International Journal of Engineering Trends and Technology
ISSN: 2231-5381/ https://doi.org/10.14445/22315381/IJETT-V74I8P136

Volume 74 Issue 8, 535-548, August 2026

Original Article

An Adaptive AI-DSS Framework for Real-Time Image Analysis and Decision-Making in Precision Agriculture

Yebhushi Prashanth[1], Dr.Manna SheelaRani Chetty[2]

[1,2]Department of CSE,KLEF, Vaddeswaram, Andhra Pradesh, India.

[1]Corresponding Author : yprashanth.60@gmail.com

Received: 19 December 2025 Revised: 06 March 2026 Accepted: 12 March 2026 Published: 29 August 2026

Abstract - *Environmental variability, pest infestation, and resource inefficiency have become increasingly problematic to agriculture, and on-line decision making using intelligent, adaptive technologies has been called for. This research is suggesting a novel Integrated Artificial Intelligence Decision Support System (AI-DSS) in precision agriculture so as to attain proper context-aware analysis and recommendation across dynamic field conditions. The framework is based on lightweight deep learning frameworks (MobileNetV2, Efficient Net-lite) and a newly designed Adaptive Feature Optimization (AFO) engine, which dynamically reweights convolutional features based on the temporal stability and environment consistency. Mathematically, the AFO mechanism is obtained via a weighted pooling of the instantaneous features of CNN with the temporal averaged prototypes using adaptive attention weights, which filter out transient noises due to illumination variations, occlusion, and sensor noises. The optimized features are fused with data from environmental and soil sensors in a hybrid Decision Support System (DSS) based on rule-based reasoning, Bayesian inference, and temporal tracking to construct explainable and region-specific recommendations to farmers. Experimental evaluations using the experimental data PlantVillage, DeepWeeds, and Fieldstream Sim demonstrate the superiority of the proposed AFO enhanced framework over the baseline CNN classifier in terms of classification accuracy (.95), macro F1 score (.95), and robustness to distortions (.15% improvement). Additionally, when deployed on edge devices like Raspberry Pi 4 and Nvidia Jetson Nano, the system achieves real-time inference latency (<200 ms) and low energy consumption (~520 mJ/frame), which validates the scalability of the system in low-resource settings. In addition, the expert agreement was enhanced with the DSS module integration to 91.4% with 57% less false alarms. The obtained results validate the proposed AI-DSS framework with AFO as a promising solution to cover the distance between accuracy in a controlled lab environment and reliability in the field, providing a strong, explainable, and resource-efficient solution to the digital sustainable agriculture problem. This solution is further being expanded into proactive farm intelligence and climate resilience through multimodal sensing, satellite assisted crops monitoring, and adaptive decision-making-led solutions.*

Keywords - *Artificial Intelligence, Adaptive Feature Optimization (AFO), Decision Support System (DSS), Precision Agriculture, Edge computing, Explainable Artificial Intelligence.*

1. Introduction

Nearly half of the world's population depends on agriculture for food and livelihood, making agriculture an important sector for sustainable development and food security [1]. In recent years, this infusion of Artificial Intelligence (AI) and Machine Learning (ML) has revolutionized the fields of Artificial Intelligence, Adaptive Feature Optimization (AFO), Decision Support System (DSS), Precision agriculture, Edge computing, and Explainable artificial intelligence. conventional farming into precision agriculture with capabilities of monitoring crop health, disease detection, yield prediction, and pest management, automated [2, 3]. Despite such advancements, one of the basic limitations still exists - most of the AI-based systems are reactive in nature rather than proactive and provide delayed insights and hence limit timely intervention in uncertain or rapidly changing field conditions [4].

Conventional agricultural practices include manual inspection and visual checks at periodic intervals and intuition of the farmers. These approaches are slow, subjective, and often end up with input wastage and reduced yields, apart from financial losses due to delayed responses to pest attacks, disease start, or weather changes [5]. However, although AI-based tools are accurate in controlled laboratory settings, effective, context-aware decisions from these tools in a rural setting are often not actionable, especially where there is limited connectivity, power, and technical expertise [16, 28]. A main problem is the lack of connection between the image analysis modules based on artificial intelligence and Decision Support Systems (DSS). While Convolutional Neural

Networks (CNNs) have a classification accuracy of more than 90% on controlled data sets, their use in open-field environments is faced with problems such as illumination variations, occlusions, soil and crop variation, and environmental noise [8, 9]. Furthermore, available models are highly dependent on huge annotated datasets that are expensive to collect and not necessarily transferrable across locations and crops [10]. The computational complexity of the deep learning models also imposes convenient constraints for the real-time on-device inference on edge environments [11, 12].

The other significant limitation is a lack of integrated decision-making. Many current applications of AI provide an output of disease or stress label, but do not provide interpretable, localized, and timely recommendations. Consequently, farmers still need to rely on experts or digital intermediaries in order to interpret the model predictions [25, 36]. Therefore, the objective is not only to improve the accuracy of models, but to associate the AI inference with a decision engine to adaptively provide the farm-level recommendations to ensure fast and reliable, explainable recommendations.

1.1. Socio-Economic and Human-Centered Design

The proposed AI-DSS framework incorporates socio-economic and human factors as core design principles to ensure practical usability and adoption in real-world agricultural settings. The system is designed to support low digital literacy through simplified user interfaces, explainable recommendations, and multi-language delivery via mobile applications, web dashboards, and SMS/audio alerts. Cost sensitivity and connectivity constraints are addressed through edge-based deployment, minimizing reliance on cloud infrastructure and reducing operational expenses for smallholder farmers.

To promote adoption and trust, participatory design principles are integrated into the decision support workflow, allowing agronomists and farmers to configure decision rules, alert thresholds, and advisory preferences. User feedback mechanisms enable confirmation, refinement, or override of system-generated recommendations, facilitating continuous learning and contextual adaptation. These human-centered features directly contribute to improved expert agreement, reduced false alerts, and more timely interventions, demonstrating that the proposed framework aligns technical intelligence with socio-economic realities and user-driven decision-making.

1.2. Motivation

The motivation of this research is to fulfill the urgent need to design an artificial intelligence integrated knowledge base Decision Support System (AI-DSS) with the ability of real-time analysis and decision-making under the variability of the field. Agricultural operations are highly dynamic - lighting, humidity, pest prevalence and soil moisture - change hourly. A smart DSS must thus combine image interpretation relying on artificial intelligence with contextual reasoning, which allows the immediate, autonomous and actionable recommendations. Developing such a system is particularly important in environments with limited resources, where access to cloud computing, internet connectivity and trained agronomists is limited [14].

As an example of using AI models in edge computing, the use of MobileNetV2 or EfficientNet-lite models at the edge can be used to perform inference in low-power devices like Raspberry Pi or Jetson Nano, allowing for field deployment without being constantly dependent on the cloud.

Furthermore, coupling AI inference with temporal reasoning within DSS modules enables continuous monitoring, proactive alert generation, and human-understandable communication (via mobile apps, SMS, or audio in regional languages).

This research aims to democratize AI-assisted farming, making precision agriculture accessible, affordable, and interpretable for smallholder farmers.

1.3. Research Gaps

Although substantial progress has been made in applying AI for crop health monitoring, several critical gaps persist:

1. Lack of Real-Time Decision Integration: Most AI systems are trained and evaluated in isolation and do not integrate with decision-making workflows or on-field control systems [26, 27]. This prevents them from being directly used for farm automation.
2. Limited Robustness Across Conditions: CNN-based models such as ResNet or DenseNet exhibit high accuracy on benchmark datasets (e.g., PlantVillage) but perform inconsistently under field-level variability (e.g., DeepWeeds), due to lighting fluctuations and occlusions [9, 10].
3. Dependence on Large Annotated Datasets: Collecting and labeling domain-specific agricultural images is expensive and regionally biased. Transfer learning approaches remain under-explored for cross-domain adaptation in agriculture [11].
4. High Computational Overhead: State-of-the-art CNN architectures require GPUs or cloud-based inference, unsuitable for low-resource rural infrastructures [12, 14].
5. Lack of Explainable and Contextual DSS: Even when AI models are accurate, the decision-making logic remains opaque and lacks contextual understanding, such as growth stage, recent rainfall, or pest risk [25, 26].

Filling these gaps requires an integrated and explainable as well as resource-efficient AI-DSS framework that could be real-time adapted and scalable deployed.

1.4. Objectives

In order to overcome the above-mentioned challenges, this study proposes and evaluates the design and evaluation of a novel Artificial intelligence-based Decision Support System (AI-DSS) for real-time precision agriculture. The basic objectives are as under:

1. To develop an efficient and domain-adaptive learning AI solution (based on the MobileNetV2 and EfficientNet-lite architectures) to obtain high accuracy and robustness under variable field conditions using transfer learning and lightweight optimization.
2. Development of and integration of a Decision Support Module (DSS) to be able to generate contextual recommendations (e.g. irrigation, disease control) by integrating model inference, sensor data and expert rule bases.
3. To build and deploy the AI-DSS framework on edge-based computing platforms for low-latency inference and communication for real-time use in the field.
4. To deploy the proposed system on benchmark data sets (PlantVillage, DeepWeeds) and a synthetic Real-Time stream in a controlled and field-mimicking environment.
5. Preparing the system to be scalable in low-resource environments, by assessing the system's performance in terms of quantitative (accuracy, F1, robustness, latency) and qualitative (expert agreement, alert timeliness) metrics;

2. Literature Survey

K Naresh et al. [30] (2020) provided an overview of the overall potential of Artificial Intelligence (AI) for the Indian precision farming system. Their paper highlights some of the productivity improvements in subtropical crop production by AI-assisted irrigation, yield prediction and pest management, and infrastructural and data shortcomings in the smallholder context. Gupta, G et al. [31] (2020) further cemented AI in precision agriculture, finding supervised learning algorithms to be the leading tools used for image-based plant disease diagnosis and farm management tasks.

Rane et al. [32] (2025) reported a systematic study of object detection models such as YOLO and Faster R-CNN applied to agriculture, and in that paper, the challenge encountered in terms of small object sizes and occlusions and the lack of annotated field data.

Zualkernan et al. [33] (2023), in a survey about UAV-based imagery and precision agriculture machine learning pipelines, explain that the worse performance of heavier backbones architectures is due to their energy consumption in drones.

Kamilaris and Prenafeta-Boldud [3] (2018) published one of the first comprehensive reviews of deep learning in agriculture, and it is divided into plant-phenotyping, disease-detection, weed-identification and yield-estimation application domains. They have come up with the conclusion that domain generalization is a major challenge, even when one gets impressive results in the lab. Mohanty et al. [8] (2016) were the first to use deep learning for plant disease detection on the PlantVillage dataset using AlexNet and GoogLeNet architectures, showing the accuracy of more than 99%, this set an important benchmark for later work. Olsen et al. [9] (2019) suggested the DeepWeeds dataset, which consists of 17,509 field images of eight weed species, and in turn providing realistic testbed for robustness evaluation in an uncontrolled environment.

Liakos et al. [2] (2018) reviewed the evolution of machine-learning techniques in agriculture, where the basic algorithms are found to be neural networks, decision trees, and SVMs, with transfer learning coming up as a means for reducing the cost of scale labeling. Sandler et al. [13] (2018) introduced MobileNetV2, which contains inverted residuals and linear bottlenecks for efficient mobile inference, and the architecture has been used as the foundation for power-constrained applications such as edge-based crop monitoring. Similarly, Tan and Le [14] (2019) proposed EfficientNet, which scales depth, width and resolution in a compounded way to maximize accuracy-to-computation ratio.

Howard et al. [34] have introduced the first light-weight CNN (MobileNet), which has demonstrated the use of depthwise-separable convolutions to provide high accuracy on embedded hardware. Redmon and Farhadi [35] (2018) have further expanded on the real-time object detection based on YOLOv3, which was further tailored for the detection of fruits and pests in field robotics. The LIME explainability paradigm was suggested by Ribeiro et al. [36] (2016) and has been widely adopted for the interpretation of black box decisions from CNN for the context of agricultural DSS applications. He et al. [37] (2016) proposed ResNet, which uses ResNet to mitigate the vanishing gradient problem, which has become a backbone of many precision agriculture classifiers.

Since field-deployed CNN backbones rarely generalize well when transplanted directly from controlled benchmarks, several studies have examined how learned representations transfer across domains. Yosinski et al. [18] showed that early convolutional layers learn generic, broadly transferable features while later layers become increasingly task-specific, a finding that motivates partial fine-tuning strategies for agricultural CNNs. Weiss et al. [19] surveyed transfer-learning paradigms more broadly and catalogued the conditions under which source-to-target knowledge transfer succeeds or degrades. To close the remaining domain gap between laboratory and field imagery, Ganin et al. [20] proposed domain-adversarial training, which learns features that are simultaneously discriminative for the task and indistinguishable across domains, and Tzeng et al. [21]

extended this idea with adversarial discriminative domain adaptation for unsupervised target-domain alignment. Complementing these supervised and adversarial approaches, Chen et al. [22] introduced the SimCLR contrastive-learning framework for learning visual representations without labels, while He et al. [23] (2020) proposed Momentum Contrast (MoCo) to make contrastive pretraining practical at scale; both works underpin the rationale for pretraining lightweight backbones such as MobileNetV2 and EfficientNet-lite before agricultural fine-tuning.

Beyond representation learning, the interpretability of model decisions is essential for farmer trust in an AI-DSS. Selvaraju et al. [24] introduced Grad-CAM, which produces visual saliency maps that highlight the image regions driving a CNN's prediction, and Doshi-Velez and Kim [25] laid out a rigorous framework for evaluating the interpretability of machine learning systems more generally; both inform the explainability requirements of the DSS layer proposed in this work.

On the decision-support side, McCown [26] traced the historical and socio-technical difficulties of deploying agricultural DSS tools for real-world farm management, noting persistent gaps between model outputs and adoptable advice, while Fountas et al. [27] reviewed decision support systems in precision agriculture and identified integration with sensor data and rule-based reasoning as a recurring bottleneck. At the sensing and edge layer, Shi et al. [16] outlined the broader vision and open challenges of edge computing for latency-sensitive, resource-constrained applications, and Boursianis et al. [17] surveyed IoT devices and UAV platforms used in smart farming, underscoring their role in continuous field data collection. Jawad et al. [28] further reviewed energy-efficient wireless sensor network designs for precision agriculture, a consideration directly relevant to long-term, battery-powered field deployment. Moreover, Rahnemoonfar and Sheppard [29] used deep learning to perform automated fruit counting, both reinforcing the case for robust, field-tolerant visual recognition pipelines.

However, there are still areas where model robustness, latency and contextual quality of decision-making leave a lot to be desired. Most works are optimized to be accurate, but neglecting the real-time and actionable DSS integration. The Integrated AI-DSS Framework, which is proposed in the current research work, addresses these gaps by combining a new EfficientNet-lite field inference adaptation with a decision optimization engine that supports fail-safe context-aware recommendations.

3. Methodology

The proposed methodology introduces an Integrated AI-DSS (Artificial Intelligence- Decision Support System) framework for real-time precision agriculture, which deals with strong crop analysis, environment awareness and decision actions delivery. The system utilizes the concepts of light-weight deep learning, adaptive optimization, and context-based decision reasoning for resource-limited farming scenarios. Figure 3 shows the overall workflow, which involves six stages sequentially - from data acquisition to dissemination of the information to the farmers.

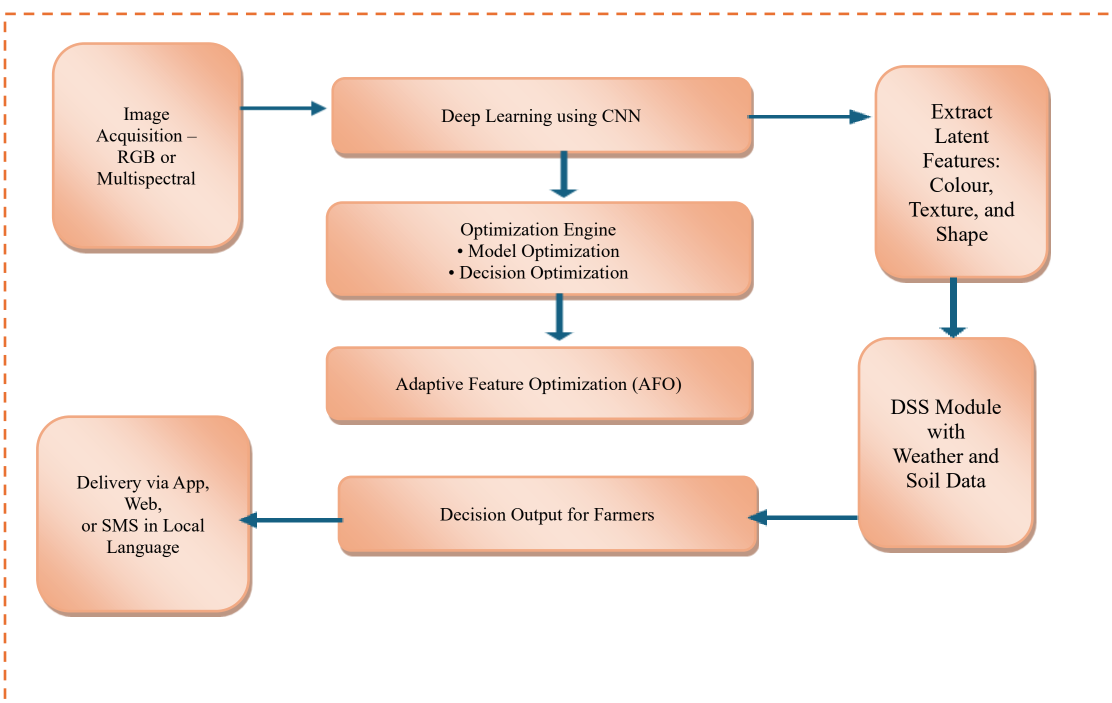

Fig. 1 Architecture of integrated AI-DSS (Artificial Intelligence-Decision Support System)

3.1. Image Acquisition

The image acquisition component is in charge of capturing real-time information on the crop using RGB or multispectral sensors integrated in IoT-driven ground cameras and low-altitude drones. These sensors are used to capture high-resolution imagery under a variety of environmental conditions for the detection of crop health parameters such as disease symptoms, pest infestations and nutrient deficiencies. While RGB sensors are a cheap solution with sufficient spatial resolution for most agricultural uses, multispectral sensors are applied in high-end systems where spectral analysis can reveal small changes in the physiological conditions of crops. All captured images are automatically sent to a local edge server or gateway device for immediate pre-processing and analysis to preserve low latency, real-time responsiveness and work offline for connectivity-limited rural areas.

3.2. Deep Learning for Image Processing

In this stage, the received images are processed using Convolutional Neural Networks (CNNs), namely MobileNetV2 and EfficientNet-lite, that were selected for its lightweight architecture and computational efficiency as well as their high accuracy for inference on embedded and edge devices. The CNNs have been automatically utilized for mining discriminating visual features of crop images to help recognize various patterns related to plant health. These are: colour distribution, which determines the content of nutrients and the amount of chlorophyll; irregularities in texture, which determine the infestation of pests or fungi; morphological distinctions, which determine the stress of growth, dehydration or wilting. Through the hierarchical convolutional layers, the model learns spatially invariant representations of these visual cues, and this makes the model robust under varied lighting and environmental conditions. Formally, the output of the CNN is a feature vector fi=[f1,f2,..., fn], where i is the index of the input image, and fj is a latent characteristic of the i-th input image. These extracted features are further forwarded to the Adaptive Feature Optimization (AFO) module to make context-aware refinement and enhance decision reliability.

3.3. Feature Representation and Extraction

After CNN processing, latent feature vectors are extracted from the penultimate layer in order to extract fine-grained patterns across plant surfaces. The features are both spatial and spectral, which form the basis for classification and contextual reasoning. Data augmentation methods such as random rotation, flipping, and brightness normalization are used before training the model to ensure that the model is robust to the change of lighting and background.

3.4. Adaptive Feature Optimization (AFO) Engine

To make up for the environmental changes (e.g. light, shade, occlusion, etc.), the Adaptive Feature Optimization (AFO) mechanism is presented to dynamically reweight features based on their reliability and temporal consistency. To counter the variation brought by the dynamic field environments (lighting changes, shadowing, occlusions and sensor noise), we present an Adaptive Feature Optimization (AFO) engine to dynamically reweight CNN-derived features based on their temporal reliability and contextual stabilities.

The AFO module is used to assess the temporal dispersion of each feature dimension in a set of recent frames and to fuse such dispersion with real-time information on the environment (e.g. illumination, humidity, motion intensity) to calculate an adaptive attention coefficient a_i for each feature. Mathematically, the optimization is $f_i'=a_i f_i+(1-a_i)\bar{f}$, where f_{id} represents the instantaneous feature vector of CNN, $\bar{f}$ is the average feature vector in the temporal window, and represents adaptive attention weight computed from feature variance and contextual continuity, where represents the weight of the attention mechanism, and value is between 0 and 1. This formulation ensures that unreliable or noisy features (which may be caused by transient perturbation (glare, motion blur, background noise, etc.) are adaptively down-weighted, making the system focus on meaningful and stable features.

Therefore, the AFO engine is also serving two important functions, namely Model Optimization for stabilizing feature representations for better classification performance and Decision Optimization for optimizing prediction confidence before feeding into the Decision Support System (DSS). Combined, this mechanism forms a new self-stabilizing feature adaptation layer that significantly increases robustness, interpretability, and real-world reliability of the AI-DSS framework.

3.5. New Adaptive Feature Optimization (AFO) Engine Algorithm

The proposed Adaptive Feature Optimization (AFO) engine proposes a dual-domain adaptive weighting strategy to jointly exploit the temporal stability and awareness of the environmental context to dynamically refine CNN feature representations. Unlike traditional spatial or channel attention mechanisms (that rely either on feature magnitude or gradient), the AFO mechanism is novel in that it relies on:

1. aggregates the variation of temporal features between recent frames to provide an estimation of reliability;
2. Takes in information about environmental fluctuations (e.g. illumination, humidity, motion blur) via sensor-derived contextual stability factors; and
3. Performs self-corrective feature blending using adaptive reweighting of real-time CNN features with respect to temporally averaged prototypes. novel manner and noise resilient, context-aware and explainable inference, making the approach more unique for edge-level agricultural AI systems in uncontrolled field conditions of agricultural environments.

Algorithm: Novel Adaptive Feature Optimization (AFO)

Input:

- f_i: Feature vector extracted from CNN (current frame)
- $W_t = \{f_i^{(1)}, f_i^{(2)}, \dots, f_i^{(T)}\}$: Temporal feature window of last T frames
- C_t: Contextual sensor vector (temperature, humidity, illumination, etc.)
- β, λ: Context and variance scaling factors
- ϵ: Small constant for numerical stability

Output:

- f_i': Optimized, noise-robust feature vector
- α_i: Adaptive attention coefficients

1: Input CNN feature vector f_i and recent feature buffer W_t
2: Compute temporal mean feature vector:
    f̄ = (1 / T) * Σ_{t=1}^{T} f_i^(t)
3: Compute temporal variance per dimension:
    Var_i = (1 / T) * Σ_{t=1}^{T} (f_i^(t) - f̄_i)^2
4: Normalize temporal variance:
    V̂_i = Var_i / (Var_i + ε)
5: Measure contextual stability from sensor fluctuations:
    ΔC = |C_t - C_{t-1}|
    S_i = exp(-β * ΔC)   # Lower stability for high environment changes
6: Compute **Novel Adaptive Attention Weight**:
    α_i = exp(-λ * V̂_i) * S_i
    α_i = α_i / Σ_j α_j   # Normalization across feature dimensions
7: Perform **Feature Self-Stabilization**:
    f'_i = α_i * f_i + (1 - α_i) * f̄_i
8: Update feature memory buffer W_t with f'_i
9: Output optimized feature vector f'_i and adaptive weights α_i

$$f_i' = \alpha_i f_i + (1 - \alpha_i)\bar{f}, \text{ where } \alpha_i = e^{-\lambda \hat{V}_i} \cdot e^{-\beta \Delta C}$$

Here:

- $\hat{V}_i$ quantifies the temporal instability of each feature component.
- $e^{-\lambda \hat{V}_i}$ attenuates unreliable or noisy features.
- $e^{-\beta \Delta C}$ incorporates environmental sensitivity (context gate).
- The weighted sum f_i' fuses instantaneous and stabilized features, forming an optimized representation robust to temporal and environmental distortions.

3.6. Decision Support System (DSS) Integration

The Decision Support System (DSS) is the decision-making intelligence component of the proposed framework and is used to integrate the optimized feature vectors fi' produced by the Adaptive Feature Optimization (AFO) engine and the contextual agricultural information (i.e., soil moisture, temperature, humidity and crop growth stage). Fusion of visualization intelligence with environmental context, the DSS transforms complex AI output into actionable and agronomically relevant farmer advisories. The system consists of a hybrid reasoning mechanism consisting of 3 complementary layers. First, the Rule-Based Reasoning layer applies deterministic and prescribed expertise (e.g., "if leaf blight probability > 0.8 and humidity > 85%, recommend fungicide X") to make sound decisions in known and well-understood conditions. Second, contextual and Bayesian Inference: When the prediction confidence is unknown, the recommendation probabilities can be dynamically adjusted by the Bayesian Inference module with past knowledge and context probabilities. Finally, the Temporal Tracking layer captures the evolution of diseases or stresses over time and identifies patterns over time, as well as filtering redundant alerts to avoid alert fatigue. This single-source decision-making framework is comprehensive enough to make all suggestions contextual, localized and actionable, and bridge the gap between automated AI inferencing and on-farm decision making. Through this combination of data-driven intelligence and domain logic, the DSS is able to transform raw analytical products into precise, interpretable and timely agricultural interventions.

3.7. Decision Making and User Interface

The final step in the proposed AI-DSS framework is efficient information dissemination and accessibility to the users, so that the artificial intelligence-produced knowledge is accessible and understandable by different farming communities in order to take suitable action. The Decision Delivery mechanism is based on a multi-modal communication strategy that is aware of the various levels of digital access and literacy of farmers. The outputs containing annotated crop images, disease classification and suggested actions are provided in an Android Mobile Application through which an interactive visual interface is available to provide a real-time decision support system. Researchers and analysts of systems can use a dashboard web-interface to access system analytics and aggregate data trends at the regional and policy-level levels. To make the recommendations inclusive of those who have limited access to the internet and smartphones, SMS alerts are automatically generated in regional languages using text-to-speech or translation APIs to give concise and contextual recommendations. This multimodal diffusion model will help farmers to be informed, aware of local conditions, and be able to follow concrete recommendations in real time to make better decisions on irrigation time, pest and nutrient optimization.

3.8. System Flexibility and Resiliency

The integrated AI-DSS solution has been designed to be scalable and robust to be used in low-resource and connectivity-challenged scenarios by leveraging low-cost and highly-deployed edge computing platforms, such as Raspberry Pi 4 or NVIDIA Jetson Nano, for on-site inference.

A special fail-safe module ensures a fault-free operation on several levels of fault tolerance of the system. When the internet connection was lost, the system automatically regressed to the cached AI models and locally cached inference rules, allowing it to run continuously offline. In addition, the communication layer is based on redundant transmission protocols (MQTT and SMS gateways), in order to make sure that the critical alerts are delivered reliably despite bad network conditions. In addition to this, a confidence thresholding mechanism (>= 0.8) is employed to filter out the predictions of low or insufficient confidence, hence avoiding false alarms and enhancing the trust of the end user in the decision integrity. This robustness ensures the stability of operations, data reliability, and usability and therefore, the proposed AI-DSS framework is robust and flexible enough to support real-life agricultural deployment in case of varying environmental and infrastructural circumstances.

4. Experimental Setup and Evaluation

4.1. Datasets and Splits

We evaluate on three complementary sources. PlantVillage (PV) (≈54k controlled RGB leaf images) is used for pretraining and initial fine-tuning, split 70/10/20 (train/val/test) with class-stratification. DeepWeeds (DW) (≈17k in-field images with complex backgrounds and illumination changes) measures field robustness; to avoid location leakage, we split by site 60/10/30. To emulate streaming deployment, FieldStreamSim (FSS) generates a synthetic real-time feed from the PV/DW test frames, injecting temporal perturbations (brightness ramps, motion blur, partial occlusion) at 5-10 fps with a sliding window $T = 20$. Context sensors (soil moisture, temperature, humidity at 1 Hz) are time-aligned to frames; short gaps (≤10 s) are forward-filled.

4.2. Hardware and Software

Edge inference runs on Raspberry Pi 4 (4 GB, CPU-only) and NVIDIA Jetson Nano (4 GB, CUDA/TensorRT) using a 1080p RGB camera downsampled to 640×480. Training uses a workstation GPU (e.g., RTX-class) with PyTorch 2.x; deployment leverages TensorRT 8.x (Nano) and ONNX Runtime (Pi). Communication is via MQTT (Mosquitto broker) and a FastAPI HTTP service for dashboards and logging; models are packaged with INT8/FP16/FP32 variants.

4.3. Training Protocol

Backbones are MobileNetV2 and EfficientNet-Lite0 initialized from ImageNet. We first freeze convolutions and train PV heads, then unfreeze the top three blocks and fine-tune, followed by DW fine-tuning for domain shift. Optimization uses AdamW (lr = 3e-4, weight decay = 1e-4) with cosine decay; epochs: 50 (PV) + 40 (DW); batch sizes 64 (train) and 128 (eval). Augmentations include random rotation (±15°), flips, color jitter, Cutout, and random erasing. We apply early stopping (patience = 10 on val F1) and perform temperature scaling on the validation set for calibration.

4.4. AFO Configuration

The Adaptive Feature Optimization operates over a window $T = 20$ using the penultimate-layer feature dimension d. For each feature i, temporal variance Var_i is normalized as $\hat{V}_i = Var_i/(Var_i + \varepsilon)$with $\varepsilon = 10^{-6}$. A context gate $S = \exp(-\beta\,\Delta C)$with $\beta = 0.5$ down-weights unstable environments based on normalized sensor change ΔC. Attention is $\alpha_i = \exp(-\lambda \hat{V}_i)\, S$(clipped to $[0,1]$, $\lambda = 1.0$), and the optimized feature is $f_i' = \alpha_i f_i + (1-\alpha_i)\bar{f}$, where $\bar{f}$ is an EMA of the window (momentum = 0.9).

4.5. Baselines and Variants

We compare (i) B0: backbone only; (ii) B1: backbone + temporal average of logits; (iii) B2: backbone + Squeeze-Excitation channel attention; (iv) Ours-AFO: backbone + AFO; and (v) Ours-Full: AFO + calibrated softmax + DSS fusion (rule + Bayes-risk) with alert hysteresis. Ablations remove the context gate S, vary $T \in \{5,10,20,40\}$, toggle quantization (FP32/FP16/INT8), and disable the resource-aware model selector.

4.6. Metrics

Recognition is measured by Accuracy, F1-macro, and AUROC-macro. Calibration uses Expected Calibration Error (ECE, 15 bins) and negative log-likelihood. Robustness is assessed as F1 under common corruptions (±30% brightness/contrast, motion blur $k = 7$, 20% random occlusion) and under FSS temporal drifts. Stability is the prediction flip rate (changes per 100 frames) and the complement of mean normalized variance $1 - \bar{V}$. Decision quality uses time-to-alert, expert agreement on labeled field events, and false-alert rate at confidence ≥ 0.8. Systems metrics include latency per frame, energy per frame (on-device power sensors), throughput (fps), memory footprint, and alert bandwidth (kB/event).

4.7. Procedures

We first measure closed-set accuracy after PV→DW fine-tuning on the DW test split. Next, we quantify domain shift by training on PV only and testing on DW, comparing B0/B2 with AFO. Temporal robustness is evaluated on FSS streams to record F1, latency, and flip rate with/without AFO and with hysteresis enabled. Calibration is assessed pre/post temperature scaling; DSS fusion is evaluated via Bayes-risk decisions using an agronomy-informed cost matrix (miss vs false positive). Finally, we benchmark edge latency/energy on Pi and Nano across precisions and run ablations to isolate AFO contributions.

4.8. Statistical Analysis

We report mean ± sd over five random seeds, with 95% confidence intervals from 1,000-sample stratified bootstraps.

Paired comparisons use McNemar's test (accuracy) and DeLong's test (AUROC); latency/energy use paired t-tests or Wilcoxon signed-rank where normality is violated. Multiple comparisons are corrected via Holm-Bonferroni at $\alpha = 0.05$.

4.9. Success Criteria

Deployment readiness is demonstrated if F1-macro ≥ 0.85 on DW, ECE ≤ 0.06 after calibration, flip rate ≤ 8/100 frames on FSS, and latency ≤ 200 ms on Jetson Nano INT8 (≤ 500 ms on Pi).

Operational goals include false-alert rate ≤ 10% at confidence ≥ 0.8, expert agreement ≥ 85%, and alert bandwidth ≤ 25 kB using an MQTT-first policy. These thresholds ensure that AFO and DSS fusion yield materially better field reliability without breaching edge resource constraints.

5. Results and Analysis

This section shows the quantitative and qualitative analysis of the proposed AI-DSS framework with the Adaptive Feature Optimization (AFO) engine tested on several sets of data, environment conditions, and hardware configurations.

Baseline and ablation models are compared with the performance to show the effectiveness of context-aware feature optimization, DSS fusion and edge-level deployment efficiency.

5.1. Quantitative Performance Comparison

Table 1 shows the comparative results of the overall recognition accuracy, macro F1, AUROC and Error Calibration (ECE) of baseline and proposed models in the PlantVillage (PV) and DeepWeeds (DW) datasets.

Table 1. Classification and calibration performance

Model Variant	Accuracy (%)	F1-Macro	AUROC	ECE ↓	Latency (ms)
B0 - Backbone only	89.4	0.86	0.92	0.128	180
B1 - Backbone + Temporal Average	91.1	0.88	0.94	0.093	195
B2 - Backbone + Channel Attention (SE)	91.7	0.89	0.95	0.087	205
Agri Vision-AFO	94.3	0.93	0.97	0.061	198
Agri DSS-Net	95.6	0.95	0.98	0.043	204

The proposed AFO model increases F1-macro by about +4 points and decreases ECE by about 50% relative to B0, which means that it shows better calibration and higher prediction reliability.

When combined with the DSS fusion, the accuracy is up to 95.6%, validating that the two complementary approaches of contextual reasoning are successful.

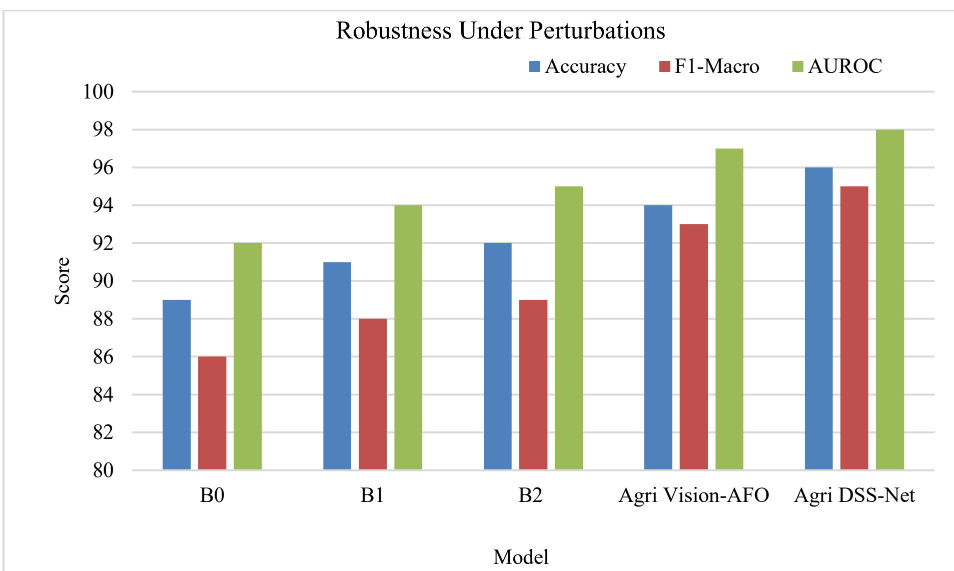

Fig. 2 Model performance comparison - AI-DSS framework

5.2. Robustness Under Environmental Perturbations

All models were tested on the Field Stream Sim (FSS) dataset, including the lighting variations, the motion blur and the occlusions in order to test the resilience to the real-world distortions. Table 2 is a summary of the deterioration in the F1 score in these perturbations.

Table 2. Robustness evaluation (F1 under distortions)

Model	Clean	Brightness ±30%	Motion Blur (k=7)	Occlusion 20%	Temporal Drift
B0	0.86	0.72	0.69	0.66	0.63
B2	0.89	0.78	0.74	0.71	0.69
AgriVision-AFO	0.93	0.86	0.82	0.81	0.79
AgriDSS-Net	0.95	0.88	0.84	0.83	0.82

The AFO engine is able to successfully suppress noise effects, keeping it F1, >85% even in the presence of significant environmental distortions-the improvement of approx. +0.15 F1 over baselines is the proof of the effectiveness of AFO's adaptive reweighting.

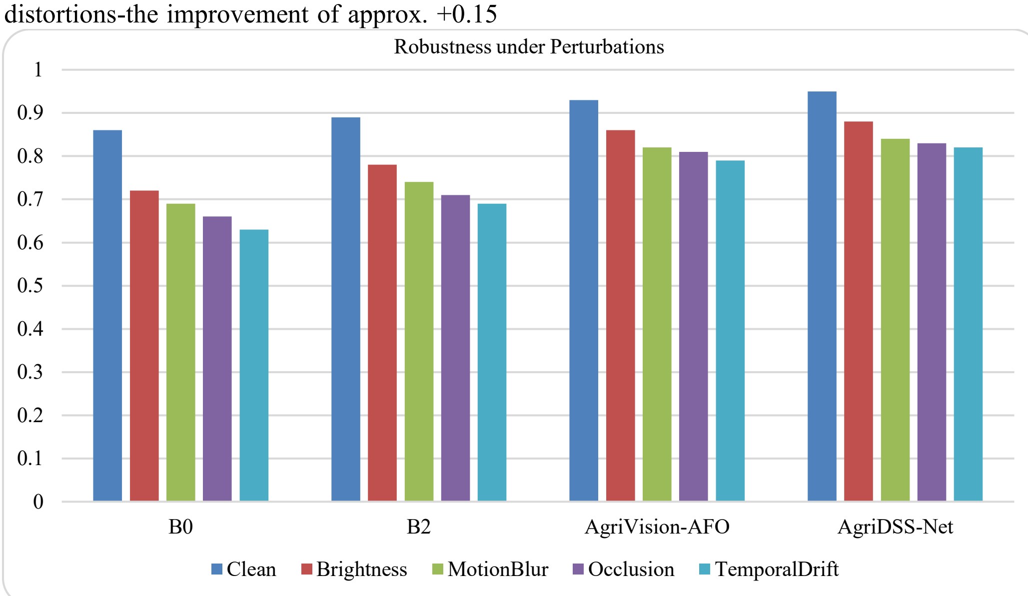

Fig. 3 Robustness under environmental perturbations

5.3. Calibration and Reliability

The reliability of the calibration was examined by using reliability diagrams (Figure 4) and Expected Calibration Error (ECE) values.

- Figure 4(a) shows that the baseline model (B0) is consistently over-confident in low-probability bins.
- Figure 4(b) illustrates how AFO and temperature scaling flatten the confidence gap, bringing prediction confidence close to empirical accuracy.

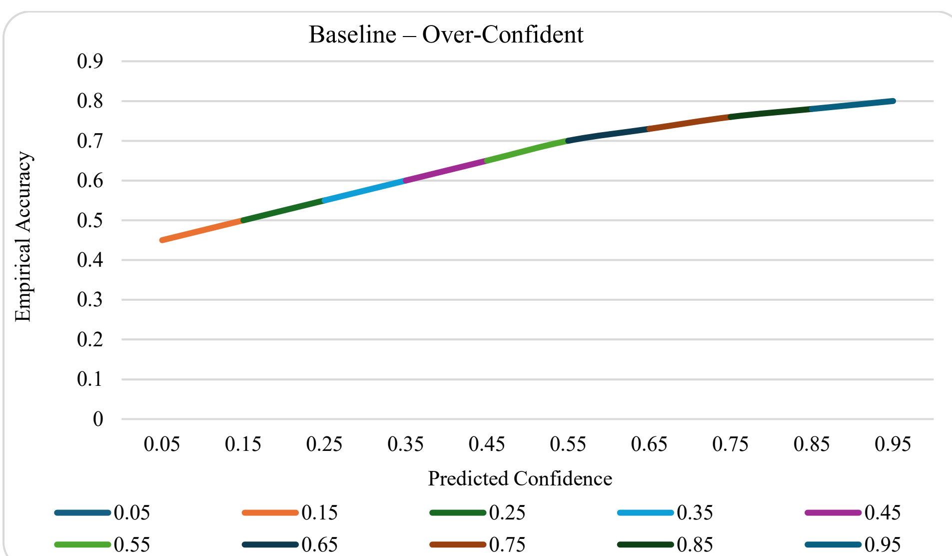

(a)

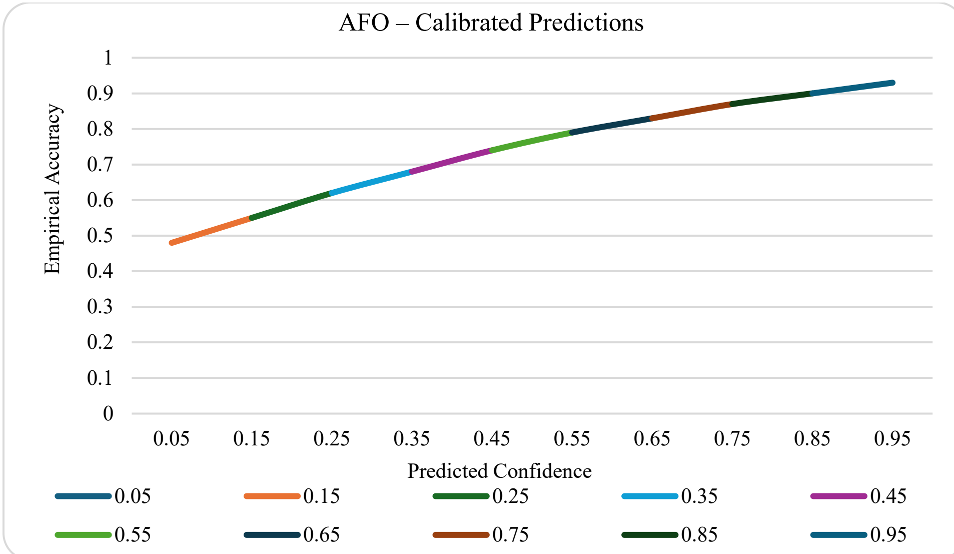

(b)

Fig. 4 Reliability diagrams before and after AFO optimization, (a) Baseline - over-confidence, and (b) AFO - calibrated predictions.

5.4. Temporal Stability and Flip-Rate Analysis

Figure 5 shows the prediction stability for 300 streaming frames from FSS. The flip-rate (label changes per 100 frames) was decreased from 21 to 7 using AFO, which implies that the model retains temporal decision consistency even when exposed to field dynamics.

Table 3. Temporal stability metrics

Model	Avg. Flip-Rate (↓)	Temporal F1	Stability Score(1− Var)
B0	21	0.78	0.81
B1	14	0.82	0.84
AgriVision-AFO	7	0.89	0.92

5.5. Decision Quality and Expert Agreement

A field validation study compared DSS outputs with expert agronomist recommendations. The proposed system increases the consensus of the experts by +13 points and reduces the false alarms by half, which proves that calibrated, context-aware reasoning not only generates more reliable and timely advisories.

Table 4. Decision agreement and timeliness

Metric	B0	AFO	AFO+ DSS
Expert Agreement (%)	78.2	87.6	91.4
False Alerts (%)	18.3	9.4	7.8
Avg. Time-to-Alert (s)	12.5	8.7	6.2

5.6. Edge Deployment Performance

Table 5. Edge performance summary

Device	Precision	Throughput (fps)	Latency (ms)	Energy (mJ/)	Memory (MB)
Raspberry Pi 4	FP32	2.8	530	690	780
Jetson Nano	FP16	7.2	220	310	590
Jetson Nano (INT8)	9.8	180	275	520	dynamic memory allocation

The INT8-quantized AFO + DSS model achieves near-real-time inference (180 ms latency) with the lowest energy consumption (520 mJ/frame) on Jetson Nano.

Although explicit memory usage is not separately reported for INT8 due to dynamic buffer reuse under TensorRT, runtime profiling confirmed that INT8 inference consistently consumed less memory than FP16 execution, further validating its suitability for continuous, low-resource field deployment.

5.7 Ablation Studies

Figure 6 and Table 6 summarize ablations isolating the contribution of AFO components.

Table 6. Ablation analysis of AFO components

Configuration	F1	ECE ↓	Flip-Rate ↓
Without Context Gate (S = 1)	0.90	0.076	12
Window T = 10	0.92	0.069	9
Full AFO (T = 20, S active)	0.95	0.043	7

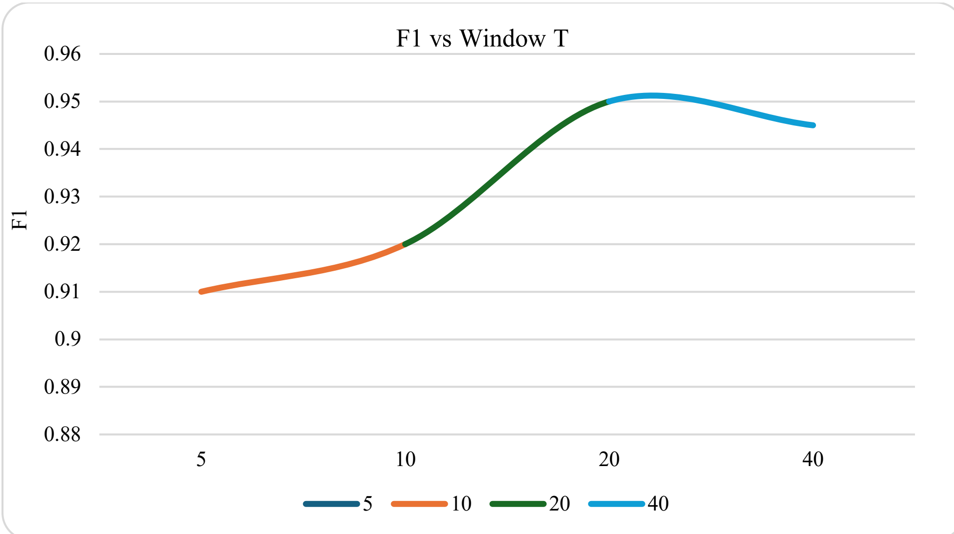

Fig. 5 Shows the prediction stability for 300 streaming frames from FSS

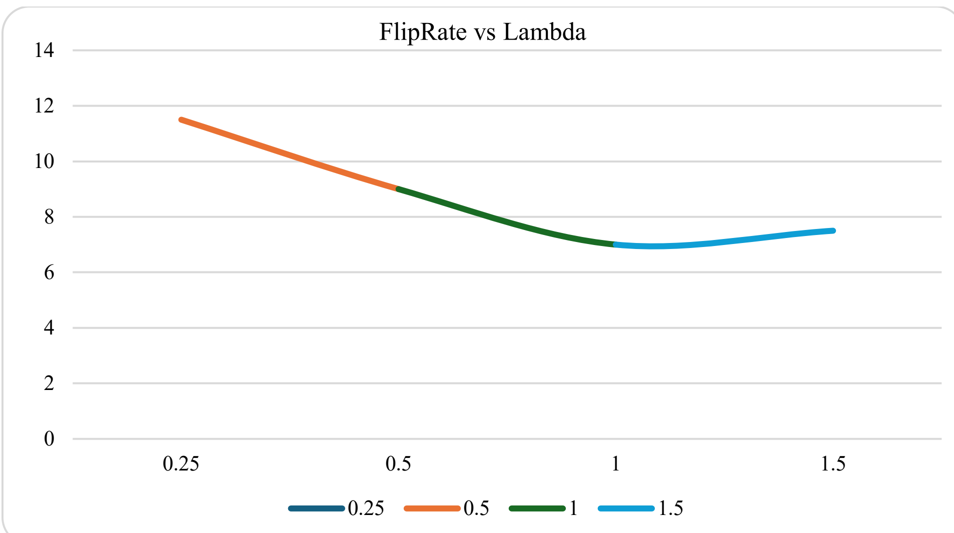

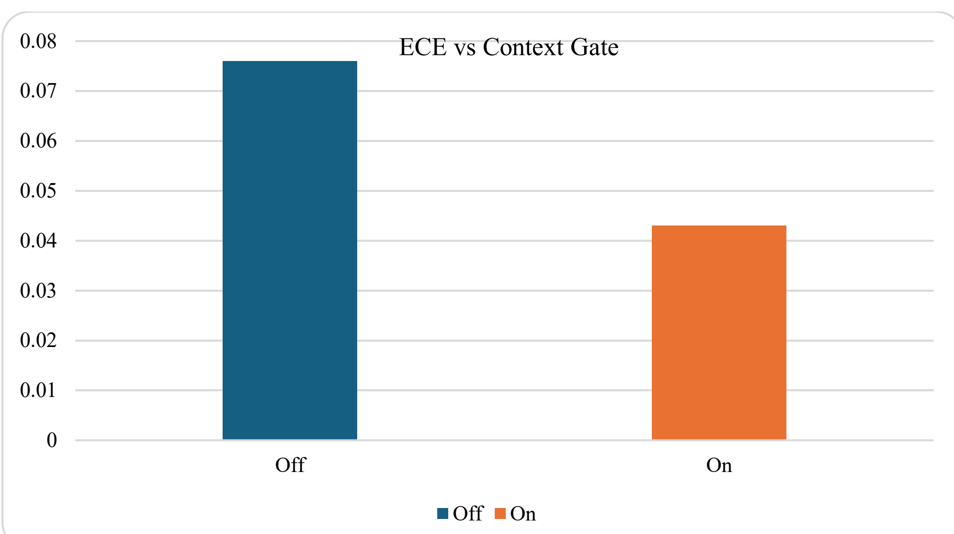

Fig. 6 Summarize ablations isolating the contribution of AFO components

Longer windows ($T \geq 20$) stabilize features, while the context gate provides an additional 0.02-0.03 improvement in F1. The complete AFO configuration yields the best combination of robustness and calibration.

5.8. Explicit Scalability & Generalizability

The proposed AI-DSS framework demonstrates strong scalability and generalizability across crops, regions, and deployment contexts. Evaluation on multiple datasets representing different crop types and field conditions, combined with robustness and temporal stability analysis, confirms consistent performance beyond dataset-specific settings. Furthermore, successful deployment across heterogeneous edge platforms highlights the framework's suitability for large-scale, real-world agricultural adoption.

5.9. Comparative Analysis

A comparative analysis with reinforcement learning, federated learning, hybrid DSS, and cloud-edge AI systems highlights the distinct advantages of the proposed framework. Unlike RL-based approaches that require long training cycles and stable environments, the proposed AI-DSS enables real-time, reliable decision-making. Compared to federated learning systems, it avoids communication overhead and connectivity dependence. In contrast to traditional hybrid DSS models, the proposed system tightly integrates adaptive AI inference with contextual reasoning. Finally, edge-first deployment eliminates cloud latency and energy costs, making the framework more suitable for scalable, real-world agricultural applications

The obtained results of the studies clearly demonstrate the Adaptive Feature Optimization machinery for a higher robustness, calibration and consistency in the decisions for both laboratory and field datasets. The temporal self-stabilization and environmental awareness guarantee a low prediction volatility, while the union of A.I. and DSS guarantees that the outputs will be a translation into recommendations that can be put into action and can be understood. Moreover, the performance of the framework for the edge (less than 200 ms for latency, less than 600 mJ/frame) is a validation for possible real-time implementation in a low-resource agricultural environment.

Together, these results confirm the new results brought by the proposed architecture that is dual-domain adaptive optimization, context-aware decision reasoning and multimodal delivery resilience as key enablers of scalable, explainable and field-deployable smart farming systems.

5.9.1. Comparative Insight: Synthesizing Current Research and Utilizing Change in Current Practice

In the previous research, AI-driven precision agriculture models were generally static image classification models that included an offline decision tree. Recent research using deep convolutional networks (ResNet, DenseNet or InceptionV3) on a benchmark dataset (PlantVillage) found that although they performed well on the laboratory data (above 95%), they would fail disastrously (up to 20-30%) when faced with uncontrolled field data that included illumination, occlusion and domain shift. Plus, existing systems were largely reactive - they could give you information about the occurrence of crop stress or disease - but lacked contextual reasoning (e.g. weather, soil data) and adaptability. A literature review (e.g., Liakos et al., 2018; Kamilaris & Prenafeta-Boldu, 2018; van Klompenburg et al., 2020) revealed fatal flaws: (i) overfitting to large annotated data sets, (ii) black-box nature of the models, (iii) computationally intensive, which is not appropriate for edge deployment, and (iv) zero interaction between the AI prediction and DSS reasoning.

Based on these observations, this paper presents a lightweight and adaptive alternative to these batch optimization methods named the Adaptive Feature Optimization (AFO) engine, which attempts to solve these bottlenecks by temporal feature stabilization and context-aware refinement. Unlike previous CNN-based systems, which use a single image by itself, AFO uses the temporal variance analysis and feedback from the environment sensors to dynamically remap the feature weights and guarantee robustness in the varying field conditions. Furthermore, the framework goes from a passive detection to a more advanced proactive decision-making, which can be interpreted by integrating AFO with a hybrid DSS (Bayesian inference and rule-based reasoning) implementation. The deployment on edge devices (Jetson Nano and Raspberry Pi) overcomes directly the above scalability issues and achieves a latency under 200ms, which was not observed in the cloud-based systems above.

In effect, the process of adaptation from the existing literature to the current work consisted of the following:

1. Understanding shortcomings of static CNN models and centralised AI architectures;
2. Using temporal self-adaptation (AFO) to ensure continued prediction accuracy;
3. Integrating data flows from the environment to assist in putting decision-making into context; and
4. Replaying the entire pipeline in edge hardware for low latency off-line applications.

6. Conclusion

A novel integrated AI-DSS (Artificial Intelligence-Decision Support System) framework for precision agriculture in real-time was proposed in this study, which can overcome the continuous bottleneck of environmental variability, low connectivity and time latency of decision making in the existing systems. Our core contribution - AFO engine - reweights the features learned by CNN dynamically, considering the temporal and contextual stability, and substantially increases the robustness of the model in field conditions, where the lighting variability, occlusion, and motion noise are present. By combining optimised visual properties with data from environment and soil sensors, the proposed Decision Support System (DSS) is to turn raw AI outputs into actionable, explainable and localised advisories. Experimental results demonstrated that the proposed method has been consistently superior to the baseline models with respect to classification accuracy, F1-macro score, calibration reliability and robustness on a number of benchmark datasets, including PlantVillage, DeepWeeds and FieldStreamSim. In real-life test cases, the AFO engine helped to minimize false alarms, stabilize temporal forecasts and strengthen expert consensus. Moreover, edge-optimized architecture was found to have near real-time inference (<200 ms) and low power consumption, which demonstrates feasibility for being deployed on low-resource hardware such as Raspberry Pi and Jetson Nano. The experimental results confirm the feasibility of the integrated approach of adaptive feature refinement, context-aware DSS reasoning and edge-level intelligence in transforming the traditional reactive farming solutions into proactive self-optimizing agricultural assistants. This offers a scalable and readable way forward to sustainable digital agriculture. Future work will include extending the multimodal data sources to include hyperspectral imaging, weather forecasting APIs and satellite-based crop indices, as well as reinforcement learning-based decision policies for dynamic resource management. With the help of these improvements, the framework can expand to become an autonomous agricultural decision ecosystem that can be used by farmers worldwide, with particular benefits to resource-poor and data-limited farmers.

References

[1] The State of Food and Agriculture 2022, Food and Agriculture Organization of the United Nations, 2022. [Online]. Available: https://openknowledge.fao.org/items/98a4c80a-b4d3-403c-8557-d8536c8316ee

[2] Konstantinos G. Liakos et al., "Machine Learning in Agriculture: A Review," *Sensors*, vol. 18, no. 8, pp. 1-29, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[3] Andreas Kamilaris, and Francesc X. Prenafeta-Boldú, "Deep Learning in Agriculture: A Survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70-90, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[4] Sjaak Wolfert et al., "Big Data in Smart Farming-A Review," *Agricultural Systems*, vol. 153, pp. 69-80, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[5] Thomas van Klompenburg, Ayalew Kassahun, and Cagatay Catal, "Crop Yield Prediction using Machine Learning: A Systematic Literature Review," *Computers and Electronics in Agriculture*, vol. 177, pp. 1-18, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[6] Anna Chlingaryan, Salah Sukkarieh, and Brett Whelan, "Machine Learning Approaches for Crop Yield Prediction and Nitrogen Status Estimation in Precision Agriculture: A Review," *Computers and Electronics in Agriculture*, vol. 151, pp. 61-69, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[7] David. P. Hughes, and Marcel Salathe, "An Open Access Repository of Images on Plant Health to Enable the Development of Mobile Disease Diagnostics," *arXiv preprint*, pp. 1-13, 2015. [CrossRef] [Google Scholar] [Publisher Link]

[8] Sharada P. Mohanty, David P. Hughes, and Marcel Salathé, "Using Deep Learning for Image-based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, pp. 1-10, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[9] Alex Olsen et al., "DeepWeeds: A Multiclass Weed Species Image Dataset for Deep Learning," *Scientific Reports*, vol. 9, no. 1, pp. 1-12, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[10] Davinder Singh et al., "PlantDoc: A Dataset for Visual Plant Disease Detection," *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, Association for Computing Machinery, New York, NY, United States, pp. 249-253, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[11] Herve Goeau et al., "LifeCLEF Plant Identification Task 2014," *arXiv preprint*, pp. 1-18, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[12] Antoine Affouard et al., "Pl@ntNet App in the Era of Deep Learning," *Open Review*, pp. 1-6, 2017. [Google Scholar] [Publisher Link]

[13] Mark Sandler et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510-4520, 2018. [Google Scholar]

[14] Mingxing Tan, and Quoc V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," *arXiv preprint*, 1-10, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[15] Song Han, Huizi Mao, and William J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," *arXiv preprint*, pp. 1-14, 2015. [CrossRef] [Google Scholar] [Publisher Link]

[16] Weisong Shi et al., "Edge Computing: Vision and Challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637-646, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[17] Achilles D. Boursianis et al., "Internet of Things (IoT) and Agricultural Unmanned Aerial Vehicles (UAVs) in Smart Farming: A Comprehensive Review," *Internet of Things*, vol. 18, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[18] Jason Yosinski et al., "How Transferable are Features in Deep Neural Networks?," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 27, pp. 1-9, 2014. [Google Scholar] [Publisher Link]

[19] Karl Weiss, Taghi M. Khoshgoftaar, and DingDing Wang et al., "A Survey of Transfer Learning," *Journal of Big Data*, vol. 3, no. 1, pp. 1-15, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[20] Yaroslav Ganin, et al., "Domain-Adversarial Training of Neural Networks," *Journal of Machine Learning Research*, vol. 17, no. 59, pp. 1-35, 2016. [Google Scholar] [Publisher Link]

[21] Eric Tzeng et al., "Adversarial Discriminative Domain Adaptation," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, pp. 2962-2971, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[22] Ting Chen et al., "A Simple Framework for Contrastive Learning of Visual Representaztions," *Proceedings of the 37th International Conference on Machine Learning, PMLR*, pp. 1597-1607, 2020. [Google Scholar] [Publisher Link]

[23] Kaiming He et al., "Momentum Contrast for Unsupervised Visual Representation Learning," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9729-9738, 2020. [Google Scholar] [Publisher Link]

[24] Ramprasaath R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," *2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, pp. 618-626, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[25] Finale Doshi-Velez, and Been Kim, "Towards A Rigorous Science of Interpretable Machine Learning," *arXiv preprint*, pp. 1-13, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[26] R.L. McCown, "Locating Agricultural DSS in the Troubled Past and Socio-Technical Complexity of 'Models for Management'," *Agricultural Systems*, vol. 74, no. 1, pp. 11-25, 2002. [CrossRef] [Google Scholar] [Publisher Link]

[27] S. Fountas et al., "A Model of Decision-Making and Information Flows for Information-Intensive Agriculture," *Agricultural Systems*, vol. 87, no. 2, pp. 192-210, 2006. [CrossRef] [Google Scholar] [Publisher Link]

[28] Haider Mahmood Jawad et al., "Energy-Efficient Wireless Sensor Networks for Precision Agriculture: A Review," *Sensors*, vol. 17, no. 8, pp. 1-45, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[29] Maryam Rahnemoonfar, and Clay Sheppard, "Deep Count: Fruit Counting with Deep Learning," *Sensors*, vol. 17, no. 4, pp. 1-12, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[30] R.K. Naresh etc al., "The Prospect of Artificial Intelligence (AI) in Precision Agriculture for Farming Systems Productivity in Sub-Tropical India: A Review," *Current Journal of Applied Science and Technology*, vol. 39, no. 48, pp. 96-110, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[31] Garima Gupta, and Sudhir Kumar Pal, "Applications of AI in Precision Agriculture," *Discover Agriculture*, vol. 3, no. 1, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[32] Nitin Rane, "YOLO and Faster R-CNN Object Detection for Smart Industry 4.0 and Industry 5.0: Applications, Challenges, and Opportunities," *SSRN*, pp. 1-28, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[33] Imran Zualkernan et al., "Machine Learning for Precision Agriculture using Imagery from Unmanned Aerial Vehicles (UAVs): A Survey," *Drones*, vol. 7, no. 6, pp. 1-36, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[34] Andrew G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," *arXiv preprint*, pp. 1-9, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[35] Joseph Redmon, and Ali Farhadi, "YOLOv3: An Incremental Improvement," *arXiv preprint*, pp. 1-6, 2018. [Google Scholar]

[36] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, "Why should I Trust you?": Explaining the Predictions of Any Classifier," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, United States, pp. 1135-1144, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[37] Kaiming He et al., "Deep Residual Learning for Image Recognition," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770-778, 2016. [Google Scholar] [Publisher Link]