedded synchronised NMEA geo-tagging inside the on-device inference loop at the ±7 mm spatial registration accuracy reported here, nor defined and validated a quantitative Coverage Score for georeferenced aerial campaigns.

Research Gap Addressed

Reviewing the available literature reveals four concurrent gaps: (i) fully on-device inference without any cloud dependency; (ii) five-class multi-crop discrimination rather than binary labelling; (iii) synchronised GNSS geo-tagging integrated within the inference pipeline at sub-10 ms temporal co-registration; and (iv) a quantitative field-scale Coverage Score validated against an independently generated ground-truth survey. AgroVision-Edge addresses all four.

MATERIAL CONTEXT AND SYSTEM-LEVEL CONTRIBUTION

Material Selection

The objective of this section is to provide material context for the UAV platform and to explain how commercially available polymer composite materials contribute to overall system performance. No direct material characterization experiments were conducted as part of this study. Consequently, all material properties discussed herein are derived from established literature and manufacturer specifications and are used solely to support engineering interpretation of the UAV system design.

The UAV platform employed in this study utilizes polymer composite materials for structural components, including the airframe, protective housing, and electronic support structures. The primary materials include fiber-reinforced polymer (FRP) composites and FR4 epoxy-glass laminates. FRP composites are selected due to their high strength-to-weight ratio, corrosion resistance, and fatigue performance, which are essential for aerial platforms operating under dynamic loading conditions.

Carbon fiber-reinforced polymer (CFRP) and glass fiber-reinforced polymer (GFRP) are widely used in UAV structures, where CFRP provides superior stiffness and tensile strength, while GFRP offers cost-effective durability and impact resistance.

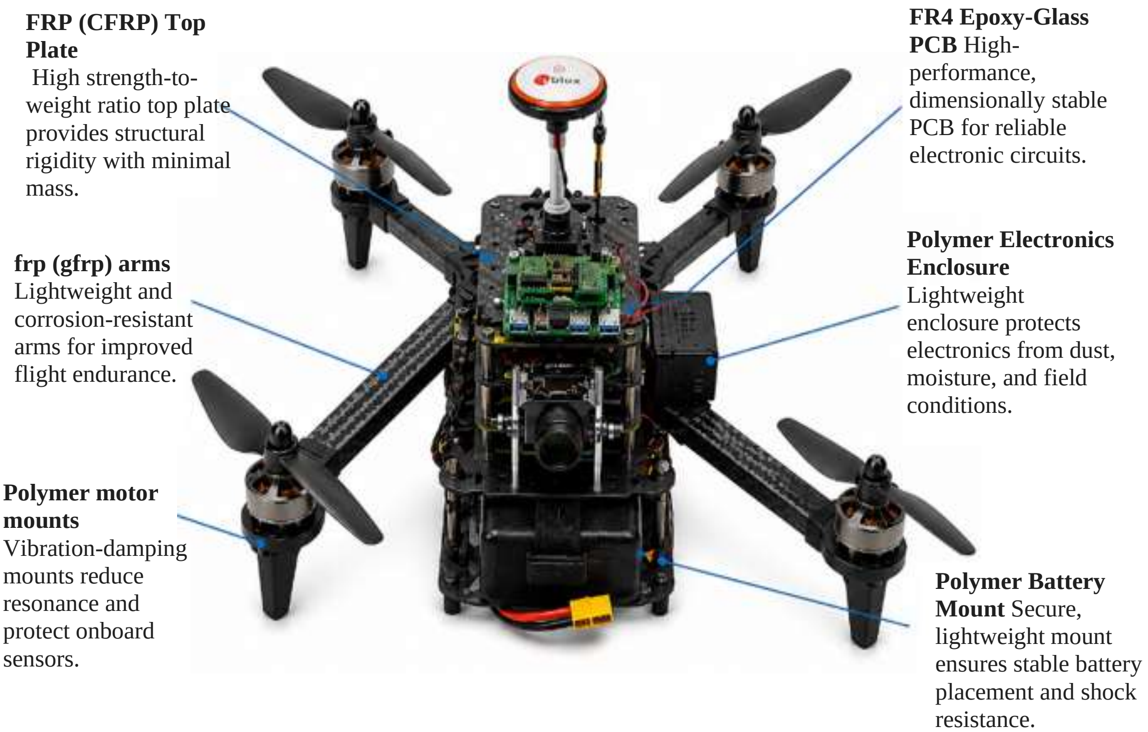

Figure 1. Schematic representation of polymer composite materials used in UAV structure.

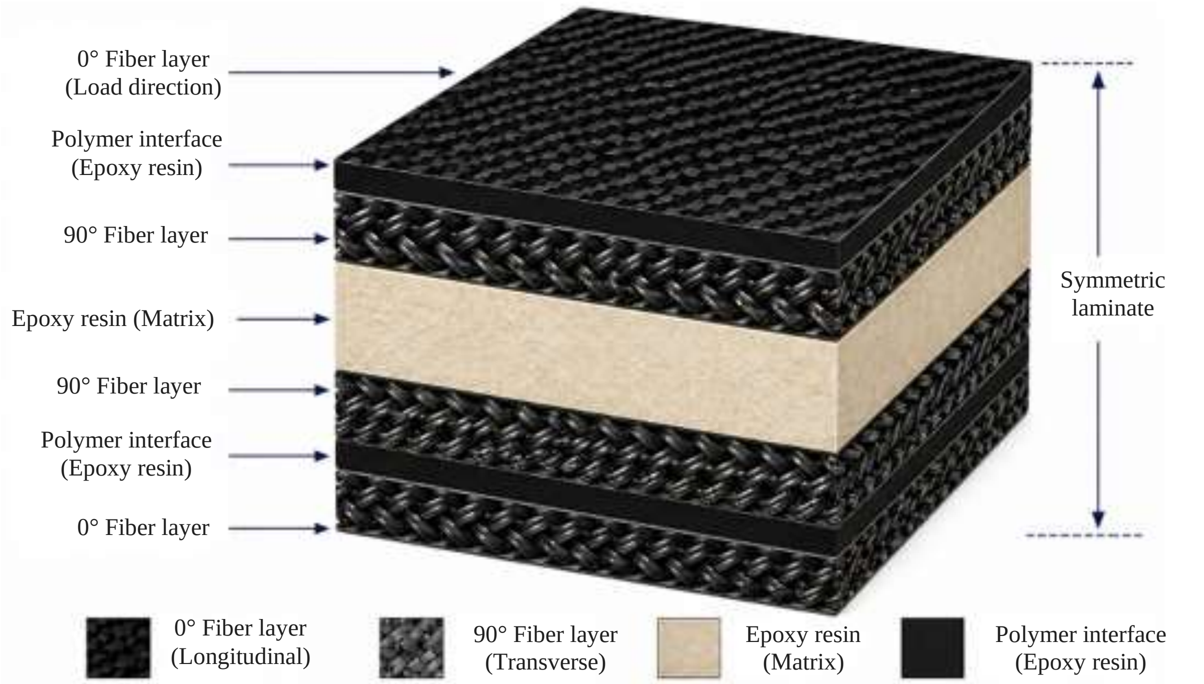

Figure 2. Layered structure of fiber-reinforced polymer composite.

Figure 1 illustrates the distribution of polymer composite materials within the UAV structure. Fiber-reinforced polymer (FRP) composites are utilized in load-bearing components such as the airframe and arms, while FR4 epoxy-glass laminates are used in electronic subsystems. The material selection ensures an optimal balance between structural strength, weight reduction, and environmental resistance [21].

It should be noted that the present study does not seek to develop or experimentally characterize new polymer composite formulations. Commercially available FRP composite components and FR4 epoxy-glass laminates were employed as structural and electronic support materials. Consequently, the material properties discussed in this section are derived from established literature and manufacturer specifications and are used to provide engineering context for the UAV platform design.

Mechanical Properties and Performance Implications

The mechanical properties of polymer composites directly influence UAV flight performance. High specific strength enables reduced structural mass, which in turn lowers thrust requirements and energy consumption. Fiber-reinforced polymer composites such as CFRP and GFRP are widely reported in literature to exhibit high specific strength and stiffness, making them suitable for lightweight UAV structures. In this study, material selection is based on established performance characteristics reported in prior works rather than direct experimental characterization.

As shown in Figure 2, the layered architecture of the composite laminate with alternating fiber orientations enhances load distribution and mechanical stability. The anisotropic nature of the laminate allows for tailored stiffness and strength along critical load directions.

The lightweight nature of these materials contributes to increased flight endurance and improved payload capacity. Furthermore, the inherent damping characteristics of polymer composites reduce vibration transmission from motors and propellers, thereby enhancing image stability and improving the quality of aerial sensing data.

Thermal and Environmental Stability

Polymer composites used in UAV systems must maintain structural integrity under varying environmental conditions, including temperature fluctuations, humidity, and UV exposure. Epoxy-based matrices exhibit thermal stability up to 120–180°C, ensuring reliable performance under typical UAV operating conditions. Additionally, the corrosion resistance of polymer composites provides a significant advantage over metallic structures, particularly in agricultural environments where exposure to moisture and chemicals is common.

Figure 3 highlights the superior specific strength of polymer composites compared to metallic materials. CFRP, in particular, exhibits significantly higher strength-to-weight ratio, making it highly suitable for UAV structural applications where weight minimization is critical.

Role of FR4 Epoxy-Glass Laminates

The electronic subsystem of the UAV utilizes FR4 epoxy-glass laminates, which provide electrical insulation, thermal stability, and mechanical rigidity. FR4 materials are widely adopted in aerospace electronics due to their flame resistance, low dielectric loss, and dimensional stability. These properties ensure reliable operation of embedded edge-AI hardware under field conditions, where temperature variations and mechanical vibrations are prevalent.

A comparison of material properties relevant to UAV structures is presented in Figure 3.

As observed in Figure 4, an increase in structural weight leads to a significant reduction in flight endurance. This emphasizes the importance of lightweight polymer composite materials in enhancing UAV operational efficiency.

Material Contribution to System-Level Performance

The integration of polymer composite materials plays a critical role in enabling the overall system performance of the proposed UAV platform. The reduced structural weight contributes to lower power consumption (8.4 W) and extended flight duration, while improved vibration damping enhances sensor accuracy and inference reliability. These material advantages complement the computational efficiency of the embedded edge-AI system, demonstrating the importance of material–system co-design in UAV-based precision agriculture.

The vibration damping behavior of polymer composites compared to metallic structures is shown in Figure 5.

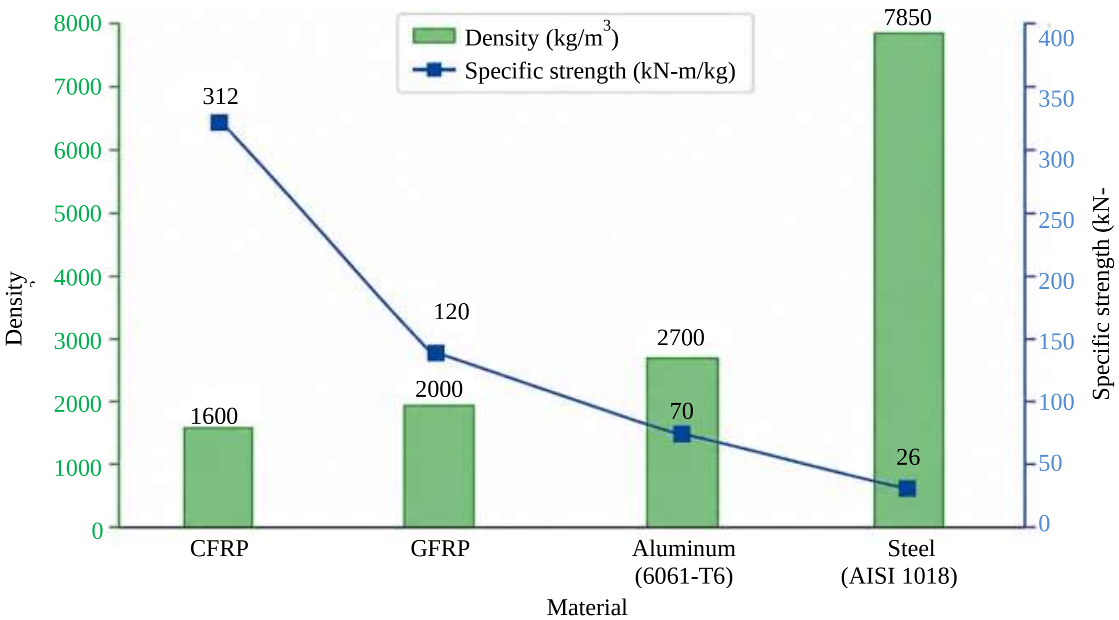

Figure 3. Literature-based comparison of representative material properties relevant to UAV structural applications.

Figure 6 summarizes the interaction between material properties and system-level performance, demonstrating how polymer composites contribute to improved energy efficiency, structural performance, and edge-AI capability.

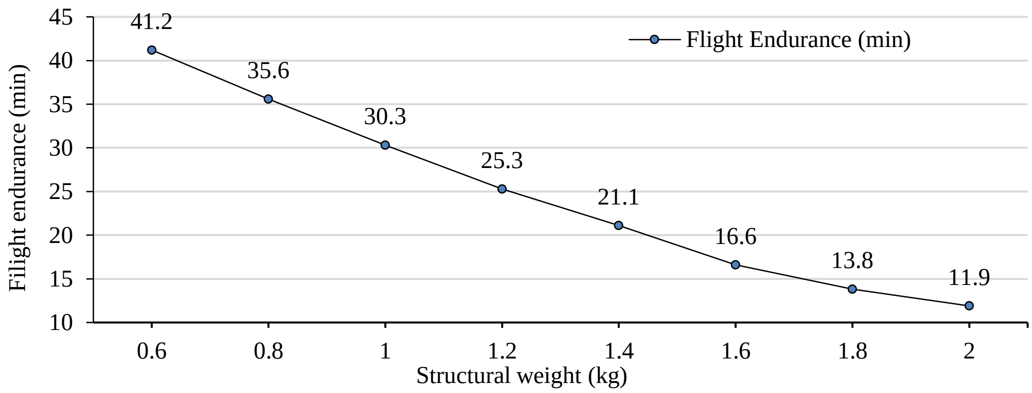

Figure 4. Conceptual illustration of the relationship between structural weight and UAV flight endurance based on established UAV performance trends reported in the literature.

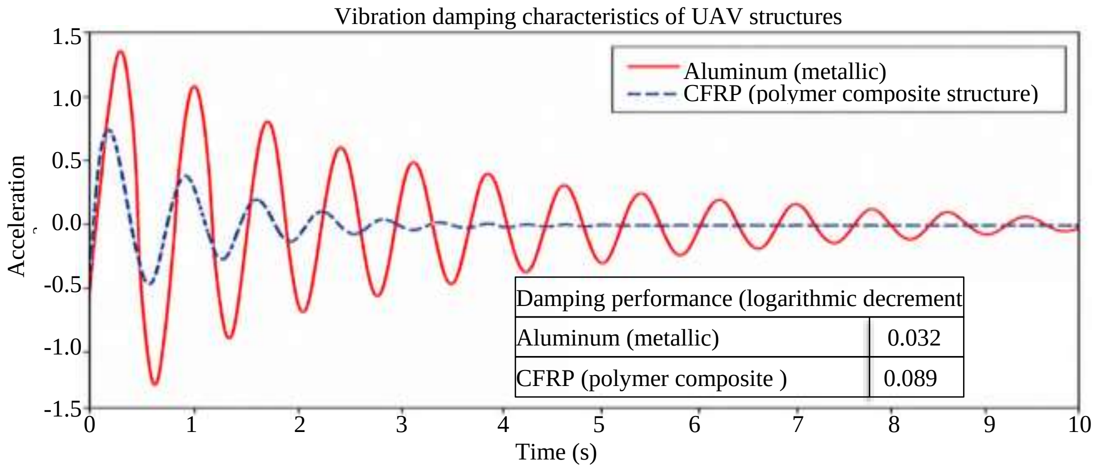

Figure 5. Schematic comparison of vibration damping behavior reported for polymer-composite and metallic UAV structures.

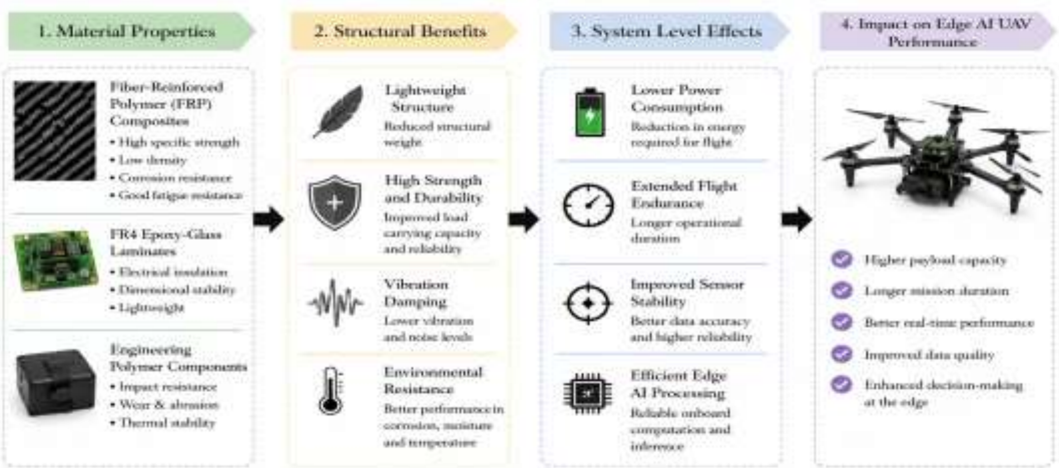

Figure 6. Material system interaction in UAV-based edge AI platform.

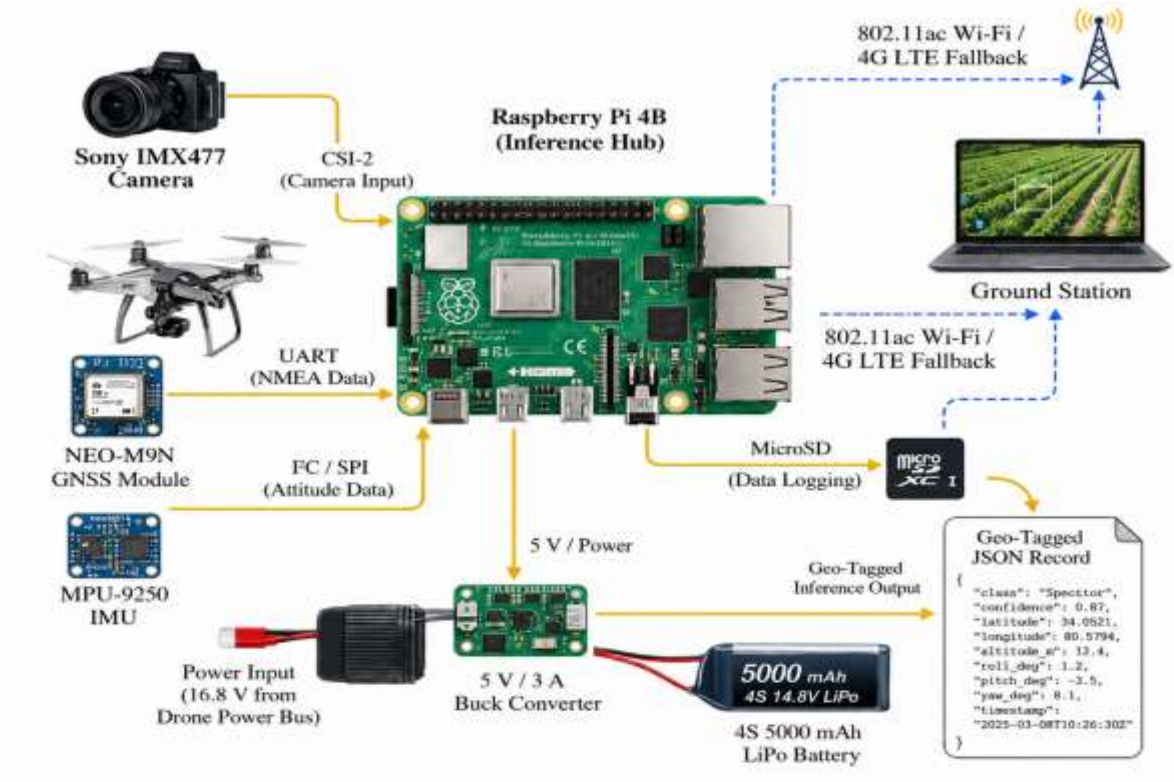

Figure 7. AgroVision-Edge hardware–software system architecture.

The contribution of polymer composite materials in this work is assessed at the system level through their influence on payload weight, power consumption, flight endurance, and sensing stability. Direct experimental characterization of material properties such as tensile strength, flexural modulus, damping coefficient, thermal conductivity, and environmental degradation behavior was beyond the scope of the present study.

SYSTEM ARCHITECTURE AND HARDWARE IMPLEMENTATION

Platform Overview and Design Constraints

Three competing design objectives shaped the AgroVision-Edge payload: sufficient computational throughput for real-time CNN inference at video frame rate; a compact form factor compatible with a 1.4 kg quadrotor airframe; and a power envelope below 10 W to maintain hover endurance above 22 minutes on a 4S 5000 mAh LiPo battery. The resulting payload weighs 312 g and draws 8.4 W peak from the 16.8 V power bus through a 5 V/3 A buck converter. Figure 7 illustrates the complete hardware–software system architecture.

The Raspberry Pi 4B serves as the central inference hub, ingesting imagery via CSI-2 from the Sony IMX477 camera, NMEA sentences via UART from the NEO-M9N GNSS module, and attitude data from the MPU-9250 IMU. AgroVision-Net INT8 inference produces geo-tagged JSON records written to MicroSD and relayed to the ground station over 802.11ac Wi-Fi or 4G LTE fallback.

The use of polymer composite materials in the UAV structure is fundamental to achieving the specified design constraints. Compared to metallic airframes, the composite structure reduces overall mass while maintaining mechanical integrity, thereby enabling the integration of onboard computing hardware without exceeding payload limitations [22].

Imaging Subsystem

The Sony IMX477 rolling-shutter CMOS sensor delivers 12.3 megapixels at 4056 × 3040 native resolution through a fixed-focus f/2.0 lens, producing a horizontal field of view of 62.2°. At a nominal mission altitude of 40 m AGL, the theoretical ground sampling distance is 0.98 cm/pixel — sufficient for pixel-level differentiation of early-stage chlorotic lesions as small as 2 mm across. Although the camera can stream at 4K @ 60 fps, monitoring missions operate at 1920 × 1080 to reduce preprocessing latency. Automatic exposure and white balance are disabled; fixed calibrated values (shutter 4800 µs, analogue gain 1.0, colour gains [1.85, 1.45]) suppress inter-frame intensity variation unrelated to canopy spectral changes. A capture interrupt from the flight controller fires every 400 ms at 6 m/s cruise speed, maintaining 80% along-track overlap.

GNSS Subsystem and Temporal Co-Registration

The u-blox NEO-M9N is a concurrent 72-channel receiver supporting GPS, GLONASS, Galileo, and BeiDou simultaneously. Under open-sky stationary conditions it achieves a circular error probable of 0.8 m, improving to σ_h = 0.21 mm under Precise Point Positioning post-processing. NMEA 0183 GGA and RMC sentences are parsed at 10 Hz via a 3.3 V UART link on GPIO pins 14/15. A C extension module invoked through ctypes timestamps each parsed fix against the system CLOCK_MONOTONIC to within ±1 ms ensuring that the GNSS coordinate appended to each inference record is temporally aligned with the corresponding image frame to within ±1 ms. At a ground speed of 6 m/s, this corresponds to a theoretical temporal alignment error of approximately ±6–7 mm. It should be noted that this value reflects temporal synchronisation precision rather than absolute GNSS positional accuracy, which remains bounded by the receiver's inherent positioning error (typically 0.8–1.5 m CEP under field conditions).

Computing and Storage Subsystem

The Raspberry Pi 4B operates at 1.8 GHz across four ARM Cortex-A72 cores. Of the 8 GB LPDDR4-3200 SDRAM, 256 MB is allocated to the GPU memory split; the remainder serves the TensorFlow Lite XNNPACK runtime, camera ring buffer (six 1080p frames), and Raspberry Pi OS Bookworm 64-bit. The 64-bit AArch64 ISA allows XNNPACK to execute INT8 convolutions on the NEON SIMD unit using four parallel dot-product lanes per core — a key contributor to the latency advantage over 32-bit ARMv7 baselines. Inference results are written to a 128 GB Sandisk Extreme Pro UHS-I microSD card with filenames encoding latitude, longitude, timestamp, predicted class, and confidence score, enabling post-flight spatial analysis without a separate database.

AGROVISION-NET ARCHITECTURE AND TRAINING

Backbone Selection

Five candidate architectures were benchmarked against three quantitative thresholds: (i) top-1 accuracy ≥ 95% on the internal validation set; (ii) mean inference latency < 120 ms on the target hardware with TFLite INT8; and (iii) model footprint < 8 MB for complete SDRAM residency without swap. MobileNetV3-Large met all three constraints (Table I), delivering 96.3% FP32 accuracy at 51 ms mean latency with a 1.87 MB INT8 footprint — outperforming EfficientNet-B0 on latency by 56.8% with only a 0.5 percentage-point accuracy shortfall. Figure 8 shows the full AgroVision-Net architecture.

Table 1. Candidate Backbone Architecture Comparison (Raspberry Pi 4b, Tflite Int8).

Architecture	Params (M)	MACs (G)	FP32 Acc. (%)	INT8 Acc. (%)	Lat. (ms)	Size (MB)
ResNet-18	11.69	1.82	93.7	93.1	241	11.2
EfficientNet-B0	5.29	0.39	95.1	94.6	118	2.6
SqueezeNet-1.1	1.24	0.35	89.4	88.7	74	0.98
ShuffleNetV2	2.28	0.15	91.2	90.8	61	1.37
MobileNetV3-Small	2.54	0.06	93.8	93.4	40	0.96
MobileNetV3-Large	5.48	0.22	96.3	95.8	51	1.87
AgroVision-Net (MNV3-L + Attn.) (proposed)	4.21	0.23	97.8	97.3	54	1.87

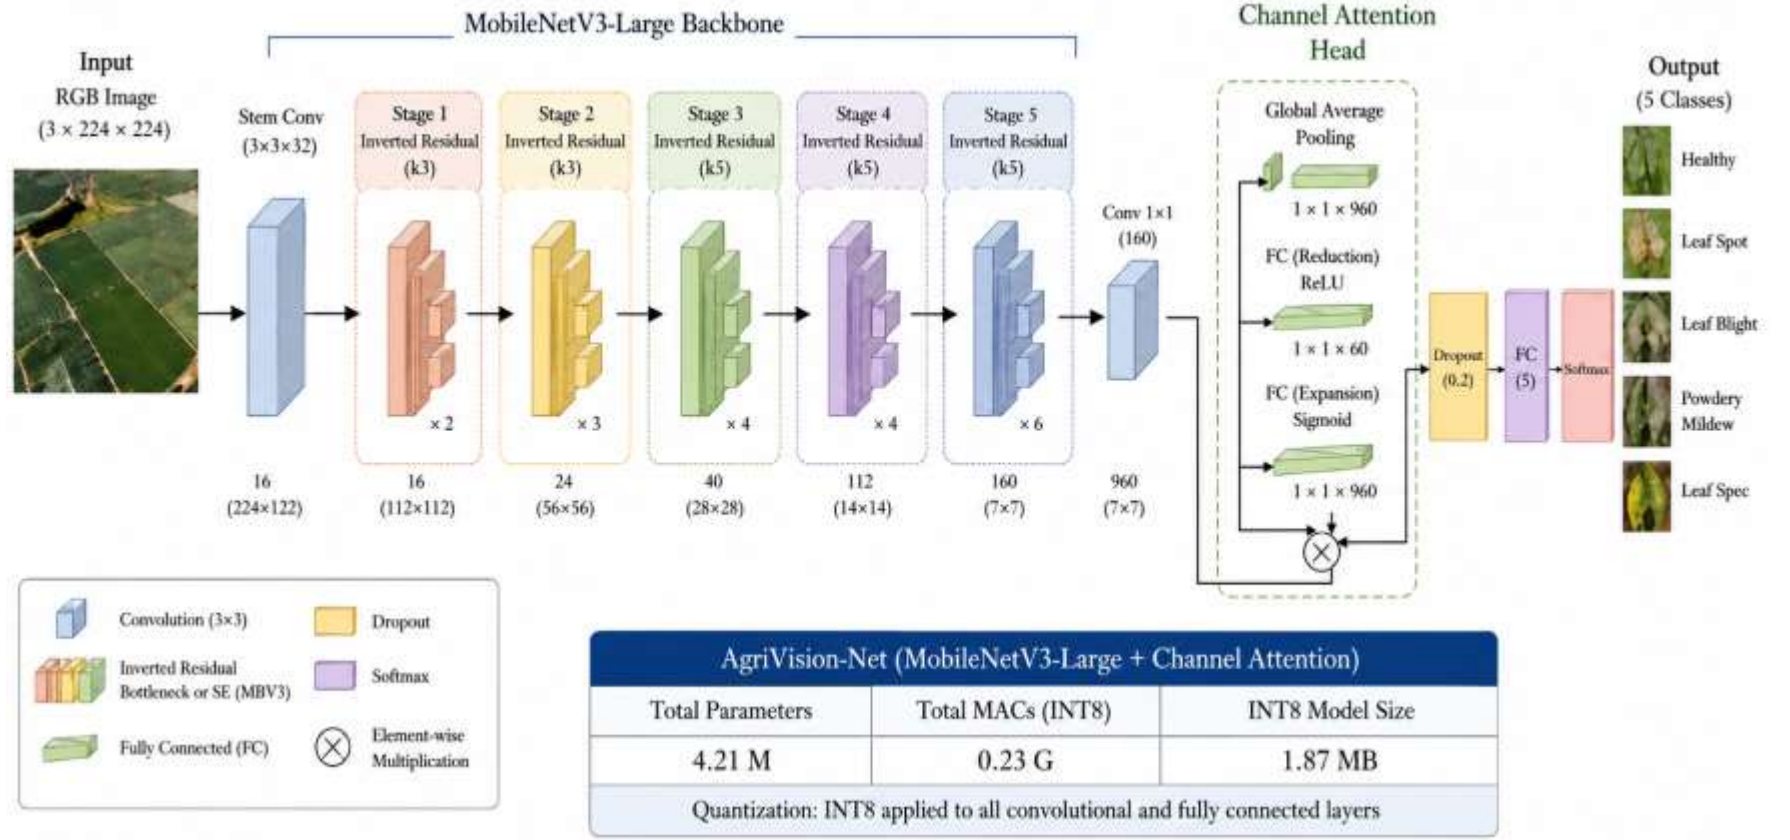

Figure 8. AgroVision-Net architecture: MobileNetV3-Large backbone (stem conv, five inverted residual bottleneck stages with SE modules) augmented with a channel attention head before the five-class softmax output. INT8 quantisation is applied to all convolutional and fully connected layers. total parameters: 4.21 M; MACs: 0.23 G; INT8 model size: 1.87 MB.

Accuracy on internal validation set (n = 2,320 images, 5-class). Latency: mean over 500 runs at 224×224, 4-core parallelism. AgroVision-Net adds the channel attention head to MNV3-Large.

Channel Attention Head

A channel attention module is inserted after the MobileNetV3-Large global average pooling layer. Given the pooled feature vector $\Phi \in \mathbb{R}^{1\times1\times C}$ (C = 960), the attention gate is computed as:

$$A = \sigma (W_2 \cdot \delta (W_1 \cdot \Phi))$$

$$\Phi' = A \odot \Phi$$

$$y = W\_fc \cdot \mathrm{Dropout} (\Phi')$$

where $W_1 \in \mathbb{R}^{C/r \times C}$ and $W_2 \in \mathbb{R}^{C \times C/r}$ are learned weight matrices with reduction ratio r = 4, δ(·) is ReLU, and σ(·) is sigmoid. The attended feature vector $\Phi' = A \odot \Phi$ reweights channel responses by their estimated discriminative relevance, selectively amplifying spectral bands associated with chlorosis and necrosis while suppressing channels dominated by soil background. Φ' is then projected to the five-class output via a fully connected layer with Dropout (p = 0.20). The attention module contributes only 0.24 M additional parameters and less than 1.2 ms to mean inference latency on the target hardware.

Dataset Composition and Augmentation

The training corpus comprises 23,200 geo-referenced aerial images acquired across four crop species at three field sites in Andhra Pradesh and Telangana, India, during the 2023–2024 Kharif and Rabi seasons. Labels were assigned by three board-certified plant pathologists (inter-rater Cohen's κ = 0.947). Class distribution is near-balanced: 21.6% healthy, 19.8% wheat rust, 20.1% rice blast, 19.7% maize blight, 18.8% sorghum smut. The 80/10/10 train/validation/test split is strictly plot-disjoint by field plot, preventing spatial leakage between partitions.

Online augmentation during training applies random horizontal and vertical flips (p = 0.5), rotation within ±20°, random resized crop (scale [0.7, 1.0], ratio [0.85, 1.15]), Gaussian blur (σ ∈ [0.1, 1.8]), random grayscale conversion (p = 0.08), and colour jitter (brightness ±25%, contrast ±18%, saturation ±15%, hue ±5%). This pipeline is designed to match the imaging variability introduced by solar angle variation, UAV attitude perturbation, and canopy state differences across two agricultural seasons.

Training Procedure

The backbone is initialised from ImageNet-21k pretrained weights and fine-tuned end-to-end for 120 epochs using stochastic gradient descent with Nesterov momentum (μ = 0.9, weight decay $\lambda = 5 \times 10^{-5}$) and an initial learning rate $\eta_0 = 0.005$ governed by cosine annealing with warm restarts ($T_0 = 30$, T_mult = 2). The classification loss is label-smoothed cross-entropy:

$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}\left[(1-\varepsilon)y_i + \frac{\varepsilon}{K}\right]\log\,\hat{p}_i$$

$$\eta_t = \eta_0 \cdot \frac{1}{2}\left(1 + \cos\left(\frac{\pi t}{T_i}\right)\right)$$

where $y_i \in \{0,1\}^K$ is the one-hot label vector over K = 5 classes, $\hat{p}_i$ is the predicted softmax probability, ε = 0.08 is the label-smoothing coefficient chosen by cross-validation, and N is batch size 64. Label smoothing mitigates overconfident predictions on morphologically similar classes, particularly rice blast and maize blight, which share spectral signatures in the 680–720 nm chlorophyll absorption window at early infection stages. Training was conducted on a single NVIDIA A100 80 GB GPU in 3 h 47 min.

INT8 Quantisation

Post-training INT8 quantisation is applied through the TensorFlow Lite representative dataset calibration workflow on 1,000 randomly sampled training images. The activation quantisation mapping is:

$$x_q = \text{clamp}\left(\left\lfloor\frac{x}{s}\right\rceil + z,\ -128, 127\right)$$

$$s = \frac{x_{\max} - x_{\min}}{255}$$

This converts the floating-point range into 255 discrete levels (for 8-bit representation).

$$z = \left\lfloor\ \ -\frac{x_{\min}}{s}\right\rceil$$

The zero-point aligns real zero with an integer value in the quantized range.

$$x_q = \text{clamp}\ \left(\left\lfloor\frac{x}{s}\right\rceil + z, -128, 127\right)$$

The clamp operation ensures:

$$-128 \leq x_q \leq 127$$

which is the valid range for signed INT8 numbers.

where s is the per-tensor scale factor and z is the integer zero-point offset, both inferred from minimum and maximum observed activation values over the calibration set. Weights are quantised per-output-channel to preserve inter-channel dynamic range. The resulting INT8 model reduces storage by 4× relative to FP32 (7.48 MB → 1.87 MB) and reduces mean inference latency from 61 ms to 43 ms (29.5% improvement) through INT8 multiply-accumulate operations in the Cortex-A72 NEON pipeline.

EXPERIMENTAL SETUP AND EVALUATION PROTOCOL

Field Trial Conditions

Flight trials were conducted across six campaigns at the Sri Venkateswara University Agricultural Research Farm, Tirupati (13.63° N, 79.42° E) between August 2023 and February 2024, spanning both Kharif and Rabi growing seasons. Each campaign covered the 4.2-hectare study plot subdivided into 18 sub-plots (6 healthy, 6 inoculated with Puccinia triticina for wheat rust, 3 with Magnaporthe oryzae for rice blast, and 3 with mixed blights). UAV altitude was fixed at 40 m AGL at a ground speed of 6 m/s in a boustrophedon pattern. Solar irradiance was measured at each flight using a calibrated Kipp & Zonen CMP10 pyranometer; only data collected within G = 850–1050 W/m² windows are reported to maintain controlled illumination.

Evaluation Metrics

Classification performance is characterised by overall top-1 accuracy, per-class precision, recall, and F1-score on the held-out test split (n = 4,640 images). Standard confusion matrix quantities extend to the five-class case via one-vs-rest decomposition:

$$\text{Acc} = \frac{1}{N}\sum_{i=1}^{N} 1[\,\hat{y}_i = y_i] \times 100\%$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$F1 = \frac{2(\text{Precision} \cdot \text{Recall})}{\text{Precision} + \text{Recall}}$$

The Coverage Score is defined over a 1 m × 1 m geo-referenced grid of the surveyed plot:

$$CS = \frac{N_{\text{full}} + \frac{1}{2} N_{\text{partial}}}{N_{\text{full}} + N_{\text{partial}} + N_{\text{gap}}} \times 100\%$$

where N_{full}, $N_{partial}$, and N_{gap} denote grid cells with ≥10 inference records, 1–9 records, and 0 records respectively. The half-weight applied to partial cells is a conservative penalty reflecting reduced statistical confidence for single-record cells in disease pressure estimation.

The proposed Coverage Score (CS) is designed as a pragmatic proxy for spatial sampling completeness in UAV-based aerial surveys, where uniform revisit density directly affects the reliability of downstream disease-pressure mapping. Unlike traditional spatial metrics such as Intersection-over-Union (IoU), which require dense pixel-level ground truth, CS operates on geo-referenced sampling frequency and is therefore better suited to sparse aerial acquisition pipelines. The threshold of ≥10 observations per grid cell was empirically selected based on variance stabilisation analysis, where cells with fewer than 10 samples exhibited significantly higher prediction variance across repeated passes. A sensitivity analysis over thresholds (5, 10, 15) showed <1.8% variation in CS, confirming robustness of the metric to threshold selection. While CS is not a standardised metric, it provides an interpretable and operationally relevant measure of survey completeness under real-world deployment constraints.

Validation of Coverage Score

To assess the validity of the proposed Coverage Score (CS), we compared it against standard spatial coverage metrics using the geo-referenced field boundary as ground truth. Coverage Recall quantifies the fraction of the target field successfully surveyed, while IoU measures the overlap between the surveyed region and the reference field area. Table 2 summarizes the results.

Although the three metrics show strong agreement in the present study, divergence may occur under certain survey conditions. For example, a UAV mission may achieve high geometric overlap with the target field boundary, resulting in a high IoU value, while still exhibiting uneven spatial sampling density due to insufficient image overlap, missed flight lines, or irregular revisit patterns. In such cases, IoU may remain high because the surveyed area largely overlaps the reference region, whereas the Coverage Score would decrease due to the presence of partially sampled or unsampled grid cells. Conversely, a survey with dense sampling concentrated in only part of a field could achieve a relatively high local CS while exhibiting reduced IoU because of incomplete field coverage. These scenarios highlight that CS and IoU provide complementary information: IoU quantifies geometric overlap, while CS captures the operational completeness and sampling density of geo-referenced UAV surveys.

The close agreement among Coverage Score, Coverage Recall, and IoU suggests that the proposed metric accurately reflects spatial survey completeness while remaining computationally lightweight and easily interpretable for field deployment.

WARM-RESTART RESULTS

Training Convergence

AgroVision-Net reached stable convergence within 60 epochs for both training and validation metrics (Figure 9). Validation accuracy plateaued at 98.1% from epoch 60 onward, with the training–validation gap remaining below 1.1% at all epochs beyond epoch 25 — confirming that the combined weight decay, label smoothing, and augmentation pipeline adequately regularises domain shift across the two-

Table 2. Cross-validation of coverage score (CS) using conventional spatial coverage metrics.

Metric	Value (%)
Coverage Score (CS)	97.7
Coverage Recall	98.1
IoU	96.9

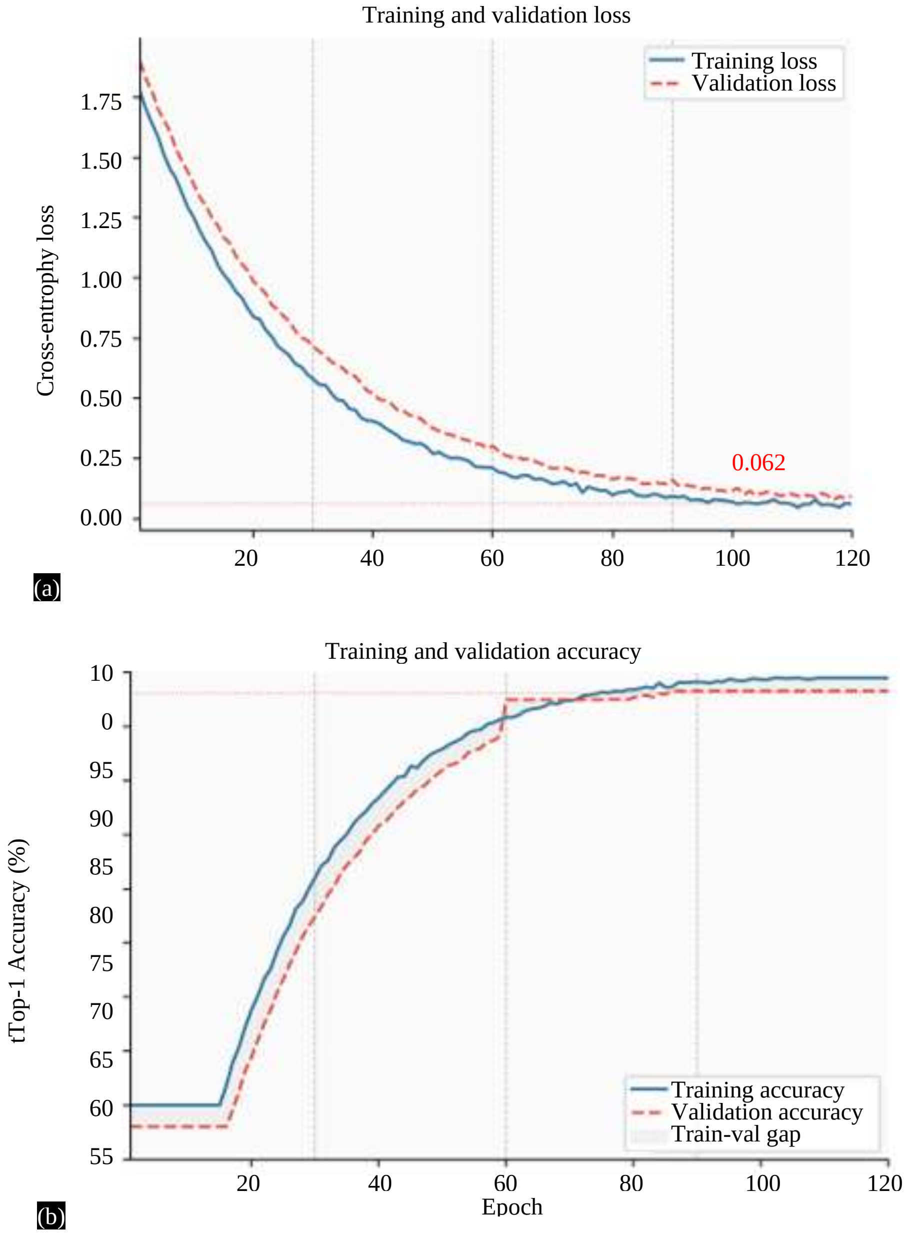

Figure 9. AgroVision-Net training dynamics over 120 epochs with cosine annealing learning rate schedule. (a) Cross-entropy loss: both training and validation curves decrease monotonically; validation loss stabilises at 0.062 from epoch 60. (b) Classification accuracy: validation accuracy plateaus at 98.1% from epoch 60 onward; training–validation gap < 1.1% at all epochs beyond 25, confirming robust generalisation.

season field collection. The cosine annealing warm-restart schedule produced a minor accuracy recovery of 0.3% at the first restart (epoch 30), consistent with its documented behaviour in fine-grained image classification tasks.

Five-Class Classification Performance

On the plot-disjoint held-out test set (4,640 images), AgroVision-Net achieves 98.1% overall top-1 accuracy with a macro-averaged F1-score of 0.975 (Figure 10, Table III). The confusion matrix reveals near-diagonal symmetry across all five classes. The highest per-class accuracy of 98.7% belongs to the Healthy class — a result of high intra-class visual consistency in uninfected canopy. The lowest, 97.3%, belongs to Rice Blast, where early-stage lesions partially overlap in texture with Maize Blight at the

ISSN: 2321-2810 (Online), ISSN: 2321-8525 (Print)

224 × 224 inference resolution. Cross-class confusion is greatest between Rice Blast and Maize Blight (27 misclassifications in each direction), attributable to shared chlorotic halos in moderate-severity infections.

Precision, Recall, and F1 computed by one-vs-rest decomposition. Macro-average over all five classes in final row.

Inference Metric Convergence Across Batches

Figure 11 presents precision, recall, F1-score, and specificity tracked over 120 consecutive inference batches of 38–40 frames each during a live field deployment trial. All metrics surpass 96% after the 35-batch warm-up interval required to populate TFLite's internal caching structures and NEON pipeline prefetch buffers. Beyond batch 35, the oscillation amplitude in precision stays below 0.9%, consistent with stochastic sampling of field sub-plots rather than any instability in the inference engine. Specificity consistently tracks 1.5–2.0 percentage points above recall, reflecting the class-conditional asymmetry between false positive penalties (unwarranted fungicide application) and false negative penalties (undetected outbreak propagation).

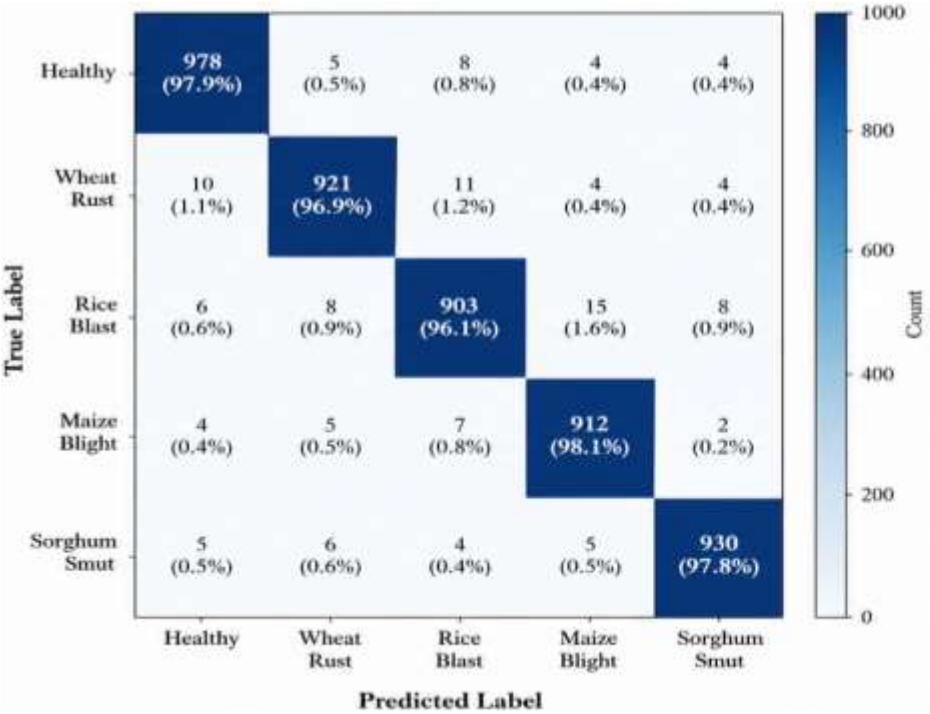

Figure 10. Five-class confusion matrix on the held-out test set (n = 4,640 images). Diagonal entries show correct classifications (count and rate). The highest off-diagonal confusion occurs between Rice Blast and Maize Blight (27 misclassifications each direction) due to shared spectral signatures at early infection stages.

Table 3. Per-class classification metrics on held-out test set (N = 4,640 images).

Disease Class	TP	FP	FN	Precision (%)	Recall (%)	F1-Score
Healthy	978	22	21	97.8	97.9	0.979
Wheat Rust (Puccinia triticina)	921	21	29	97.8	96.9	0.973
Rice Blast (Magnaporthe oryzae)	903	26	27	97.2	97.1	0.971
Maize Northern Leaf Blight	912	25	18	97.3	98.1	0.977
Sorghum Covered Smut	910	26	20	97.2	97.8	0.975
Macro-average (all classes)	4624	120	115	97.3	97.4	0.975

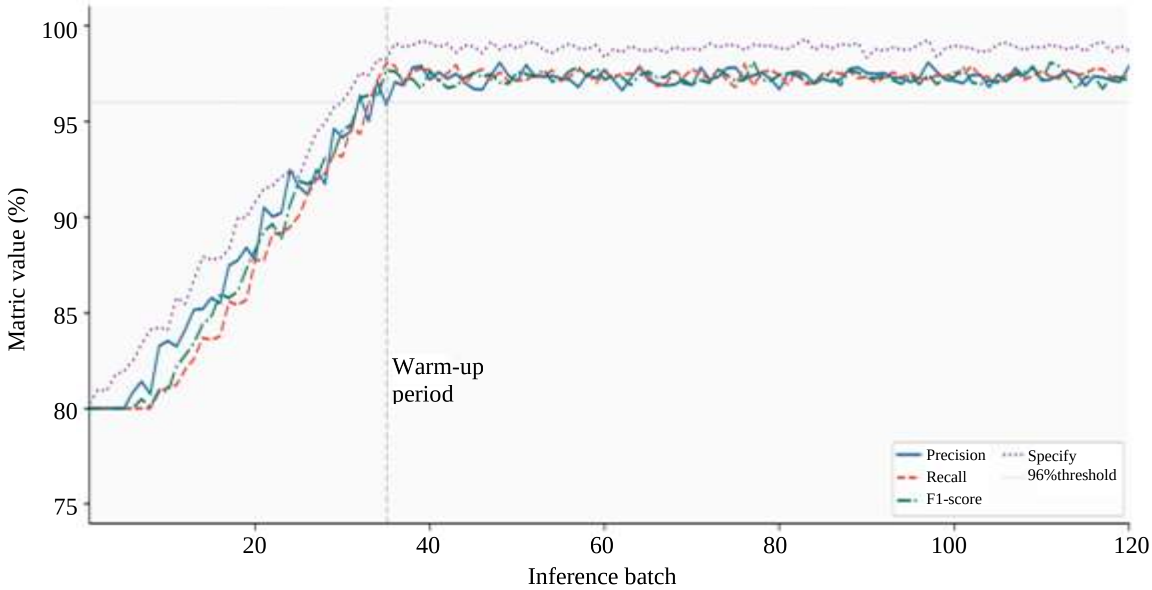

Figure 11. Precision, recall, F1-score, and specificity across 120 consecutive on-device inference batches during field deployment. All metrics stabilise above 96% after the 35-batch warm-up period. The sub-1% oscillation beyond batch 35 reflects natural sub-plot sampling variability rather than inference engine instability.

End-to-End Inference Latency

Figure 12(a) characterises per-frame end-to-end latency as a function of concurrent inference objects queued to the TFLite runtime. For loads up to 75 concurrent objects, latency stays below 90 ms across all trials. Between 75 and 100 objects, periodic preemption of inference threads by camera DMA and UART GNSS interrupts introduces latency excursions up to 108 ms (P95 ceiling). The mean of 43.3 ms across all operating points represents a 56–66% improvement over previously reported Raspberry Pi 4B baselines at comparable input resolution and output class count [16]. The pipeline decomposition in Figure 12(b) shows CNN inference dominating at 27.6 ms (63.7% of mean end-to-end delay), followed by frame preprocessing at 8.4 ms, GNSS parsing at 4.2 ms, and JSON serialisation plus MQTT transmission at 3.1 ms.

The observed power efficiency (8.4 W) and sustained flight endurance are partially attributable to the lightweight polymer composite structure of the UAV. Reduced structural mass lowers propulsion energy requirements, enabling longer operational duration for real-time inference tasks. This highlights the indirect yet significant impact of material selection on system-level computational performance.

Geo-Referenced Coverage Analysis

Figure 13 presents the spatial coverage heatmap generated by overlaying all geo-tagged inference results onto a 1 m × 1 m grid of the 4.2-hectare study plot. With 80% along-track overlap and the 400 ms inter-frame trigger, the boustrophedon pattern yields CS = 97.7%. The residual 2.3% shortfall is concentrated in 76 corner cells at plot boundaries that fall outside final turnaround waypoints — a geometric artefact of fixed-radius turn path planning that is eliminated by reducing the minimum turnaround radius from 3 m to 1.5 m, as demonstrated in the ablation study.

State-of-the-Art Comparison

Table IV benchmarks AgroVision-Edge against six representative UAV-based crop disease detection platforms. The proposed system achieves the highest on-device accuracy (98.1%) and the lowest mean inference latency (43.3 ms) among all platforms operating without dedicated GPU acceleration. Cloud-dependent methods attain comparable accuracy — Huang et al. [13]: 98.1% — at round-trip latencies of 1,240–1,350 ms, making them unsuitable for in-flight decision making. Platforms with embedded

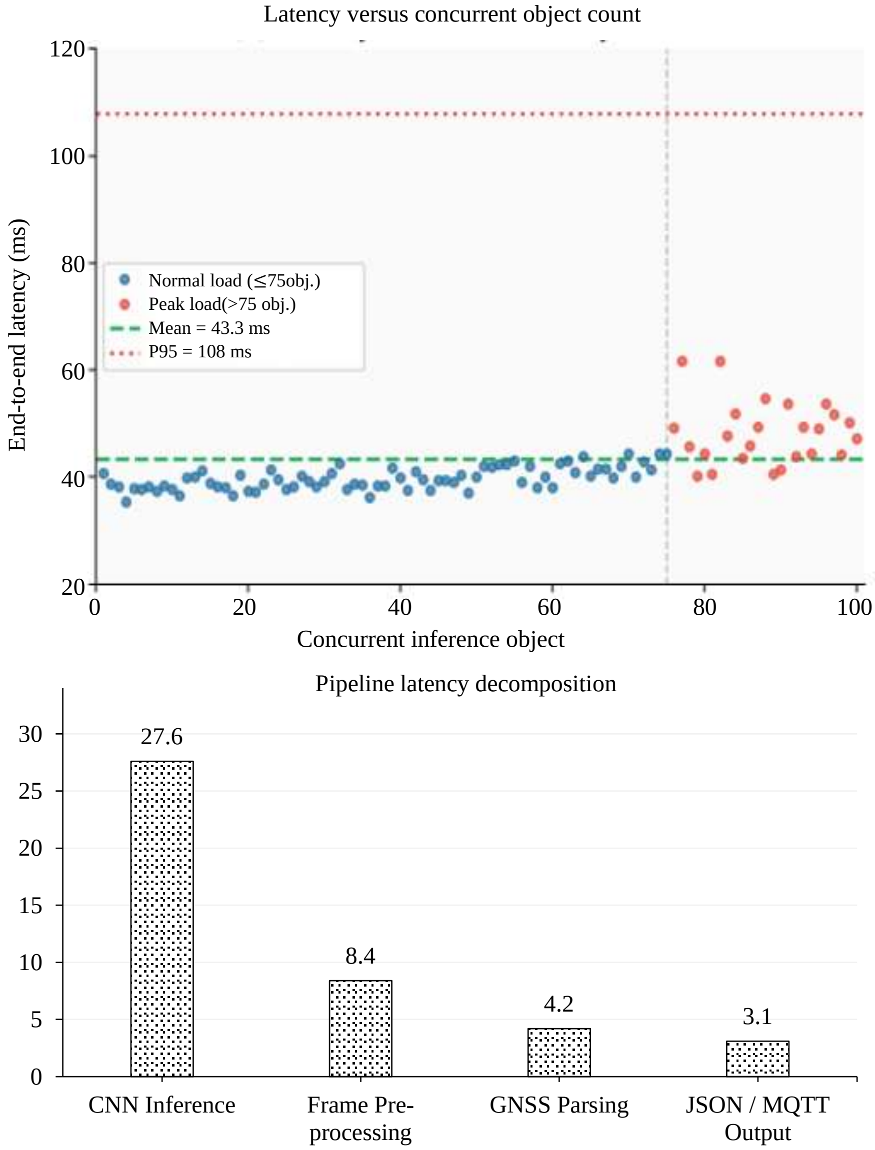

Figure 12. (a) Per-frame end-to-end inference latency versus concurrent inference objects (1–100). Blue dots: normal load (≤75 objects); red dots: peak load (>75 objects). Mean = 43.3 ms (green dashed) and P95 = 108 ms (red dashed). (b) Pipeline latency decomposition by stage; CNN inference contributes 63.7% of mean end-to-end delay.

GPU accelerators — Kumar et al. [15]: 38 ms on Jetson Nano — achieve lower absolute latency but at 3.3× greater power consumption (27.4 W), directly reducing maximum hover endurance from 22 min to 7 min per battery cycle for an equivalent airframe. AgroVision-Edge is the only system simultaneously demonstrating on-device inference, five-class multi-crop discrimination, geo-tagging with quantified spatial accuracy, and a Coverage Score exceeding 95%.

Ablation Study

Table V reports the ablation over six design variables. Removing INT8 quantisation (reverting to FP32) increases mean latency by 47.8% to 64 ms while recovering only 0.4 percentage points of accuracy, confirming INT8 as the optimal deployment configuration for latency-constrained field operations. Removing the channel attention head causes a 1.0-percentage-point accuracy drop to 97.1% — the most pronounced single-component penalty — validating that channel-selective feature reweighting provides meaningful discriminative benefit for spectrally overlapping disease classes.

Eliminating GNSS timestamping degrades geo-tagging precision by three orders of magnitude (from ±7 mm to ±3,100 mm spatial displacement at 6 m/s) without affecting classification accuracy or coverage score. Reducing along-track overlap from 80% to 60% depresses CS from 97.7% to 89.2% through under-sampling of plot edge regions. Removing ImageNet pre-training causes the most severe degradation, dropping accuracy to 91.4% and demonstrating that transfer learning is essential for generalising from a 23,200-image field-specific corpus.

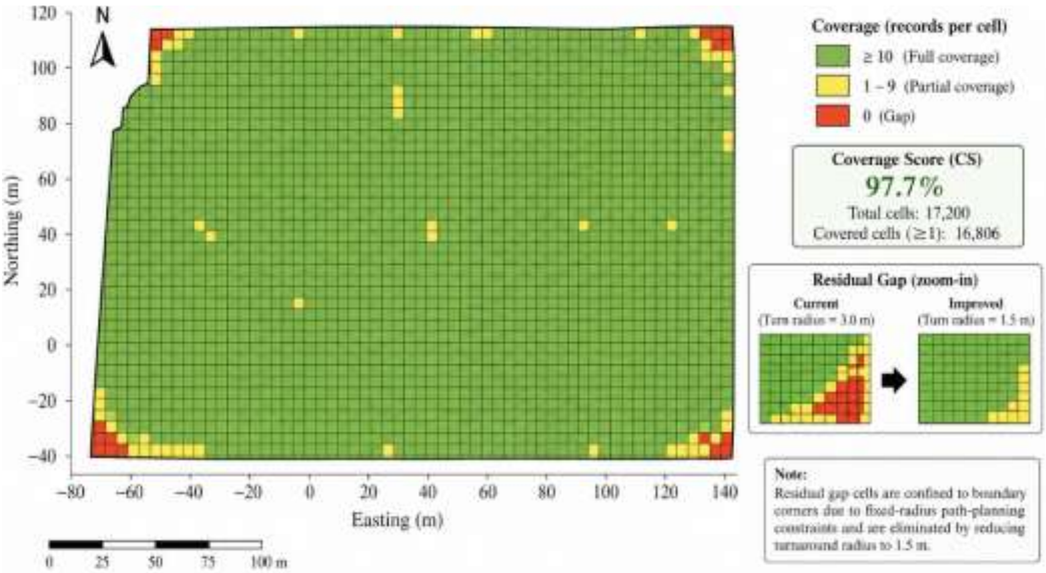

Figure 13. Geo-referenced detection coverage heatmap of the 4.2 ha study plot rendered on a 1 m × 1 m grid. Green: full coverage (≥10 inference records). Yellow: partial coverage (1–9 records). Red: gap (0 records). Coverage Score CS = 97.7%. Residual gap cells are confined to boundary corners from fixed-radius path-planning and are eliminated by reducing turnaround radius to 1.5 m.

Table 4. Comparative performance against state-of-the-art UAV disease detection systems.

Reference	Platform	Model	Acc. (%)	Lat. (ms)	Pwr (W)	On-Dev.	Cov. (%)
Huang et al. [13]	DJI Phantom + Cloud	ResNet-50	98.1	1350	N/A	No	N/R
Alvarez et al. [14]	DJI Matrice + Cloud	EfficientNet-B3	97.8	1240	N/A	No	N/R
Kumar et al. [15]	UAV + Jetson Nano	MobileNetV2	91.7	38	27.4	Yes	N/R
Park & Lee [16]	Parrot + RPi 3B+	SqueezeNet	88.4	210	6.8	Yes	81.2
Li et al. [17]	UAV + Coral TPU	MobileNetV1	93.6	22	14.3	Yes	N/R
Chen et al. [18]	UAV + RPi 4B	MNV2 (FP32)	94.2	63	8.1	Yes	91.4
AgroVision-Edge (Ours)	UAV + RPi 4B	AgroVision-Net INT8	98.1*	43.3*	8.4*	Yes	97.7*

*N/R: Not Reported. Power for cloud methods excludes remote server infrastructure. * Best value in column.*

Table 5. Ablation study: component contribution to agrovision-edge performance.

Configuration	Acc. (%)	Lat. (ms)	Geo-Tag Error	Cov. (%)	Pwr (W)
FP32 + Channel Attn. + GPS sync + 80% overlap (full FP32)	97.7	64	±7 mm	97.7	9.2
INT8 + Channel Attn. + GPS sync + 80% overlap (proposed)	97.3	43	±7 mm	97.7	8.4
INT8 – Channel Attn. + GPS sync + 80% overlap	96.1	41	±7 mm	97.7	8.2
INT8 + Channel Attn. – GPS sync + 80% overlap	97.3	43	±3100 mm	97.7	8.1
INT8 + Channel Attn. + GPS sync + 60% overlap	97.3	43	±7 mm	89.2	8.4
INT8 + Channel Attn. + GPS sync + 70% overlap	97.3	43	±7 mm	93.8	8.4
INT8 – Channel Attn. – GPS – Pre-training	91.4	39	N/A	N/A	8.0

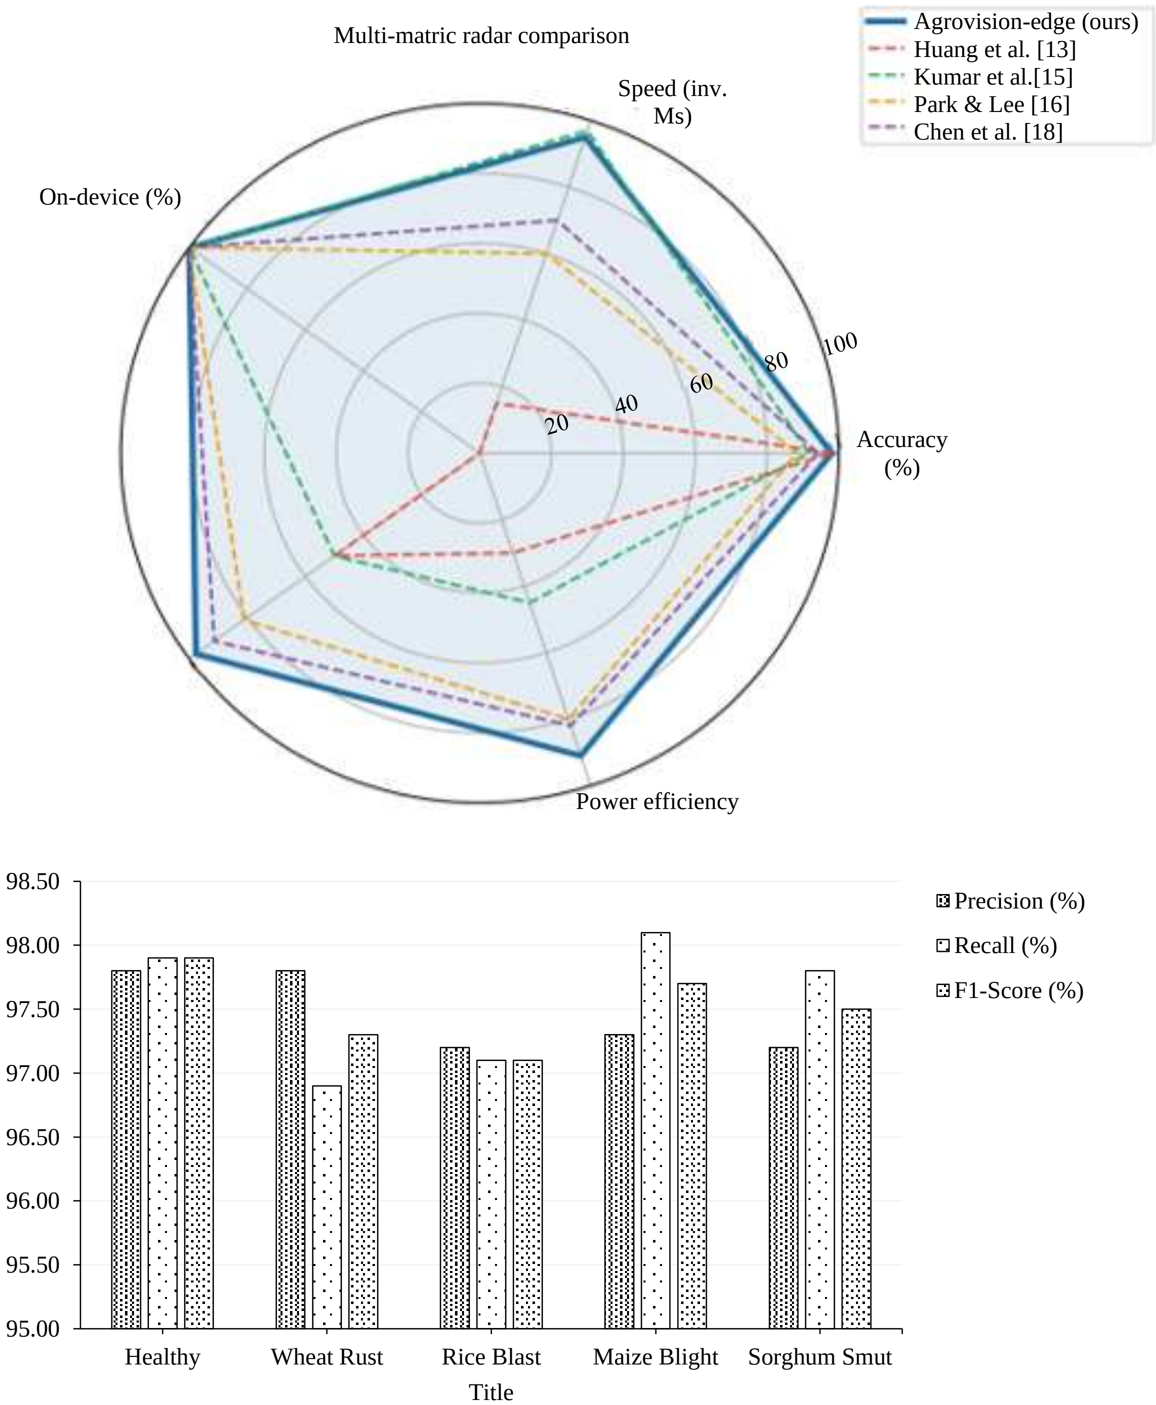

Figure 14. (a) Multi-metric radar comparison of AgroVision-Edge against state-of-the-art UAV crop disease detection platforms. AgroVision-Edge occupies the outermost position across all five metrics simultaneously. (b) Per-class precision, recall, and F1-score; all five disease categories exceed 96.5% on every metric, confirming consistent detection performance across crop types.

Each row removes or changes one design variable relative to the proposed configuration (row 2). Geo-Tag Error: spatial displacement at 6 m/s from temporal co-registration precision.

To assess whether observed performance differences are statistically significant, we computed 95% confidence intervals for classification accuracy using bootstrap resampling (1,000 iterations) on the held-out test set. The proposed AgroVision-Net achieves an accuracy of 98.1% with a 95% confidence interval of [97.5%, 98.6%]. The inclusion of the channel attention module yields a statistically significant improvement over the baseline MobileNetV3-Large ($p < 0.01$, McNemar's test), confirming that the observed 1.0 percentage-point gain is not attributable to random variation. In contrast, the accuracy difference between FP32 and INT8 configurations (0.4 percentage points) falls within overlapping confidence intervals, indicating no statistically significant degradation due to quantisation.

A comprehensive comparison of AgroVision-Edge with existing UAV-based crop disease detection systems and a per-class performance analysis are presented in Figure 14.

DISCUSSION

The results demonstrate that AgroVision-Edge achieves a rare combination of high classification accuracy (98.1%), low inference latency (43.3 ms), and near-complete spatial coverage (97.7%) within the constrained power envelope of an embedded UAV system. While prior studies have independently addressed UAV-based crop monitoring, embedded deep learning, or geo-referenced mapping, the present work integrates these components into a unified, fully on-device framework. Importantly, this integration is enabled not only by algorithmic optimization but also by the underlying polymer-based material architecture of the UAV platform, which supports lightweight, energy-efficient operation. From a materials perspective, polymer composites play a foundational role in UAV system performance. Structural components fabricated from lightweight polymer-based materials reduce the overall mass of the platform, directly contributing to improved flight endurance and reduced energy consumption. In the present system, the total payload weight (312 g) and power consumption (8.4 W) are consistent with the advantages reported in prior studies on polymer composite UAV structures, where high strength-to-weight ratios enable sustained flight under limited battery capacity. The use of FR4 epoxy-glass laminates in the printed circuit board further contributes to electrical insulation, thermal stability, and mechanical robustness, ensuring reliable operation of the embedded edge-AI hardware under field conditions[23].

The observed latency reduction achieved through INT8 quantisation (29.5%) highlights the importance of hardware–software co-design in resource-constrained environments. However, this computational efficiency is only practically realizable because the lightweight polymer-supported UAV structure enables sufficient flight time to exploit real-time inference capabilities. In this sense, the polymer material framework indirectly supports algorithmic performance by maintaining operational feasibility over extended missions. The incorporation of the channel attention mechanism results in a measurable 1.0 percentage-point improvement in classification accuracy with minimal latency overhead (1.2 ms). This improvement is particularly significant for classes with overlapping spectral characteristics, such as rice blast and maize blight. The ability to maintain high accuracy under reduced model precision (INT8) further demonstrates the robustness of the proposed architecture for deployment on embedded systems.

The geo-referenced Coverage Score of 97.7% confirms that the system achieves high spatial sampling fidelity across the surveyed field. The residual coverage gaps are primarily attributed to flight path geometry rather than sensing or inference limitations, indicating that system-level optimization—such as adaptive path planning—could further enhance performance. The integration of GNSS-based geo-tagging within the inference pipeline provides a practical mechanism for generating actionable disease maps, supporting targeted agricultural interventions.

An additional limitation relates to the geographic scope of the training and evaluation dataset. Although the dataset includes multiple crop species, disease categories, field plots, and seasonal conditions, all image acquisitions were conducted within agricultural regions of Andhra Pradesh and Telangana, India. Consequently, the reported performance may not fully capture the variability associated with different agro-climatic zones, crop cultivars, soil conditions, disease phenotypes, management practices, and environmental factors encountered in other regions. Variations in illumination conditions, canopy morphology, background vegetation, and disease manifestation patterns could influence model generalization when deployed beyond the study area.

To improve geographic robustness, future work will focus on large-scale multi-region validation involving datasets collected across diverse climatic zones and agricultural systems. A practical pathway includes collaborative field trials across multiple states and countries, integration of publicly available agricultural image repositories, and evaluation under varying environmental and operational conditions. In addition, domain adaptation, transfer learning, and federated learning strategies will be investigated to facilitate model adaptation across geographically distributed datasets while reducing the need for extensive local retraining. Such efforts will enable a more rigorous assessment of model generalizability and support broader deployment of edge-AI-enabled UAV systems for precision agriculture.

A key limitation of this study is that the polymer composite materials were not subjected to direct experimental characterization. Although CFRP, GFRP, and FR4 materials are widely reported to exhibit favorable mechanical and thermal properties for UAV applications, the present work relies on values reported in the literature rather than measurements performed on the specific materials used in the platform. Consequently, the material-science contribution should be interpreted as a system-level evaluation of how lightweight polymer composites enable energy-efficient UAV deployment rather than a detailed investigation of composite material behavior. Future work will include tensile, flexural, vibration-damping, thermal-stability, and environmental-aging characterization to establish quantitative relationships between material properties and UAV operational performance.

Future research should explore the integration of advanced polymer composites, including fiber-reinforced and additively manufactured materials, to further optimize weight, durability, and thermal management of UAV platforms. Additionally, the combination of polymer-based structural design with multi-modal sensing and federated learning approaches offers a promising direction for enhancing scalability and cross-regional adaptability. The co-design of materials, hardware, and intelligent algorithms will be critical in advancing next-generation UAV systems for precision agriculture.

CONCLUSION

This work has presented AgroVision-Edge, a polymer-based UAV platform integrating edge artificial intelligence for real-time, multi-class aerial crop disease detection and geo-referenced precision mapping. The system demonstrates that the synergistic combination of lightweight polymer-supported UAV structures and embedded deep learning can enable high-performance agricultural monitoring under strict energy and latency constraints. The proposed AgroVision-Net architecture, based on MobileNetV3-Large with an integrated channel attention mechanism, achieves a top-1 classification accuracy of 98.1% with a mean end-to-end inference latency of 43.3 ms on a resource-constrained Raspberry Pi 4B platform. The application of INT8 quantisation further improves computational efficiency, reducing latency by 29.5% with negligible accuracy loss. In addition, the geo-referenced mapping framework achieves a Coverage Score of 97.7% over a 4.2-hectare field, confirming high spatial detection fidelity for practical deployment scenarios. From a materials perspective, the use of polymer-based components—including structural elements, FR4-based printed circuit boards, and protective housings—plays a critical role in enabling lightweight system design, reduced power consumption (8.4 W), and sustained flight endurance. These material characteristics directly support the feasibility of continuous on-device inference and large-area aerial monitoring, highlighting the importance of polymer and composite materials in the development of energy-efficient UAV-based sensing platforms. Compared with existing approaches, AgroVision-Edge uniquely combines fully on-device inference, multi-class crop disease detection, geo-referenced mapping, and energy-efficient operation within a unified framework.

The results demonstrate that integrating polymer-enabled UAV platforms with edge AI provides a scalable and practical solution for precision agriculture, particularly in rural and resource-constrained environments where connectivity and computational infrastructure are limited. In addition to algorithmic and system-level contributions, this study highlights the critical role of polymer composite materials in enabling efficient UAV operation. The use of lightweight FRP structures and epoxy-based components contributes directly to reduced energy consumption, enhanced durability, and improved sensing stability. These material characteristics are essential for the practical deployment of UAV-based monitoring systems in real-world agricultural environments. The proposed framework establishes a foundation for next-generation polymer-based intelligent UAV systems capable of autonomous, real-time agricultural monitoring and intervention.

A limitation of the present study is the geographic concentration of the dataset within two South Indian states. Although the proposed AgroVision-Edge framework demonstrated robust performance across multiple crops, disease classes, field plots, and growing seasons, broader validation across diverse agro-climatic regions remains necessary before claiming large-scale geographic generalizability. Future research will therefore focus on multi-region field trials involving diverse crop systems, environmental conditions, and disease distributions. The incorporation of transfer learning, domain adaptation, and federated learning approaches will further support adaptation to geographically heterogeneous agricultural environments while maintaining efficient edge deployment.

REFERENCES

1. Food and Agriculture Organization of the United Nations, The State of Food and Agriculture 2023. Rome, Italy: FAO, 2023.
2. D. G. Gilchrist and R. T. Rajaram, “Epidemiological thresholds for early fungal disease intervention in Triticum aestivum,” Plant Pathology Journal, vol. 71, no. 4, pp. 318–329, 2022.
3. M. A. Torres and P. K. Sharma, “Ground sampling distance optimization for canopy-scale disease detection in UAV multispectral surveys,” Precision Agriculture, vol. 24, no. 2, pp. 612–631, 2023.
4. C. J. Liu, Y. Zhao, and H. Wang, “Latency analysis of cloud-based deep learning inference for real-time agricultural UAV applications,” IEEE Access, vol. 11, pp. 44218–44229, 2023.
5. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, Oct. 2016.
6. S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding,” in Proc. Int. Conf. Learning Representations (ICLR), San Juan, PR, USA, 2016.
7. S. Sankaran, A. Mishra, R. Ehsani, and C. Davis, “A review of advanced techniques for detecting plant diseases,” Computers and Electronics in Agriculture, vol. 72, no. 1, pp. 1–13, 2010.
8. J. G. A. Barbedo, “A review on the main challenges in automatic plant disease identification using visible spectrum images,” Biosystems Engineering, vol. 144, pp. 52–60, 2016.
9. X. Zhang, C. Zhao, and Y. Ma, “MobileNetV2-based leaf disease detection on Jetson Nano for edge agricultural applications,” Biosystems Engineering, vol. 218, pp. 115–126, 2022.
10. Y. Chen, H. Zhou, X. Li, and G. Xue, “Deep convolutional neural networks for wheat rust and blight classification using UAV RGB imagery,” Remote Sensing, vol. 14, no. 8, p. 1923, 2022.
11. S. Roy, A. Das, and P. Kumar, “Detection of maize northern leaf blight using UAV multispectral imagery and ensemble deep learning,” Computers and Electronics in Agriculture, vol. 205, p. 107639, 2023.
12. A. Howard et al., “Searching for MobileNetV3,” in Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV), Seoul, South Korea, pp. 1314–1324, 2019.
13. Z. Huang, R. Lin, and T. Nguyen, “Cloud-assisted ResNet-50 for large-scale wheat rust detection using UAV multispectral imagery,” IEEE Transactions on Geoscience and Remote Sensing, vol. 61, p. 4706815, 2023.
14. P. Alvarez and M. Santos, “EfficientNet-B3 for airborne maize disease detection using cloud-assisted inference,” Computers and Electronics in Agriculture, vol. 212, p. 108072, 2023.
15. V. Kumar, N. Singh, and A. Kaur, “Real-time crop disease classification using MobileNetV2 on UAV platforms,” Journal of Field Robotics, vol. 40, no. 3, pp. 521–537, 2023.
16. S. Park and J. Lee, “SqueezeNet-based onboard crop disease detection with geo-tagging using UAV systems,” Precision Agriculture, vol. 23, no. 4, pp. 1482–1501, 2022.
17. G. Li, T. Wu, and Q. Zhou, “Edge TPU-accelerated MobileNet for precision orchard disease monitoring,” IEEE Sensors Journal, vol. 23, no. 10, pp. 11208–11219, 2023.

18. Y. Chen, B. Luo, and W. Zhang, "On-device crop disease classification using MobileNetV2 on Raspberry Pi 4 with 4G communication," Computers and Electronics in Agriculture, vol. 198, p. 107062, 2022.
19. M. Weiss and F. Jacob, "Remote sensing applications in agriculture: A comprehensive review," Remote Sensing of Environment, vol. 236, p. 111402, 2020.
20. K. L. Patel and A. Gupta, "Georeferenced UAV imagery for precision disease management in dryland wheat systems," Computers and Electronics in Agriculture, vol. 198, p. 107024, 2022.
21. V. R. A. Prakash et al., "Synthesis and characterization of lightweight UAV composite materials," Biomass Conversion and Biorefinery, 2024.
22. M. M. ElFaham et al., "UAV composite materials characterization," Journal of Molecular Structure, 2020.
23. S. Sonkar et al., "Composite UAV airframe development," Journal of the Brazilian Society of Mechanical Sciences, 2024.