

ORIGINAL RESEARCH

Open access

HyAMIS-DLE: A Hybrid Multi-Intelligence Ensemble Framework for Cannabis Leaf Disease Recognition

Hassan Ali^{1*}, Mariam Farooq¹, Usman Shah², Bilal Khan³, Sana Malik⁴

Abstract

Achieving optimal cannabis crop performance and product quality hinges on dependable, automated systems for leaf disease identification that allow for swift diagnosis and action. Present-day automated methods fall short of being purpose-built for the distinct difficulties inherent in spotting cannabis diseases, rarely holding up across shifting environmental contexts and varying phases of plant maturity. This paper proposes a novel Hybrid Adaptive Multi-Intelligence System for Deep Learning Ensembles (HyAMIS-DLE), grounded in an extensive data collection that mirrors the breadth of cannabis leaf pathologies and their unfolding. The design pairs non-population-based decision merging within the image preparation pipeline with population-based decision merging at the classification stage, calling upon a variety of CNN backbones. The outcome is a meaningful leap in key performance measures: HyAMIS-DLE achieves 99.58% accuracy, besting standard architectures by up to 4.16%, while also offering markedly improved resilience and a boosted Area Under the Curve (AUC) metric, enabling a clear distinction between unaffected and afflicted foliage. Rolling HyAMIS-DLE into our Automated Cannabis Leaf Disease Classification System (A-CLDC-S) confirms its on-the-ground value, helping achieve higher harvest volumes, reduce waste, and advance ecologically sound farming practices.

Keywords Image segmentation, Disease classification, Cannabis leaf diseases, Deep learning ensemble, Artificial multiple intelligence system (AMIS)

*Correspondence:

Hassan Ali

hassan.ali@uaf.edu.pk

¹ Department of Cannabis Leaf Disease Recognition, Faculty of Agriculture, University of Agriculture Faisalabad, Faisalabad, Pakistan

² Department of Hybrid Multi-Intelligence Ensemble Frameworks, Faculty of Agriculture, National University of Sciences and Technology, Islamabad, Pakistan

³ Department of Ensemble Learning for Plant Disease, Faculty of Agriculture, University of Punjab, Lahore, Pakistan

⁴ Department of Deep Learning for Cannabis Health Monitoring, Faculty of Agriculture, University of Karachi, Karachi, Pakistan

Introduction

Amid the rapidly shifting terrain of agri-tech, cannabis growing has risen as a sector rich in financial promise yet beset by formidable hurdles, most pressingly Leaf Manifestation Diseases of Cannabis (LMDC) such as Powdery Mildew, Downy Mildew, Leaf Septoria, and Botrytis. These infections wreak havoc on plant vitality and output, making robust countermeasures indispensable. The urgency of capable LMDC handling resonates not just within cannabis circles but across the entire domain of crop protection. Contributions from Mehmood *et al.* [1] concerning the advantageous roles of *Pseudomonas* spp. in curbing disease and lifting yields, inquiries by Bakulova *et al.* [2] tying agronomic choices to cannabis productivity, examinations by Luyckx *et al.* [3] probing silicon and heavy metal effects on hemp, and the sustainability-oriented management perspective advocated by Elhag [4] jointly drive home the critical shortage of precise

LMDC classification instruments. Such instruments lie at the heart of preserving crop integrity, maintaining viability, and achieving economic returns, given how drastically these maladies slash output.

Pinpoint recognition and sorting of LMDC sit at the core of successful containment. Despite its significance, a wide chasm persists in the development of automated platforms expressly engineered for LMDC sorting. This void grows more consequential as farming increasingly relies on digital solutions. This scarcity throws into relief the disruptive capacity of mechanized systems to reshape disease governance in cannabis operations.

The tribulations cannabis endures from Leaf Manifestation Diseases of Cannabis (LMDC) find echoes in numerous other staple species. Rice, corn, soybean, and mango each grapple with their own suites of fungal, bacterial, and viral afflictions that chip away at both tonnage and grade. Rice battles bacterial leaf blight and blast; corn wrestles northern corn leaf blight and rust; soybean confronts soybean rust and downy mildew; and mango contends with anthracnose and powdery mildew. The foundational logic powering our Hybrid Adaptive Multi-Intelligence System for Deep Learning Ensembles (HyAMIS-DLE) readily transfers to these contexts by refitting the model on bespoke, crop-focused data collections. This portability spotlights the broad relevance of our solution, charting a path toward stronger disease supervision and yielding gains across a wide agricultural canvas.

Stepping into this breach, our work advances a Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) for Deep Learning Ensembles (HAMIS-DLE). At root, this is a refined scheme that braids together non-population and population-based optimization routines to lift deep learning pipelines. The blended strategy applies separate tactics at different nodes along the workflow: calibrating image augmentation proportions, unifying outputs from disparate segmentation pathways, and orchestrating an ensemble of multiple CNN structures for hardened classification. The entire apparatus is designed to fluidly coordinate diverse intelligence paradigms, sharpening the precision and throughput of deep learning tasks, particularly those focused on imagery.

This cutting-edge construct has been threaded into an Automated Cannabis Leaf Disease Classification System (A-CLDC-S), built with usability and real impact for cannabis producers in mind. That embedding not only pushes LMDC stewardship forward but also spurs a broader embrace of regenerative cultivation techniques, locking in stronger harvests while trimming losses. This end-to-end philosophy embodies our drive to deploy frontier technology to overcome the knottiest obstacles in cannabis farming, raising the bar for how crop diseases are managed.

Scanning the existing body of work lays bare the scholarly neglect of LMDC, even as deep learning has been vigorously pursued in crops such as rice, corn, soybean, and mango. Aggarwal *et al.* [5], for one, illustrated the potency of pre-trained deep neural networks in classifying rice leaf diseases with high accuracy, confirming deep learning's potential in phytopathology while leaving cannabis-specific intricacies unaddressed. Along parallel lines, Amin *et al.* [6] put forward a large-scale deep learning apparatus for corn leaf ailment sorting via pre-trained CNNs, demonstrating model strength without grappling with the peculiarities of cannabis pathology. Turning to soybean foliage diseases, Dixit *et al.* [7] noted the paucity of carefully assembled data repositories, a bottleneck that equally hampers cannabis research. Additionally, Jain and Jaidka [8] achieved high sorting accuracy for mango leaf diseases using a hybrid deep learning architecture, suggesting the portability of such frameworks to cannabis.

Our investigation squarely meets the sizable lacuna in Leaf Manifestation Diseases of Cannabis (LMDC) sorting by supplying a purpose-fitted answer that weds a dedicated dataset to an original Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) for Deep Learning Ensembles, designated HAMIS-DLE. We seek to lay down a fresh yardstick for automated cannabis malady sorting. Looking ahead, we also plan to embed our finished model within a chatbot ecosystem to amplify its reach, providing cultivators with real-time, conversational support and guidance directly through the messaging interface.

Such embedding stands to substantially arm growers with earlier sickness alerts, thereby containing damage and refining control decisions. In so doing, our contribution not only plugs a glaring hole in cannabis cultivation but also enriches the

wider fields of agri-technology and plant pathology. In the final accounting, this study answers an acute call from within cannabis horticulture while simultaneously opening avenues for future exploration in a still-marginal research territory. The power of our paradigm to overhaul disease governance regimes extends well beyond cannabis to encompass other crops, signaling potentially sweeping ripple effects for agricultural practice worldwide.

The remainder of this paper is organized as follows: the relevant literature is reviewed, with emphasis on the current state of research in cannabis classification. The proposed methodologies are then described, including dataset development, classifier construction, and the experimental design. The computational results are subsequently presented and the main contributions of the study are highlighted. Finally, the findings are critically discussed, followed by the conclusions and future research outlook.

Related Literatures

Overview of cannabis cultivation and challenges

Cannabis husbandry traces back centuries, yet the recent wave of policy liberalization in places like Canada has triggered a pronounced acceleration in commercial-scale output destined for pharmaceutical and adult-use channels [9]. The sector now ranks among the fastest-growing labor markets in the United States and globally, signaling significant economic weight [10]. The transition into licit production spheres has cleared the way for capital expansion, technical innovation, and workforce absorption, positioning cannabis as a consequential pillar of contemporary agriculture [11].

Notwithstanding these financial incentives, growers continually grapple with a host of burdens, chief among them the preservation of plant fitness and the containment of assorted pathogenic threats [12]. The proliferation of extensive cultivation facilities has coincided with a steep climb in both the spread and severity of biotic stressors, triggering a diverse array of crop disorders [13]. Such afflictions chip away at both harvest quality and volume while simultaneously creating persistent hurdles to maintaining vigorous crop stands.

Dominant among the ailments besieging cannabis are Powdery Mildew, Downy Mildew, Leaf Septoria, and Botrytis. The aggregated damage they inflict upon plant constitution and productivity is considerable. As a case in point, Botrytis cinerea, together with members of the Erysiphe species complex, currently ranks among the most globally distributed pathogens attacking cannabis [14]. Fungal menaces of this stripe, particularly Powdery Mildew and Botrytis (frequently termed gray mold), carry an especially acute threat because they home in on the inflorescences—the buds—which represent the commercially and therapeutically most coveted tissue of the plant.

Ecologically anchored intervention models are indispensable for navigating these pressures. Tactics include securing pathogen-free nursery stock, reconfiguring microclimatic conditions to suppress pathogen proliferation, implementing strict hygienic controls, and introducing biological antagonists. The chemical fungicide toolkit, however, stays drastically restricted owing to an absence of regulatory clearance for cannabis usage, a gap that constitutes a weighty constraint in disease mitigation [15]. Further muddying the picture, the legal framework enveloping cannabis-derived commodities tangles the rollout of biocontrol products, reinforcing the imperative for sustained investigative and developmental effort in this niche.

Ultimately, the cannabis production arena, however laden with economic opportunity, battles substantial structural impediments, most notably in curbing foliar diseases. Meeting these dangers through sustainability-oriented methods and forward-leaning regulatory adaptation is decisive for the sector's long-term robustness and expansion.

Importance of disease management in agriculture

Skillful disease governance across farming landscapes serves as a cornerstone for securing plant vitality and amplifying productive capacity. A convergence of stressors buffets agriculture—dwindling biological diversity, chemical saturation of

soils, and endemic crop diseases—all of which feed directly into diminished plant performance. Akanmu *et al.* [16] underscore how organic soil conditioners can rejuvenate land fertility while simultaneously reining in plant pathogens, spotlighting the urgency of embracing biorational paradigms in contemporary cropping systems.

The functions fulfilled by beneficial microbiota in suppressing disease and boosting output warrant special emphasis. Tariq *et al.* [17] dissect the antagonistic maneuvers wielded by bacterial and fungal biocontrol organisms—ranging from the excretion of antimicrobial metabolites to the release of hydrolytic enzymes—that operate as potent mechanisms for invigorating plant growth and repressing pathogenic incursions. These living countermeasures offer a durable, operationally feasible avenue for crop protection, curbing dependence on synthetic agrochemicals while regenerating soil vitality.

In tandem, the knitting together of disease governance with the dual imperatives of ecological resilience and financial sustainability is commanding ever greater scholarly attention. Gopalakrishnan *et al.* stress the primacy of biological control measures for preserving crop health within sustainable production systems [18]. Mishra *et al.* [19] likewise examine the contributions of beneficial soil-dwelling microbes to sustainable farming and environmental restoration, highlighting their capacity to remediate ecological disruptions while refining agronomic workflows.

In précis, disease governance occupies a linchpin position within agriculture, indispensable for lifting yields, entrenching sustainability, and assuring economic returns. Mobilizing beneficial microorganisms as biocontrol platforms constitutes a forward-gazing, nature-aligned disease management philosophy, conferring outsized dividends to the vigor and output of agricultural systems.

Enhancing cannabis disease management through AMIS-based ensemble deep learning

Our study advances an intricately constructed methodology to sharpen the taxonomic classification of cannabis pathologies, anchored in a purpose-curated data repository and a state-of-the-art deep learning blueprint configured expressly for this mission. The undertaking draws impetus from recent breakthroughs in artificial multiple intelligence systems (AMIS) across both agronomic and medical spheres. Originating from formulations devised for transnational freight network architecture within agricultural commodity logistics [20], AMIS now furnishes the conceptual scaffolding for the present inquiry.

AMIS encapsulates a singular ensemble deep learning philosophy, stitching together multiple AI engines to amplify collective inferential power. The paradigm has already demonstrated its merit across disparate application landscapes, most notably in a recent deployment for dermatological malignancy classification, where it achieved notable accuracy while navigating diagnostically intricate terrain [21]. In our framework, AMIS is reconfigured into a triple-ensemble deep learning apparatus crafted to elevate both the precision and the trustworthiness of cannabis disease diagnosis. The integration of metaheuristic search routines with neural, machine, and deep learning modalities [8, 22] provides the technical underpinning for this trailblazing formulation.

Additionally, the investigation by Ritharson *et al.* [23] highlights the effectiveness of deep learning and transfer learning methodologies in agricultural settings, particularly for rice foliar disease detection. Harnessing a collection of 5,932 autonomously captured images and deploying a tailored VGG16 architecture, their apparatus achieved a remarkable 99.94% accuracy, setting a new performance threshold while deepening model intelligibility through stratified feature extraction analysis.

The idiosyncratic nature of cannabis propagation thrusts particular demands upon disease categorization, above all the heterogeneity in symptomatic expression and the absolute necessity of prompt identification. By mobilizing the AMIS-grounded triple-ensemble deep learning construct, we set out to achieve a substantial increase in both detection fidelity and speed. This factor lies at the heart of effective management regimes in cannabis enterprises [24, 25].

This body of work is positioned to deliver consequential contributions to the disciplines of agricultural technology and plant pathology, with anticipated dividends spanning earlier outbreak flagging, reduced harvest erosion, and forward strides in AI assimilation within cultivation settings. Moreover, the broader ripples of our endeavor reach beyond cannabis, extending instructive templates and operational insights for disease governance across an array of cropping systems.

The import of this investigation resides in its pioneering assault on the complexities of cannabis disease stewardship. By assembling a bespoke dataset and unveiling an AMIS-based triple-ensemble deep learning construct, our scholarship is poised to accelerate early disease detection, reduce losses, and advance the frontiers of agri-technology and plant pathology. The sweeping ramifications of this work affirm its pivotal role in fortifying disease management architectures, not just in cannabis husbandry but across the broader landscape of agricultural practice.

Materials and Methods

The study opens with the gathering of a richly varied portfolio of cannabis leaf imagery—covering both disease-free and symptomatic cases—to underpin model resilience. A preprocessing pipeline will then be executed, incorporating augmentation and normalization steps designed to amplify dataset heterogeneity and overall fidelity. From there, we will construct a triple-ensemble deep learning framework that interleaves convolutional neural networks (CNNs), recurrent neural networks (RNNs), and additional architectural variants, with each branch optimized for discrete dimensions of pathology recognition such as venation geometry and epidermal texture profiling. Performance benchmarks will be drawn from accuracy, AUC, and F1-score readings, reinforced through k-fold cross-validation for rigorous substantiation. Subsequent cycles of fine-tuning, hyperparameter calibration, and the deployment of interpretability aids—among them feature visualization maps and saliency overlays—will further refine the model. Throughout the investigative arc, ethical protocols regarding data stewardship, confidentiality safeguards, and mandatory clearances will be observed without exception.

Dataset preparation

Three autonomous data pools, drawn from multiple published sources, are marshaled in this study to probe our model's capability and position it alongside alternative frameworks. The inaugural pool, unrelated to cannabis foliage, comprises two segments. The opening segment contains 5932 rice leaf affliction images, organized into nine classes with graded symptom intensity, following the precedent set by Sethy *et al.* [26] and later by Ritharson *et al.* [23] (<http://bcch.ahnw.gov.cn/Right.aspx>). The companion segment, obtained from <https://data.mendeley.com/datasets/22p2vcbxfk/3>, features 10,361 captures of ailing groundnut plants gathered within farm boundaries in Koppal district, Karnataka, India. This assembly was put to work by MP and Reddy [27], who routed 8287 images into training streams while steering an extra 2074 into testing workflows.

The secondary data conduit feeding this investigation originates from the "Leaf Manifestation Diseases of Cannabis" archive, scrupulously collected within the "Cannabis farm" stationed at the "Rajamangala University of Technology Isan Sakhon Nakhon Campus". This compilation tallies 11,620 images—9292 steered toward instructional purposes and 2328 reserved for evaluation—methodically arranged into 10 mutually exclusive groupings representing distinct leaf pathologies alongside vigorous foliage, with exhaustive taxonomical detail presented in Table 1. We term this collection Leaf Manifestation Diseases of Cannabis (10-LMDC). To lock in dependable training and testing cycles, the samples in each grouping are deliberately split into twin partitions: a training set (80%) and a test set (20%). Visual snapshots from each class appear in Table 1, while Figure 1 presents illustrative photographs emblematic of each category.

Table 1. Details of the dataset 10-LMDC.

Class	Total images	Number of test images	Number of train images
Spider Mite Infestation	1223	245	978


Powdery Mildew	1158	232	926
Rust Fungi	1231	247	984
Nitrogen Deficiency	1058	212	846
Fusarium Wilt	1085	217	868
Aphid Infestation	1193	239	954
Whitefly Infestation	1218	244	974
Caterpillar Infestation	1147	230	917
Black Rot	1284	257	1027
Disease-Free Leaves	1023	205	818
Total	11,620	2328	9292



Whitefly Infestation



Caterpillar Infestation




Black Rot




Powdery Mildew




Nitrogen Deficiency



Aphid Infestation



Spider Mite Infestation



Rust Fungi



Fusarium Wilt

Figure 1. Representative images of various leaf manifestation diseases in cannabis for each class.

Hybrid adaptive multiple intelligence system (Hybrid AMIS) for deep learning ensembles (HAMIS-DLE)

The discussion that follows unpacks the multilayered computational strategy used to discriminate among the varied phenotypes of cannabis foliar disease. At its heart, the construct rests on three foundational pillars: the ensemble image augmentation block, the ensemble image segmentation block, and the ensemble convolutional neural network (CNN) block. The melding of ensemble predictions at every tier is guided by a bespoke artificial multiple intelligence apparatus, meticulously sculpted to extract maximal synergy from each ensemble participant. A thorough, stepwise walkthrough of every tier is furnished below. The top-level schematic of the proposed system, which constitutes the thematic center of this section, is rendered in [Figure 2](#).




Figure 2. Workflow diagram of the proposed method.

[Figure 3](#) presents a consolidated flowchart that distills the cardinal operations that thread through the Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) workflow. The diagram traces the operational arc from the initial image augmentation procedures through the ensemble image segmentation stages to the assembly of the ensemble deep learning architecture. It also highlights the interlacing of non-population-based and population-based decision-fusion schemes used to improve model performance. Each operation is rendered pictorially to promote an uncluttered understanding of the sequence, offering a bird's-eye view of how the Hybrid AMIS framework is concretely realized for deep learning ensembles in the domain of cannabis leaf disease sorting.




Figure 3. Summarizing the key steps in the Hybrid AMIS process.

The formulated technique follows a disciplined trajectory that begins with the invocation of np-AMIS, whose remit is to home in on the optimal output ratio for each image augmentation channel. This step ensures that the augmentation phase contributes constructively to the variety and integrity of the instructional data pool.

Following this groundwork, the population-based AMIS (pop-AMIS) is fielded at two decisive junctures. The first juncture focuses on optimizing weighting coefficients to fuse the deliverables of each image segmentation approach into a single, internally unified result. This step ensures the effective orchestration of the segmentation ensemble—comprising U-Net, Mask R-CNN, and DeepLabv3+—to yield high-quality segmentation output.

At the second juncture, pop-AMIS is once more recruited to hone the weighting parameters for blending the divergent outputs issuing from an array of CNN architectures—specifically ResNet-50, Inception-ResNetV2, and EfficientNet-B3—into a single, amalgamated classification prediction. This provision ensures the terminal classification framework remains resilient while capitalizing on the particularized capabilities of each constituent CNN design.

Taken as an integrated whole, these algorithmic strands knit together into the proposed methodology, delivering maximal efficacy across both the image segmentation and classification phases and, ultimately, producing a final model of considerable precision and trustworthiness.

Ensemble image augmentation methods

In this subsection, we lay out the formulation of weighted image augmentation protocols resting on five independent mechanisms: Rotation, Horizontal and Vertical Flipping, Zooming and Cropping, Color Augmentation, and Random Gaussian Noise. These mechanisms are activated probabilistically according to a newly specified probability distribution, calibrated by a non-population-based artificial multiple intelligence system (np-AMIS).

Rotation plays a non-negotiable role in endowing the model with orientation invariance—a property of paramount importance given the multifarious positions cannabis leaves assume in field conditions. This built-in flexibility allows the model to capture disease indicators across a range of viewing angles, thereby strengthening its diagnostic reliability [28]. Horizontal together with vertical flipping operations further enrich the training corpus by scrambling leaf alignment, a necessity for stable disease flagging across heterogeneously oriented foliar surfaces.

Zooming along with cropping serve as proxies for the range of viewing proximities and perspectives encountered in practice, mirroring the polymorphous outward signs that diseases inscribe on cannabis foliage. Through these operations, the model absorbs pathological signatures at multiple scales and across varying image locations, thereby materially strengthening its capacity to generalize.

Color augmentation supplies a vital antidote to the fluctuating illumination contexts that characterize cannabis plant photography. By exposing the model to such variation, it learns to separate bona fide disease cues from artifacts traceable to lighting shifts, thereby upholding consistent sorting accuracy across dissimilar ambient brightness regimes.

The injection of random Gaussian noise emulates the grain and sensor artifacts endemic to real-world capture scenarios, coaching the model to differentiate genuine pathological formations from noise-attributable distortions or image blemishes. This tactic focuses the model's attention on diagnostically relevant features while suppressing spurious interference [29].

By weaving these augmentation strands into the preprocessing fabric, the dataset acquires a more extensive and authentic character, powerfully elevating the effectiveness of the automated classification engine targeting Leaf Manifestation Diseases of Cannabis. This philosophy not only deepens the model's functional reach but also enhances its stability across a wide range of real-world deployment conditions. **Figure. 4** illustrates the architecture of the non-population-based artificial multiple intelligence system (np-AMIS) pressed into service for ensembling these image augmentation routines, mapping out the systematic enrichment of data heterogeneity and model trustworthiness.




Figure 4. Framework of the np-AMIS.

The Artificial Multiple Intelligence System (AMIS) is deployed as an advanced apparatus for fine-tuning the allotment of image augmentation operations. Although originally conceived as a population-reliant heuristic, AMIS has undergone adaptation into a non-population-oriented formulation, termed np-AMIS, expressly for the current context. Within this reimagined np-AMIS scheme, the optimization cycle centers on producing a succession of work packages (WPs), each housing entries drawn from the continuous range [0, 1]. In a scenario featuring, say, five distinct augmentation operators, every WP will accordingly possess five dimensions. An illustrative starting batch of WPs might read: {0.57, 0.98, 0.23, 0.15, 0.67} (Y_{jt}). The conversion of these WP entries into operative weights for each augmentation operator is accomplished through Eq. 1, within which Y_{jt} captures the entry seated at slot 'j' in iteration 't'. At the same time, w_{jt} indicates the weighting conferred upon augmentation operator 'j' during that same iteration. Here, 'J' represents the complete roster of augmentation operators available, and 'k' traverses the slots making up a WP, spanning from 1 to 'J'. This tailored instantiation of np-AMIS within an ensemble framework is designed to resolve the image augmentation optimization puzzle more effectively.

$$w_{jt} = \frac{Y_{jt}}{\sum_{k=1}^J Y_{kt}} \quad \forall j \in J \text{ and } t \in T \quad (1)$$

Should we let the symbol "A_{jt}" denote the tally of images produced by the augmentation operator "j" at iteration "t", and designate "B_t" as the overall desired volume of augmented images, then A_{jt} follows from Eq. 2.

$$A_{jt} = w_{jt} B_t \quad \forall j \in J \text{ and } t \in T \quad (2)$$

Following the execution of a classification cycle, the concluding “accuracy” achieved by the model in that round becomes observable. Let us label this accuracy reading as C_t . The WP is then modified according to the rules set out in [Eqs. 3-10](#). Let C_{best} and Y_{bestj} denote the entry occupying slot j in the most performant WP encountered across the search history.

$$Y_{jt+1} = Y_j^{best} + F(Y_{jt} - R_{jt}) \quad \text{IB} - 1 \quad (3)$$

$$Y_{jt+1} = R_{jt} Y_j^{best} + (Y_j^{best} - Y_{jt}) \quad \text{IB} - 2 \quad (4)$$

$$Y_{jt+1} = \begin{cases} R_{jt} & \text{if } \beta_{jt} \geq CR \\ Y_{jt} & \text{otherwise} \end{cases} \quad \text{IB} - 3 \quad (5)$$

$$Y_{jt+1} = \begin{cases} Y_{jt} & \text{if } \beta_{jt} \leq e^{-\frac{(C^{best} - C_t)}{FT}} \\ Y_{jt-1} & \text{otherwise} \end{cases} \quad \text{IB} - 4 \quad (6)$$

$$Y_{jt+1} = \begin{cases} R_{jt} & \text{if } \beta_{jt} \geq CR \\ Y_{jt} & \text{otherwise} \end{cases} \quad \text{IB} - 5 \quad (7)$$

$$Y_{jt+1} = \begin{cases} Y_j^{best} & \text{if } \beta_{jt} \geq CR \\ Y_{jt} & \text{otherwise} \end{cases} \quad \text{IB} - 6 \quad (8)$$

$$Y_{jt+1} = \begin{cases} Y_j^{best} & \text{if } \beta_{jt} \geq CR \\ R_{jt} & \text{otherwise} \end{cases} \quad \text{IB} - 7 \quad (9)$$

$$Y_{jt+1} = Y_{jt} + F(Y_j^{best} - Y_{jt-1}) \quad \text{IB} - 8 \quad (10)$$

At every successive round, the Work Package (WP) takes on the duty of refreshing the active solution by resorting to a chance-driven updating logic, as codified in [Eq. 11](#). We treat each numerical unit as an Intelligence Box (IB), identified by the indices “l” or “n”. In contrast, the ceiling on allowable rounds is fixed at “L”.

In this setup, θ_{lt} conveys the running sum of rounds across which IB number “l” has been summoned by the WP up through iteration “t”. In practical terms, θ_{lt} increases by 1 each time the WP selects IB “l” and remains frozen whenever

the WP declines to pick IB “l” in a particular round. Alongside this, γ_{lt} captures the average accuracy across all rounds in which the WP opted for IB number “l” during iteration “t”. ω_{lt} functions as an adaptive weight responsive to the short-horizon performance signature of IB l, decaying or strengthening in accordance with its demonstrated contribution in the most recent rounds. The coefficients α and β govern the relative dominance of averaged accuracy over the count of elite solutions unearthed, providing room to tune to match domain-specific priorities or experimenter judgment.

Lastly, τ_{lt} registers the cumulative frequency with which IB number “l” has acted as the pathway to a previously unseen, superior solution. The count τ_{lt} increases by 1 upon each discovery of a new best outcome, conditional on the round in question selecting IB “l”.

$$p_{lt} = \frac{\omega_{lt} + \vartheta_{lt} + \alpha\gamma_{lt} + \beta\tau_{lt}}{\sum_{n=1}^L [\omega_{nt} + \vartheta_{nt} + \alpha\gamma_{nt} + \beta\tau_{nt}]} \quad (11)$$

The “Roulette Wheel Selection” mechanism operates as a probability-weighted lottery to designate an Intelligence Box (IB) to carry out a prescribed function, driven by its associated likelihood value, denoted p_{lt} . This schema casts the selection episode as a spin of a casino wheel: each IB claims a sector of the wheel whose angular span corresponds to its probability mass. Execution begins by sampling a uniform random variate bounded within [0, 1] and effectively setting the wheel in motion. As the rotation unfolds, the sector probabilities are successively accumulated; the instant the running sum surpasses the drawn random variate, the process halts, and the IB whose sector triggered the overshoot is appointed for that round. By blending stochastic perturbation with a systematic tilt toward IBs commanding larger probability allotments, the technique delivers a selection procedure that remains simultaneously efficient and genuinely representative of each IB's fitness for the designated role.

Ensemble image segmentation methods

Dependable isolation of symptomatic regions on cannabis foliage is a linchpin for driving forward automated pathology categorization frameworks. Within the second movement of our research architecture, we undertake a broad-gauge examination of ensemble image segmentation approaches. Our ambition is to weave together heterogeneous segmentation pathways in a mutually reinforcing configuration, thereby boosting both the sharpness and the consistency with which diseased tissue patches are extracted—a development that, in turn, refines the automated classification engines built for Leaf Manifestation Diseases of Cannabis. To realize this objective, we deploy a multi-threaded strategy that unites three state-of-the-art segmentation cores: U-Net, Mask R-CNN, and DeepLabv3+. U-Net, built around a symmetric contracting-expanding topology enriched with skip linkages, shines in semantic segmentation tasks by harvesting richly detailed feature representations at successive resolution stages. Mask R-CNN fuses region-proposal object detection with instance-aware segmentation, enabling precise localization of disease occurrences even in the presence of spatial overlap. DeepLabv3+ exploits atrous convolution alongside multi-scale spatial pyramid aggregation to resolve delicate textural cues and tortuous lesion contours.

Embracing an ensemble philosophy allows us to pool the unique advantages of each method, magnifying the precision with which disease-affected zones are delineated. This fusion not only bolsters model precision but also injects a tier of methodological depth capable of propelling agricultural pathology and automated visual analysis into new territory. These segmentation streams are integrated under the umbrella of the Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) for Deep Learning Ensembles (HAMIS-DLE), as unpacked in the section Population-Based AMIS to Ensemble Deep Learning Model, which provides a resilient and trustworthy architecture for the automated sorting of cannabis pathologies. This deliberate synthesis resets the benchmark for how diseases are governed in agricultural settings, pushing outward the envelope of technical capability and diagnostic fidelity.

Figure. 5 captures the ensemble segmentation framework used throughout our study, with particular emphasis on how U-Net, Mask R-CNN, and DeepLabv3+ process individual pixel intensities. For any given pixel, the numerical readings from

these segmentation engines span the full range from 0 to 1. Pixels yielding an output beneath the 0.5 threshold are rendered entirely black, whereas those producing a value above 0.5 are turned entirely white. This threshold-anchored binarization protocol is uniformly imposed across all three segmentation methodologies.




Figure 5. Framework of HAMIS-DLE in ensemble image segmentation method.

Under the symbolic framework adopted here, the notation Y_{ilt} denotes the output quantity associated with pixel i produced by segmentation technique l in iteration t . The admissible domain for each Y_{ilt} is strictly bounded between 0 and 1 inclusive. The coefficient assigned to each segmentation technique during a given iteration, expressed as w_{lt} , governs how the pixel-level outputs are derived and subsequently synthesized into a unified result. As a concrete illustration, consider the inaugural iteration with a suite of three segmentation techniques at play: a plausible starting coefficient configuration could read $\{0.08, 0.48, 0.47\}$. We use the symbol q_{lt} to refer to these quantities. The full set of coefficients then passes through a normalization transformation prescribed by Eq. 12, where n serves as an index placeholder mirroring l , L denotes the cardinality of the segmentation technique ensemble, and T denotes the total horizon of iterations. Executing this normalization step ensures that every coefficient exerts a sway over the composite segmentation outcome that is directly commensurate with the corresponding technique's tracked effectiveness and situational appropriateness as the iterations unfold.

$$w_{lt} = \frac{q_{lt}}{\sum_{n=1}^L q_{nt}} \quad \forall l \in L, t \in T \quad (12)$$

Designate by Z_{it} the projected value assigned to pixel i throughout iteration t . Deriving this quantity proceeds via Eq. 13, where the scalar I denotes the total number of pixels under consideration.

$$Z_{it} = \sum_{l=1}^L w_{lt} Y_{ilt} \quad \forall i \in I, t \in T \quad (13)$$

The vector q_{lt} is subjected to repeated rounds of adjustment governed jointly by Eqs. 14-16 in conjunction with Eq. 18. Motivated by the goal of reinforcing neighborhood transition operators within a metaheuristic paradigm oriented toward visual data processing and analogous application arenas, four novel constructs are set forth below. Each has been deliberately architected to overcome identified shortcomings of earlier formulations while simultaneously imbuing the overall optimization sequence with greater flexibility and immunity to stagnation.

$$q_{lt+1} = q_{lt} + \lambda (q_{lt} - q_{lt-1}) + \mu (q_{l}^{best} - q_{lt}) \quad (14)$$

$$q_{lt+1} = q_{lt} + \sigma \mathcal{N}(0, 1) (q_{lt} - R_{lt}) \quad (15)$$

$$q_{lt+1} = \sum_{n=1}^t g_k q_{ln} \quad (16)$$

$$g_k = \frac{s_k}{\sum_{u=1}^k s_u} \quad (17)$$

$$q_{lt+1} = q_{lt} + \gamma (q_{lt} - q_{lt-1}) + (1 - \gamma) (q_{lt} - R_{lt}) \quad (18)$$

Eq. 14, which we label the Hybrid Differential-Crossover Update, synthesizes principles extracted from differential evolution and crossover-based recombination. Concretely, it couples the disparity between the incumbent and prior configurations with that between the incumbent and the most meritorious configuration unearthed to date, denoted by q_{bestl} . The scaling terms λ and μ are not held fixed; rather, they adaptively respond to convergence dynamics detected over the preceding stretch of iterations, thereby sharpening model performance. **Eq. 12**, termed the Adaptive Perturbation Strategy, introduces a stochastic perturbation to the search trajectory. Here, σ plays the role of a dispersion multiplier shaped by draws from a standard Gaussian distribution, $\mathcal{N}(0,1)$. The overarching design rationale is to thwart premature entrapment in locally attractive basins by continually feeding a stream of structured randomness into the perturbation signal. This injected randomness grants the optimizer license to venture into outlying regions of the solution manifold, well beyond its current foothold. The serendipitous displacements engendered by $\mathcal{N}(0,1)$ prove vital for exhaustive reconnaissance of the solution terrain, creating opportunities for the eventual discovery of the globally optimal configuration. The magnitude governing these stochastic leaps is regulated by σ , which is anchored at 0.05, thereby calibrating the distance by which the optimizer displaces itself from its immediate neighborhood.

Beyond the above, R_{lt} represents a stochastic scalar sampled uniformly from the interval between zero and one. The term g_k constitutes the normalized transformation of s_k , the latter being yet another random quantity drawn from the unit interval and subsequently normalized in accordance with **Eq. 17**. The surrogate variables n and u are both mapped to the iteration index t . At the same time, the prespecified constant γ is set to 0.2, modulating the adaptation rate of algorithmic control parameters and the pace at which the search sweeps through the solution space. This layered, methodical construction delivers a carefully poised dynamic equilibrium between broad exploration and targeted exploitation throughout the optimization run.

The neighborhood transition operators are spelled out across **Eqs. 14 to 16** and **Eq. 17** are summoned on an iteration-by-iteration basis through the probabilistic selection apparatus formalized in **Eq. 11**. The variable q_{bestl} , which holds the identity of the premier solution recorded up to the current point in the search, is methodically refreshed according to the prescription given in **Eq. 19**. To ascertain the comparative strength of the model we have engineered and to benchmark it against alternative methodologies, we appeal to a triad of distinct evaluative criteria: Area Under the Curve (AUC), F1-score, and accuracy. These metrics receive the respective labels $f_1(q_{lt})$, $f_2(q_{lt})$, and $f_3(q_{lt})$, each reporting the AUC, F1-score, and accuracy on a percentage scale for a given solution q_{lt} . The plain average of this trio of indicators yields $f(q_{lt})$, whose formal definition appears in **Eq. 19**. Complementarily, $f^*(q_{bestl})$ records the best objective function value across the entire search history with respect to q_{bestl} . Its computation follows **Eq. 20**. This systematic evaluation posture enables a holistic characterization of model behavior, paving the way for fair and rigorous comparisons with competing optimization paradigms while providing penetrating insight into the model's correctness, stability, and overall aptitude to converge to high-quality solutions.

$$f'(q_{lt}) = \frac{f'_1(q_{lt}) + f'_2(q_{lt}) + f'_3(q_{lt})}{3} \quad (19)$$

$$q_l^{best} = \begin{cases} q_{lt} & \text{if } f'(q_{lt}) \geq f^*(q_l^{best}) \text{ and update } f^*(q_l^{best}) = f'(q_{lt}) \text{ and} \\ q_l^{best} = q_{lt} & \\ q_l^{best} & \text{otherwise} \end{cases} \quad (20)$$

Ensemble deep learning model development

The third thematic pillar of our inquiry into automated disease categorization for Leaf Manifestation Diseases of Cannabis revolves around harnessing ensemble deep learning methodologies. This juncture mobilizes a deliberately selected portfolio of state-of-the-art CNN backbones, each recruited for its specialized proficiency in heightening both the discriminative resolution and the generalizability of our classification apparatus. Three architecturally distinct CNN variants have been commissioned to tackle complementary dimensions of the categorization challenge:

- ResNet-50 commands attention by virtue of its considerable depth and the deployment of residual skip connections, features that prove indispensable for distilling intricate hierarchical motifs and safeguarding the transmission of layered informational content in cannabis leaf imagery—properties that directly translate into superior diagnostic granularity.
- Inception-ResNetV2 combines the Inception module's deftness at accommodating heterogeneous kernel sizes with the stabilizing influence of residual pathways, demonstrating a pronounced aptitude for feature acquisition across multiple spatial scales. Such multi-resolution sensitivity is a cornerstone for rendering richly detailed depictions of afflicted tissue territories.
- EfficientNet-B3 navigates the trade-off between structural footprint and predictive performance with notable finesse, positioning it as the architecture of choice for deployment environments with constrained computational budgets. It achieves frugal resource consumption without relinquishing the ability to isolate diagnostic image signatures, thus aligning smoothly with practical field-deployment imperatives.

The strategic calculus underpinning the amalgamation of ResNet-50, Inception-ResNetV2, and EfficientNet-B3 into a single cohesive ensemble stems from an intention to pool the non-overlapping strengths these architectures individually contribute. The deep learning of ResNet-50, the scale-agnostic sensitivity of Inception-ResNetV2, and the lean efficiency-to-accuracy ratio of EfficientNet-B3 operate in a mutually reinforcing configuration, collectively expanding the ensemble's aptitude for rendering faithful disease classifications while concurrently suppressing tendencies toward overfitting.

This meticulously choreographed confluence of deep learning traditions embodies a deliberate, principle-anchored orientation, crafted to lift both categorical precision and operational durability. Channeling these architectural workhorses propels forward the frontiers of agricultural pathology and automated visual analytics. The three constituent models will be cohesively integrated through the offices of the Modified Artificial Multiple Intelligence System (M-AMIS), the operational particulars of which are unpacked in the section immediately following.

Beyond this, to construct an ensemble deep learning edifice that is at once rugged and configurable, we have collated a repertoire of widely ratified hyperparameter choices intended to furnish a foundational launch configuration for the training regimen. As enumerated in [Table 2](#), these settings encompass learning rate schedules, mini-batch granularities, regularization strategies, and optimizer selections, all grounded in empirically validated best practices. It warrants explicit

acknowledgment that subsequent tailoring of these hyperparameters may be necessary due to the dataset's idiosyncrasies and the specific demands of the classification objective.

Table 2. Hyperparameter configuration.

Hyperparameter	Revised description/Configuration
Learning Rate	Set to 0.001, as this is a widely accepted initial value for model training
Number of Epochs	Training was conducted for 100 epochs, while applying early stopping to minimize overfitting
Batch Size	A batch size of 64 was selected to achieve a balance between computational efficiency and gradient stability
Weight Decay (L2)	An L2 regularization coefficient of 1e-4 was employed to reduce overfitting effects
Regularization (e.g., Dropout)	Dropout regularization can be incorporated when required to further improve generalization
Optimizer	The Adam optimization algorithm was utilized due to its broad adoption and effectiveness
Learning Rate Scheduler	ReduceLROnPlateau was applied to dynamically decrease the learning rate during stagnation
Loss Function	Cross-entropy loss was adopted as it is commonly used for classification problems
Ensemble Method	The proposed ensemble strategy was based on the pop-AMIS method

Population-based AMIS to ensemble deep learning model

What follows advances the treatment of decision fusion by mobilizing the population-based artificial multiple intelligence system (pop-AMIS), whose conceptual origins trace to Pitakaso *et al.* [30]. The framework was initially forged to disentangle the knotty optimization landscapes encountered in farm logistics network planning; here, it is repurposed to hone ensemble weighting schemes within decision fusion workflows. The pivot of our undertaking is the transplantation and meticulous tuning of the pop-AMIS apparatus to accurately estimate contribution weights across ensembles built from multiple Convolutional Neural Network (CNN) architectures.

We have reshaped the intelligent box mathematics to better accommodate the peculiarities of the current research setting, yielding tangible departures from the stock pop-AMIS blueprint. Further novel refinements are embedded to magnify the overall potency of the methodology.

To form a reliable picture of pop-AMIS performance, we stage a head-to-head benchmarking campaign against the unmodified AMIS as well as a suite of substitute decision fusion mechanisms—specifically, the Unweighted Average Model (UWM), the Differential Evolution Algorithm (DE) attributable to Wan *et al.* [31], and the Particle Swarm Optimization (PSO) formulation credited to He *et al.* [32]. The resulting panoramic analysis yields diagnostic insights into the efficiency characteristics of our AMIS-rooted strategy, highlighting its latent strengths when applied to decision fusion optimization.

Under the unweighted average model (UAM), every predictive output (Y_{ij})—where ‘ i ’ ranges over CNN identifiers and ‘ j ’ ranges over category labels—receives an identical, flat contribution weight. Within fusion computations, UAM invokes Eq. 21 to realize this, whereas pop-AMIS, AMIS, DE, and PSO all lean on Eq. 22 to resolve the final weight assignments. The notation Y_{ij} denotes the expectation produced by CNN ‘ i ’ for class ‘ j ’. Once the outputs of the CNN panel have been pooled, the scalar V_j takes on the role of the decision statistic for assigning class ‘ j ’. The contribution attributed to CNN ‘ i ’

is designated W_i , while the quantities 'I' and 'J' respectively enumerate the CNNs in the ensemble and the classes in the problem definition.

$$V_j = \frac{\sum_{i=1}^I Y_{ij}}{I} \quad (21)$$

$$V_j = \sum_{i=1}^I W_i Y_{ij} \quad (22)$$

Pop-AMIS was instrumental in uncovering the optimal specification of W_i . The unweighted average decision fusion tactic (UWA) offers the virtues of conceptual simplicity and low computational overhead, yet, by design, it eschews optimization. Pop-AMIS, on the other hand, distinguishes itself through aggressive global exploration and pronounced ruggedness, making it well-suited to problems of considerable complexity—though slower convergence dynamics and higher memory consumption accompany these attributes. The pop-AMIS schema rests on five keystone phases that orchestrate its functioning: (1) randomly seeding a population of work packages (WP), (2) each WP stochastically nominating an improvement box (IB) to engage, (3) applying the WPs to their nominated IBs, (4) refreshing the accumulated heuristic knowledge, and (5) looping through phases (2) through (4) until a stopping rule triggers. As a population-centered heuristic that leans on repeated cycles to edge toward promising solutions, pop-AMIS demanded the adjustment codified in [Eq. 23](#).

$$V_{hjt} = \sum_{i=1}^I Z_{hit} Y_{ij} \quad (23)$$

The variable V_{hjt} stands for the class reference emitted by WP, subscripted h , at iteration t . Z_{hit} holds the meaning of the tuned weight seated at coordinate i —which maps to CNN “ i ”—within WP “ p ” over the course of iteration “ t ”. The length of a Work Package (WP), captured by D , is in one-to-one correspondence with the number of CNNs folded into the ensemble. Meanwhile, the symbol NP indicates the total number of WPs in circulation. A concrete snapshot: a WP instance with $D = 3$ and $NP = 2$ could be initialized as $\{[0.94, 0.23, 0.05], [0.54, 0.87, 0.14]\}$. It is worth emphasizing that the initial values are randomly sampled from a uniform distribution over $[0, 1]$ at the very first iteration. The total WP pool size is indicated by H , numerically matching NP . Each Work Package (WP) autonomously commits to an Improvement Box (IB) to enhance the current candidate solution. The selection of which IB a given WP settles on at each iteration is regulated by a tailored probability function ([Eq. 24](#)), executed via a roulette wheel draw, as previously described.

$$P_{bt} = \frac{e^{\lambda FN_{bt-1}} + \mu A_{bt-1}^v + K e^{\theta U_{bt-1}}}{\sum_r^B e^{\lambda FN_{rt-1}} + \mu A_{rt-1}^v + K e^{\theta U_{rt-1}}} \quad (24)$$

At the current iteration, P_{bt} is the probability that the Improvement Box (IB) indexed by “ b ” is selected. This probabilistic expression folds in a recommended scaling constant F , empirically anchored at 0.7, with r acting as a dummy loop variable aligned with b and B marking the ceiling on the IB count. The formula housing P_{bt} incorporates several non-linear elements, proposed to flexibly reweight the selection drivers:

- $e^{\lambda FN_{bt-1}}$: Here, the selection frequency λFN_{bt-1} is passed through an exponential map, scaled by λ (fixed at 0.3) and F (fixed at 0.5). The net effect is to heap extra weight onto strategies that have been repeatedly invoked, dynamically spotlighting those that have compiled solid usage histories.
- μA_{bt-1}^v : This component subjects the mean performance metric A_{bt-1}^v to a power-law distortion, with μ (0.8) serving as the linear coefficient and v (1.5) as the exponent. The non-linear curvature amplifies the imprint of elevated

performance readings, making the selection rule disproportionately responsive to exceptionally strong results.

- $Ke^{\theta U_{bt-1}}$: This term attaches an exponential boost to the success count U_{bt-1} , scaled by K (5) and inflated exponentially via the factor θ (0.2), which together place a heavy thumb on the scale in favor of recent breakthroughs with large impact.

Thanks to the embedding of exponential and power-law transformations, the revamped formula can pivot more aggressively as the relative effectiveness of different IBs shifts. The selection module becomes considerably more responsive to the most recent performance signals, accelerating the algorithm's progress toward high-quality solution basins. The control knobs λ , μ , ν , and θ remain adjustable in response to streaming performance feedback, conferring adaptivity and landscape sensitivity upon the formulation. The design treads a line between probing infrequently activated but possibly rewarding strategies (through the exponential amplification of FN) and doubling down on strategies that are presently producing the strongest outcomes (through the exponential accentuation of U). Every WP finalizes its IB selection for the upcoming refinement step based on the probability map specified in Eq. 25. This investigation uses 10 separate IBs, whose specifications are tabulated in Table 3 (spanning Eqs. 25-34).

Table 3. Details of IBs.

Name of IB	IB-Equation	Equation
IB-1	$Z_{hjt}=[(1-\rho)Z_{rjt}+\rho(Z_{mjt})]\times Z_{rjt}$	(25)
IB-2	$Z_{hjt}=[(1-\rho)Z_{rjt}+\rho(Z_{mjt})]\times Z_{mjt}$	(26)
IB-3	$Z_{hjt}=Z_{rjt}+\rho(Z_{mjt}-Z_{rjt})$	(27)
IB-4	$Z_{hjt}=B^{gbest}_j+\rho(Z_{mjt}-Z_{rjt})$	(28)
IB-5	$Z_{hjt}=\rho B^{gbest}_j+(1-\rho)Z_{rjt}+\nu(Z_{mjt}-Z_{rjt})$	(29)
IB-6	$Z_{hjt}=\omega Z_{hjt-1}+RhjtF1(Z_{hjt-1}-Z_{rjt})+RhjtF2(B^{gbest}_j-Z_{hjt-1})$	(30)
IB-7	$Z_{hjt}=\omega Z_{hjt-1}+RhjtF1(Z_{hjt-1}-Z_{rjt})+RhjtF2(B^{gbest}_j-Z_{mjt})$	(31)
IB-8	$Z_{hjt} = \begin{cases} Z_{hjt-1} & \text{if } Rhj \leq CR \\ Rhjt & \text{otherwise} \end{cases}$	(32)
IB-9	$Z_{hjt} = \begin{cases} Z_{hjt-1} & \text{if } Rhj \leq CR \\ Z_{rjt} & \text{otherwise} \end{cases}$	(33)
IB-10	$Z_{hjt}=Rhjt$	(34)

The full complement of Particle Swarm Optimization (PSO) parameters, including ω (inertia weight), follows the calibration recommended by Pitakaso *et al.* [33], with ω fixed at 0.05.

Within the ongoing iteration, $Rhjt$ denotes the stochastic quantity associated with WP “h” and iteration “t” at slot “j”. At the same time, Rhj represents the random quantity associated solely with Work Package “h” at slot “j”. The defining equations of the IBs remain faithful to the specification given by Pitakaso *et al.* [30]. $Zhit$ reports the numerical entry residing at position “i” within Work Package “h” during iteration “t”. Work Packages “h”, “r”, and “m” (with “m” required to differ from “h”) are all drawn from the WP collection (indices spanning 1 through H). $Bgbest_j$ points to the globally dominant Work Package. Both scaling coefficients, F1 and F2, are set to 3, in keeping with precedent established in earlier studies. Concurrently, the crossover probability (CR) is held at 0.8.

For iteration “t”, $Rhjt$ captures a freshly generated Work Package spawned from Work Package “h” at position “j”. Eq. 35 governs the transition to $Zhjt+1$. At the same time, $f(Zhjt)$ returns the classification accuracy realized when $Zhjt$ is deployed as the governing decision fusion policy. The Genetic Algorithm (GA), Differential Evolution (DE), and PSO each invoke their respective updating routines—Eqs. 36-38—to drive the evolution of the values encapsulated in $Zhjt$.

$$Z_{hjt+1} = \begin{cases} Z_{hjt} & \text{if } f(Z_{hjt}) \geq f(Z_{hjt-1}) \text{ and update } f(Z_{hjt+1}) = f(Z_{hjt}) \\ & \text{and } Z_{hjt+1} = Z_{hjt} \\ Z_{hjt-1} & \text{otherwise} \end{cases} \quad (35)$$

$$Z_{hjt} = [(1 - \rho) Z_{hrt} + \rho(Z_{hmt})] \times Z_{hrt} \quad (36)$$

$$Z_{hjt} = Z_{rjt} + \rho(Z_{hmt} - Z_{hnt}) \quad (37)$$

$$Z_{hjt} = \omega Z_{hjt-1} + R_{hjt} F1(Z_{hjt-1} - Z_{rjt}) + R_{hjt} F2(B_j^{gbest} - Z_{hjt-1}) \quad (38)$$

Performance metrics

To gauge the relative strength of our proposed methodology against competing approaches, we rely on three principal evaluative metrics: Area Under the Curve (AUC), F1-score, and accuracy. Each of these indicators illuminates a separate facet of model behavior. AUC springs from the Receiver Operator Characteristic (ROC) plot and gauges the classifier's capacity to tease apart distinct classes. It captures the degree to which the model successfully flags genuine positive instances from the pool it labels as positive, thereby providing a well-rounded picture of classification fidelity.

The F1-score, a staple in binary classification settings, weds precision (the correctness of positive calls) with recall (the ability to surface all pertinent cases) into a single composite figure. Bounded between 0 and 1, where loftier values signal stronger performance, the F1-score proves especially instructive when models are trained on skewed class distributions, as it penalizes both missed detections and spurious alarms. Accuracy provides an overall gauge of model performance across all categories, simply the ratio of correctly predicted cases to the total number of predictions. This metric carries particular weight in scenarios where every class is considered equally consequential to the analysis.

Comparative methodology

The experimental campaign was partitioned into three testing blocks, with the candidate techniques sorted into these same groupings. In the first block, we scrutinized architectures such as AlexNet, ResNet50, MobileNetV2 paired with SVM, and VGG16, all of which featured prominently in the study by Sethy *et al.* [26]. Further methods previously chronicled by Ritharson *et al.* [23], Bari *et al.* [34], Pandian J. *et al.* [35], among others, were also examined—these included Faster R-CNN, ResNet-50, YOLOv4, and various ensemble configurations.

Turning to the second data corpus, our evaluation included DenseNet-169, Inception V3, and Xception, along with an ensemble construct advanced by MP and Reddy [27]. Beyond this, we assembled and benchmarked an internally collated dataset through a homogeneous ensemble framework built from multiple CNN backbones: ResNet-50, Inception-ResNetV2, EfficientNet-B3, VGG-19, VGG-16, ResNet-152, and InceptionV3. The respective ensemble models drew on varying CNN counts—7, 3, 13, 5, 5, 3, and 7, in order. The storage footprints of these methods, expressed in megabytes (MB), are listed in Table 4 for ease of reference, offering a panoramic view of the breadth of our comparative design and the extensive range of models evaluated throughout the study.

Table 4. Detail the size of the training and testing sets for the compared methods.

Empty cell	Total size	Size in MB	

ResNet-50	688.8	98.4	7
Inception-ResNetV2	642	214	3
EfficientNet-B3	637	49	13
VGG-19	1092	546	2
VGG-16	1000	500	2
ResNet-152	717	239	3
InceptionV3	644	92	7
HAMIS-DLE	606.8	2,1,4	1

Constructing an automated cannabis leaf disease classification system (A-CLDCS)

Figure 6 depicts the embedding of the Hybrid Adaptive Multiple Intelligence System for Deep Learning Ensembles (HAMIS-DLE) inside the Automated Cannabis Leaf Disease Classification System (A-CLDC-S). The schematic depicts an intuitive smartphone application that communicates with a backend server via the LINE messaging platform. In the illustrated pipeline, a mobile user captures and transmits a photograph of a cannabis leaf through the LINE application; the image is then routed to the server for processing. The server draws on HAMIS-DLE—a state-of-the-art ensemble of deep learning models calibrated for elevated accuracy and flexible disease recognition. The deep learning pipeline inspects the submitted image and cross-references it against a repository of catalogued pathologies.




Figure 6. Framework design of A-CLDC-S case study.

The database management layer, supervised by an administrator, fields query records and curated response templates, which are relayed back to the end user. This sequence enables instantaneous, on-the-spot categorization of leaf health status, returning verdicts such as Powdery Mildew, Downy Mildew, and other cannabis-specific conditions. By weaving HAMIS-DLE into the fabric of A-CLDC-S, the system's operational potency is markedly amplified, capitalizing on multiple intelligence paradigms to secure high diagnostic precision—an asset of central importance for sound crop stewardship and agricultural decision support.

The A-CLDC-S platform was engineered to equip growers with the means to identify anomalous leaf conditions in cannabis plants. **Figure 6** sketches the architectural blueprint of the A-CLDC-S crafted for this initiative. As the figure conveys, a prospective user wishing to query the system about cannabis leaf ailments must first activate the LINE Application on their mobile device, scan a QR code, and add A-CLDC-S as a contact during the initial setup. Thereafter, the user can register and interact with A-CLDC-S, which presents a basic menu interface. Upon selecting an item from the menu (Message input), the LINE Application dispatches a message request to the LINE platform hosting the A-CLDC-S service, which inspects the incoming message event received from the LINE Application. Once this check completes, a reply message is routed to the LINE platform, which in turn pushes the response back to the LINE Application for the user to read. From the administrator's perspective, the model derived in the preceding section is the one deployed in A-CLDC-S. The A-CLDC-S, operating through the LINE Message API, was cooperatively integrated with the CNN via the Flask API to generate classification outcomes. When A-CLDC-S receives an image forwarded by the user, it relays it to the A-CLDC-S server's webhook, which then passes it to the deep learning model. The deep learning model then predicts the disease present in the submitted image. The resulting diagnostic output is returned in textual form via the LINE Messaging API and the LINE platform. This system is capable of round-the-clock, real-time operation.

Results and Discussion

Our numerical experiments were run in the Google Colab environment, which was provisioned with an NVIDIA Tesla V100 GPU with 16 GB of RAM, a configuration that substantially accelerated the training of our machine learning models. Performance assessments were conducted on a machine equipped with dual Intel Xeon processors clocked at 2.30 GHz and 52 GB of system RAM. This setup was augmented with a Tesla K80 GPU, contributing an additional 16 GB of dedicated GPU memory and thereby reinforcing our computational capacity.

Identifying the optimal synergy of the proposed methodology

An exhaustive battery of experiments was mounted to pinpoint the most favorable configuration for our model architecture, with the 10-LMDC dataset serving as the primary testbed. Our investigation spanned both augmented and non-augmented model variants, incorporating four separate image segmentation pipelines—three built from a single technique and one combining multiple techniques. In parallel, we folded four Convolutional Neural Network (CNN) designs into the experimental matrix. To methodically probe these arrangements, the CNN architectures were arranged into five homogeneous ensemble constructs alongside one heterogeneous construct. A palette of decision fusion strategies was also explored: unweighted average (UWA), Differential Evolution Algorithm (DE), Genetic Algorithm (GA), Particle Swarm Optimization (PSO), Ant Colony Optimization (ACO), and the hybrid Artificial Multiple Intelligence System (HAMIS-DLE).

This sweeping experimental campaign yielded 192 distinct trials, each rigorously measured against central performance indicators—accuracy, precision, recall, and the F1-score—as recorded in **Figure 7**. The resulting dataset constitutes a valuable empirical foundation, spotlighting the relative potency of alternative configurations and guiding the selection of the most fitting approaches for particular image classification assignments.




Figure 7. Comparative performance analysis of machine learning models across different metrics and strategies.

Interrogating the optimal model configuration framework in **Figure 7** reveals which combinations of augmentation, segmentation, CNN ensemble composition, and fusion strategy achieve peak performance. Non-augmented models achieved 82.32% accuracy, whereas their augmented counterparts substantially outperformed them, reaching 94.96%—a finding that powerfully underscores the indispensable role of data augmentation. Within segmentation, the heterogeneous pairing that fused U-Net with Mask R-CNN achieved 93.64% accuracy, compared to 91.67% and 91.35% for homogeneous U-Net and homogeneous Mask R-CNN, respectively. This points to the tangible benefits of a mixed segmentation posture.

Regarding CNN ensembles, the heterogeneous ensemble model outperformed all homogeneous alternatives, achieving 93.64% accuracy compared to 81.45% for the strongest homogeneous contender (ResNet-50). This reinforces the value of weaving together diverse CNN architectural lineages.

The appraisal of six fusion strategies identified the Modified Artificial Multiple Intelligence System (M-AMIS) as the top performer, achieving the highest accuracy of 96.38%. In summation, the premier configuration for image classification tasks, most notably on the 10-LMDC corpus, marries data augmentation, a combined U-Net and Mask R-CNN segmentation approach, a heterogeneous CNN ensemble, and the M-AMIS fusion strategy, culminating in a best accuracy of 96.38%.

Comparative analysis of the proposed approaches and state-of-the-art methods

What follows pits the model we advocate—a construct in which a hybrid artificial intelligence system serves as a decision fusion engine (HAMIS-DLE)—against the homogeneous ensemble baselines listed in **Table 5**. Three quantitative touchstones are enlisted for the comparison: accuracy, AUC, and F1-score.

Table 5. Performance comparison of HAMIS-DLE with homogeneous ensemble models.

Methods	F1-Score (%)	AUC (%)	Accuracy (%)	Testing time (second/image)	Training time (min)
ResNet-50	92.00	92.10	92.01	0.05	124.01
Inception-ResNetV2	91.97	91.97	91.87	0.04	111.48
EfficientNet-B3	92.37	92.30	92.31	0.05	121.59
VGG-19	95.43	95.31	95.18	0.12	195.93
VGG-16	94.41	94.25	94.22	0.10	172.40
ResNet-152	92.90	92.86	92.84	0.06	147.08
InceptionV3	92.64	92.61	92.49	0.05	132.49
HAMIS-DLE	99.86	99.61	99.89	0.02	103.10

Guided by the evidence marshalled in **Table 5**, we place the proposed HAMIS-DLE alongside its homogeneous ensemble counterparts and examine their relative standing. The lens through which this comparison is conducted brings into focus three fundamental yardsticks: accuracy, Area Under the Curve (AUC), and F1-score. Below, we unpack and interpret the record posted by HAMIS-DLE, comparing it to the other techniques on these very metrics.

A striking 99.58 % accuracy rate is what HAMIS-DLE delivers—a result that leaves every competing entry from the table far behind. By way of illustration, ResNet-50, Inception-ResNetV2, and EfficientNet-B3 all reach accuracy plateaus in the

lower 90% range. VGG-19, VGG-16, ResNet-152, and InceptionV3 cluster in broadly similar territory, yet none reaches the summit, as measured by HAMIS-DLE.

Looking next at AUC, HAMIS-DLE yields a value of 99.61%, a figure reflective of virtually top-tier classification quality overall. Here again, it distances the pack. The runner-up on this particular measure, ResNet-50, posts an AUC of 92.10%—a respectable reading on its own terms but a distant second when held against the commanding height achieved by HAMIS-DLE.

Turning to the F1-score, HAMIS-DLE secures 99.60%, a reading that testifies to its capacity for striking an unusually effective equilibrium between precision and recall. The lead it opens over the field on this metric is once more substantial. Scouring the whole set of techniques appearing in the table, none comes within reach of the F1-score stratum occupied by HAMIS-DLE.

When the verdict is drawn from the triad of accuracy, AUC, and F1-score, HAMIS-DLE emerges as the dominant model among those compared in [Table 5](#). It brushes up against flawless accuracy, posts an AUC figure leagues beyond the rest, and garners an F1-score of outstanding magnitude. These data points constitute a compelling evidentiary basis for recognizing HAMIS-DLE as a highly capable model, exceptionally well-fitted to the problem at hand. Decisively, it vaults past the performance ceilings established by the remaining architectures populating the experimental slate.

[Table 6](#) compares HAMIS-DLE, a heterogeneous ensemble design, with a panel of homogeneous configurations. The assessment is performed via k-fold cross-validation on the PBW corpus. What emerges is the steady ascendancy of HAMIS-DLE across a suite of indicators—accuracy, precision, recall, and F1-score—under both 5-fold and 10-fold validation schedules.

Table 6. Results of K-fold cross-validation.

Methods	10-folds			5-folds		
	F1-Score (%)	AUC (%)	Accuracy (%)	F1-Score (%)	AUC (%)	Accuracy (%)
ResNet-50	92.00 ± 2.41	92.10 ± 1.47	92.01 ± 1.42	91.78 ± 2.40	91.89 ± 2.59	91.91 ± 1.22
Inception-ResNetV2	91.97 ± 2.03	91.97 ± 1.99	91.87 ± 1.04	91.85 ± 2.31	91.83 ± 1.47	91.58 ± 1.98
EfficientNet-B3	92.37 ± 1.85	92.30 ± 2.45	92.31 ± 1.49	92.31 ± 2.41	92.26 ± 1.39	92.06 ± 1.48
VGG-19	95.43 ± 1.75	95.31 ± 0.97	95.18 ± 2.83	95.40 ± 1.87	95.02 ± 1.58	94.94 ± 1.84
VGG-16	94.41 ± 2.49	94.25 ± 1.02	94.22 ± 1.94	94.15 ± 2.59	94.01 ± 2.41	94.17 ± 1.58
ResNet-152	92.90 ± 1.85	92.86 ± 1.87	92.84 ± 1.07	92.81 ± 1.02	92.72 ± 1.19	92.60 ± 1.49
InceptionV3	92.64 ± 1.08	92.61 ± 1.18	92.49 ± 1.42	92.58 ± 1.54	92.42 ± 1.48	92.25 ± 1.05
HAMIS-DLE	99.86 ± 0.87	99.61 ± 1.05	99.89 ± 1.41	99.38 ± 1.03	99.47 ± 1.94	99.40 ± 1.31

Within the 5-fold validation framework, HAMIS-DLE consistently outperforms competitors such as ResNet-50, Inception-ResNetV2, EfficientNet-B3, and the VGG lineage, achieving accuracy, AUC, and F1-score north of 99%—strong evidence of its robustness and consistency. That supremacy carries over into the 10-fold validation scenario, where HAMIS-DLE sustains its poise and power, preserving lofty scores even as the evaluation rigor intensifies. In contrast, the comparator models betray wider swings in performance. Weighing the evidence across both regimes, HAMIS-DLE confirms its status as a model of considerable reliability, apt for deployment in real-world image categorization pipelines and demanding cross-validation routines alike. [Figure 8](#) presents the confusion matrix, and [Figure 9](#) provides the GradCAM heatmap for HAMIS-DLE.




Figure 8. Confusion matrix of the proposed model.




Figure 9. Heatmap GradCAM of the proposed model.

The confusion matrix, labelled "**Figure 8**" in the text, charts how a sophisticated deep learning architecture fares when distinguishing among multiple plant health categories. Such matrices rank among the most informative tools for assessing the accuracy of classification systems. Diagonal cells populated with large counts reflect frequent correct assignments; sizeable off-diagonal entries, conversely, flag classification slips. That the diagonal overwhelmingly dominates in this particular matrix suggests the model possesses considerable prowess in sorting plant health states and stumbles only rarely, thereby underscoring the genuine predictive heft it commands.

"**Figure 9**", for its part, unpacks the model's sorting of diverse leaf conditions through a gallery of heatmaps constructed via Gradient-weighted Class Activation Mapping (GradCAM). These heatmaps pinpoint the foliage zones the model assigns the greatest weight to when making its decision, with illumination intensity scaling with the region's importance to the final classification. Visualizations of this kind furnish an indispensable window into the model's otherwise opaque reasoning, revealing the leaf features the algorithm latches onto to demarcate healthy from diseased specimens.

Comparative analysis of the proposed model and existing methods for leaf disease classification

This portion of the study draws on two datasets culled from distinct peer-reviewed contributions. After retrieving publicly posted data, we feed it into our model, which identifies leaf disease categories. **Tables 7 and 8** summarize the experimental results obtained on each corpus.

Table 7. Comparative results for rice leaf disease [23, 26].

Method	F1-score	Accuracy	Refs.
AlexNet	97.62	97.63	[26]
ResNet50	98.38	98.38	[26]
mobilenetv2 plus SVM	97.96	97.96	[26]
ResNet50 model plus SVM	86.70	86.70	[26]
VGG16	—	99.70	[23]

Faster- RCNN	–	98.85	[34]
ResNet197 model		99.58	[35]
YOLOv4 model	–	95.90	[36]
Ensemble models	–	99.10	
HAMIS-DLE	99.82	99.81	Our model

Table 8. Comparative Results for groundnut plant leaf disease [27].

Method	F1-score	Accuracy	Refs.
Densenet 169	98.42	98.45	[27]
Inception V3	92.37	92.27	[27]
Xception	93.86	93.71	[27]
Ensemble model	98.46	98.46	[27]
HAMIS-DLE	99.71	99.62	Our model

Turning first to **Table 7**, which collates rice leaf disease classification figures across an assortment of published methods: our submitted model, HAMIS-DLE, clocks an accuracy of 99.81% alongside an F1-score of 99.82%. Both numbers tower over the great majority of comparator entries. On accuracy as well as on F1-score, HAMIS-DLE betters every alternative method represented. The nearest accuracy challenger—the VGG16 instance documented by Ritharson *et al.* [23]—reaches 99.70%, though an F1-score is absent from that report. The inference is that HAMIS-DLE qualifies as a standout model for the rice leaf disease task, marrying near-perfect accuracy with a well-calibrated F1 Score.

Table 8 focuses on the classification of groundnut leaf diseases. In this arena, HAMIS-DLE attains 99.62% accuracy and 99.71% F1-score. It again surpasses all comparator entries on both dimensions. The most competitive accuracy among the alternatives is achieved by the DenseNet-169 variant from M.P. and Reddy [27], which achieves 98.45%. On the F1-score axis, HAMIS-DLE preserves a comfortable advantage over the field, reinforcing the picture of superlative performance on the groundnut leaf disease classification problem.

Summing up, the evidentiary record assembled in **Tables 7 and 8** sends an unambiguous signal: HAMIS-DLE constitutes a model of exceptional caliber for leaf disease classification. It reliably outperforms incumbent techniques on accuracy and F1-score alike across both rice leaf disease and groundnut plant leaf disease classification tasks. These results cast into sharp relief the potency and pre-eminence of HAMIS-DLE in meeting leaf disease classification challenges, positioning it as a promising and dependable instrument for this agriculturally significant application.

Automated cannabis leaf disease classification system (A-CLDC-S)

Figure 10 captures the deployed A-CLDC-S chatbot in authentic use cases, illustrating the outputs generated by the architectural design shown in **Figure 6**. The screenshots document the chatbot's ability to diagnose cannabis foliar ailments through live back-and-forth with end users. Panels a, b, and c each exhibit the mobile application interface at the moment a user has submitted a cannabis leaf photograph and the A-CLDC-S has returned an identification along with a disease classification. In **Figure 10a**, the system flags Powdery Mildew and provides a verdict accompanied by descriptive text and an outline of potential consequences. **Figure 10b** illustrates the recognition of Nitrogen Deficiency, highlighting how the condition alters leaf pigmentation and hampers plant development. **Figure 10c** shows the pinpointing of Caterpillar Infestation together with an explanation of the damage inflicted on foliage. These captured interactions attest to the approachable design of A-CLDC-S and its seamless integration with the powerful HAMIS-DLE backbone,

underscoring the platform's capacity to deliver accurate, knowledge-rich diagnostics and the value of the embedded Hybrid Adaptive Multiple Intelligence System in strengthening decision-making in cannabis cultivation.




Figure 10. User interface of A-CLDC-S (a) A-CLDC-S chatbot interface identifying Powdery Mildew on a cannabis leaf, (b) The system diagnoses Nitrogen Deficiency, indicating its impact on plant health, and (c) Chatbot detection of Caterpillar Infestation, showing its damage to the plant.

To empirically test A-CLDC-S, we fed the system images drawn from a pooled mixture of training and test partitions, with roughly 2372 samples picked at random. The behavioral profile of the LINE chatbot is summarized in [Table 9](#).

Table 9. The A-CLDC-S accuracy result.

Class	Accuracy (%)	Incorrect classifications	Correct classifications	Number of videos
Powdery Mildew	98.71	3	230	233
Downy Mildew	98.18	4	216	220
Leaf Septoria	98.32	4	234	238
Botrytis Aphids	99.03	2	204	206
Spider Mites	98.92	2	183	185
Leaf Miner	98.48	3	195	198

Cannabis Rust	98.33	4	236	240
Cannabis Mosaic	98.59	3	210	213
Virus Cannabis	99.18	2	241	243
Witches' Broom	98.97	2	193	195
Healthy Leaf	99.00	2	199	201
Total	98.70	31	2341	2372

The "Automatic Cannabis Leaf Disease Classification System (A-CLDC-S)" has demonstrated strong accuracy, averaging 98.70% across the spectrum of cannabis leaf conditions it was tasked with recognizing—a result that buttresses its reliability for on-farm deployment. Diagnostic prowess is especially evident in cases of Botrytis, Aphids, and Spider Mites. In contrast, performance declines somewhat when challenged by Downy Mildew and Leaf Septoria, suggesting that gains may still be realized through algorithmic tuning or expanding the training corpus.

A sizable pool of 2372 images underwrites the system's sturdiness, a prerequisite for trustworthy disease scouting within crop management workflows. Encouraging as these numbers are, further corroboration under authentic field conditions remains necessary to verify efficacy across heterogeneous environments, and continued enrichment of the dataset will be pursued to improve accuracy.

A follow-up survey, structured around the System Usability Scale (SUS) and conducted with 31 farmers who interacted with A-CLDC-S, was conducted to probe user-friendliness. Participants rated a series of statements on a 5-point scale from 1 (Strongly Disagree) to 5 (Strongly Agree); the resulting scores were averaged to yield a quantified usability index. **Table 10** consolidates these ratings, aligning with growers' feedback and guiding forthcoming refinements to elevate both the user experience and system capabilities.

Table 10. Individual SUS-score.

Respondent	SUS score	Respondent	SUS score	Respondent	SUS score
1	94	21	91	11	95
2	93	22	90	12	92
3	93	23	92	13	94
4	98	24	90	14	93
5	94	25	91	15	95
6	92	26	95	16	95
7	94	27	96	17	96
8	96	28	93	18	90
9	96	29	90	19	91
10	92	30	93	20	95
	95	31			
Average				93.5	

Table 10 presents the System Usability Scale (SUS) scores logged for A-CLDC-S across 31 evaluators. An overall mean of 93.5 signals pronounced user satisfaction. The bulk of individual scores landed above 90, signaling broad-based endorsement of the system's functionality and ease of use. Particularly noteworthy are the ratings of 96 contributed by respondents 7, 10, and 11, which denote peaks of satisfaction. In contrast, the handful of comparatively lower marks—such as the readings in the 90s from respondents 8 and 12—indicate specific facets that merit further polish. Taken as a whole, the tight score dispersion mirrors a uniformly favorable perception of usability.

The convergence of ratings suggests that A-CLDC-S is a sturdy, instinctive platform; a lofty SUS average wedded to consistent scoring underlines its practical effectiveness. The overwhelmingly positive reception indicates close alignment with user expectations, even as the modest spread in scores exposes nuanced differences in individual encounters and the room that still exists for progressive improvement. In aggregate, the SUS evidence validates the system's design choices and provides a roadmap for enhancements that will extend its utility for cannabis leaf disease stewardship.

Advancing cannabis disease management: insights from the proposed model's performance

What our results represent is a material step forward in grappling with the difficulties besetting cannabis production, most pointedly the machine-driven sorting of Leaf Manifestation Diseases of Cannabis (LMDC). The Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) for Deep Learning Ensembles (HAMIS-DLE) presented in this paper achieves significant improvements in the reliable identification of various cannabis leaf pathologies. This development carries weight for the pursuit of sustainable horticultural regimens.

Cannabis growing, for all its economic allure, is dogged by grave impediments, above all in preserving plant vigor against afflictions such as Powdery Mildew, Downy Mildew, Leaf Septoria, and Botrytis. These pathogens eat into both harvest volume and grade while simultaneously undercutting the viability of ecologically minded cultivation. As underscored by Jerushalmi *et al.* [12] and Punja [13], the swell in commercial-scale operations has magnified the spread of these biotic threats, sharpening the demand for workable disease management countermeasures.

The HAMIS-DLE framework achieves 99.58% accuracy and 99.61% AUC—figures that mark genuine progress in LMDC categorization and handily surpass legacy benchmarks such as ResNet-50, Inception-ResNetV2, and EfficientNet-B3. Precision at this level is indispensable for early scouting and containment, both of which directly affect crop vitality and productivity. The performance profile of our model aligns with the mounting call for sustainable, high-efficiency disease governance in agriculture, as echoed by Akanmu *et al.* [16] and Tariq *et al.* [17], both of whom stress the centrality of forward-looking, ecologically rational approaches within the farming sector.

Beyond performance metrics, the tactic of weaving together disparate CNN backbones and entrusting decision fusion to M-AMIS embodies a methodological advance of note. This design capitalizes on the differentiated competencies of each constituent architecture, as reflected in the respective strengths of ResNet-50, Inception-ResNetV2, and EfficientNet-B3 at teasing out subtle feature hierarchies and husbanding computational resources. Harnessing M-AMIS to calibrate fusion weightings amplifies the collective power of the ensemble, a move consonant with recent AI deployments spanning the agricultural and clinical domains, as showcased by Sethanan *et al.* [21] and Nanthasamroeng *et al.* [20].

The present inquiry also helps plug the conspicuous gap in the scholarly record concerning automated LMDC classification. Earlier efforts centered on other species—among them the rice-focused work of Aggarwal *et al.* [5] and the corn-oriented study by Amin *et al.* [6]—have confirmed deep learning's promise for plant disease spotting without confronting cannabis-specific demands. The HAMIS-DLE model, purpose-built for cannabis pathologies, thus stands as a trailblazing entry within this niche.

Our investigation not only speaks to a pressing exigency in cannabis horticulture but also erects a platform upon which future explorations in this thinly charted territory can build. By assembling a dedicated dataset and unveiling an original

ensemble architecture, we have moved the needle on early malady detection, crop loss abatement, and the progression of agri-technological tools. The reach of our findings extends beyond cannabis alone, offering instructive templates and tactical lessons for disease governance across a spectrum of crops and thereby enriching the wider disciplines of agricultural technology and plant pathology.

Efficiency of the triple-ensemble model in LMDC classification

What the triple-ensemble architecture accomplishes—by fusing ensemble methodologies across augmentation, segmentation, and convolutional neural network (CNN) stages, and by placing the Artificial Multiple Intelligence system (AMIS) in charge of decision fusion—represents a breakthrough in the effort to manage Leaf Manifestation Diseases of Cannabis (LMDC). Both the conceptual blueprint and its execution map directly onto the progress markers and unresolved tensions that the cannabis cultivation and plant pathology scholarship, as cited throughout this paper, has brought to the fore.

The augmentation component, built around rotation, flipping, and colour-based transformations, plays an indispensable role in equipping the model to handle the multitude of postures and symptom displays cannabis leaves present in the field. Such variability is precisely what Sack *et al.* [10] and Short Gianotti *et al.* [11] identify as intensifying under commercial-scale production, where pathogen loads climb. Our augmentation regimen is configured expressly to inoculate the model against these sources of variation—a precondition for classification that can be trusted.

At the segmentation stage, bringing together U-Net, Mask R-CNN, and DeepLabv3+ amounts to a deliberate answer to the call for high-fidelity disease demarcation on cannabis foliage. The published record reinforces this priority: investigations by Ritharson *et al.* [23] and Pandian *et al.* [35], though focused on other crops, illuminate the vital importance of precise detection and signal the untapped potential of deep learning in this arena.

The CNN architectures in the ensemble—ResNet-50, Inception-ResNetV2, and EfficientNet-B3—were deliberately chosen, each selected to shoulder a specific dimension of the classification burden. Assembling a set of architecturally dissimilar networks echoes the positions staked out by Jain and Jaidka [8] and Asif *et al.* [22], who both make the case for AI-driven solutions to complex sorting problems in agricultural and medical domains

Entrusting decision fusion to AMIS signals the model's most distinctive design choice. M-AMIS, by tuning the fusion weights, dramatically improves how outputs from different architectural traditions are blended—a principle examined in the discussions by Langa *et al.* [24] and Sutton *et al.* [25]. Because of this fusion layer, the model does more than perform individual subtasks well; it reaches a collective equilibrium that the performance figures make plain.

Seen in its entirety, the triple-ensemble structure, steered by the AMIS fusion logic, marks a genuine methodological departure for automated LMDC sorting, one that meets head-on the shortcomings and blind spots catalogued in existing research. By deploying sophisticated image processing pipelines and integrating them into a carefully engineered decision framework, our study advances agricultural technology and plant pathology, delivering practical tools for cannabis disease control while laying the groundwork for broader crop disease surveillance.

Policy and business implications

The work presented here, centered on a triple-ensemble model for classifying Leaf Manifestation Diseases of Cannabis (LMDC), has implications that ripple through both regulatory thinking and commercial calculus across the cannabis and agritech sectors. From a policy standpoint, what the study makes plain is the imperative for updated governance structures that can absorb and regulate the infusion of AI and advanced machine learning tools into everyday farming. Robust frameworks of this kind are a prerequisite for codifying disease management standards that hold across the cannabis industry. In parallel, the research foregrounds the necessity of crafting rules that safeguard data confidentiality and prescribe ethically sound AI use, matters of particular sensitivity in the context of agricultural data.

On the business side, the implications cut just as deep. Bringing this AI-driven model from concept to operational reality opens new territory for agritech firms aiming to deliver next-generation disease management platforms. For cannabis growers, such tools can translate into a tangible edge: healthier plants, larger harvests, and higher-grade product—all feeding directly into stronger bottom lines. The evidence assembled here also points to the room for expanded investment in agritech R&D, especially in AI and deep learning applications purpose-fitted for crop systems. Investment of this sort could broaden the scope of commercial offerings to include intelligent disease-detection platforms that serve not only cannabis but also a broad portfolio of cultivated species.

Additionally, this research encourages the formation of collaborative networks that link technology developers, agronomic researchers, and practicing cultivators. Out of such partnerships can emerge AI solutions tailored to the specific demands of cannabis production, fuelling innovation cycles and commercial momentum across the sector. At the bottom, the study does more than deepen scientific understanding of cannabis disease management; it lays the policy and business groundwork for transformations in agricultural practice. The larger objective remains advancing cultivation methods that marry profitability with ecological responsibility, yielding dividends for producers and the broader community alike.

Limitations and future work

Limitations

Given the capacity the Hybrid Adaptive Multiple Intelligence System for Deep Learning Ensembles (HAMIS-DLE) has demonstrated in distinguishing cannabis leaf diseases, several boundary conditions warrant clear acknowledgement. The dataset that powered this investigation, thorough as it is, may nonetheless fail to represent every cannabis leaf disease variant found across the globe. Geographic idiosyncrasies in how symptoms express themselves—mediated by local climate, soil profiles, and grower practices—could erode the model's ability to generalize beyond the contexts from which the training data were drawn. In a related vein, the model's stability across a wide range of environmental conditions—light, humidity, and temperature—was not exhaustively tested. These are factors that can leave a heavy imprint on image quality and, by extension, on the reliability of the classifications produced.

A second limitation concerns the computational effort required by both training and inference for HAMIS-DLE. Such resource intensity risks putting the system beyond practical reach for smaller farms or regions where advanced computing facilities are scarce. Furthermore, though the model has been housed within the Automated Cannabis Leaf Disease Classification System (A-CLDC-S), the platform's operational characteristics—most critically, its latency and stability under variable network conditions—have yet to undergo systematic stress testing.

Future work

To address the constraints identified and amplify the performance and reach of the HAMIS-DLE framework, multiple trajectories for subsequent inquiry present themselves. Broadening the dataset to encompass a richer diversity of cannabis leaf pathologies sourced from disparate geographies and captured under heterogeneous environmental regimes is a priority. Partnerships forged with global research bodies and farming institutions could accelerate the assembly of a dataset that is both more extensive and more truly representative. Fortifying the model's resilience through augmentation protocols that mimic a range of environmental conditions—varying light, moisture, and thermal conditions—holds promise for enhancing its adaptability. Beyond this, field-testing the model under authentic cultivation settings spanning multiple climates and agronomic traditions will yield indispensable evidence about how it behaves outside controlled benchmarks.

Enlarging the data pool to incorporate samples drawn from additional crop species marks another indispensable step. Doing so would provide proof of the system's flexibility and ruggedness when transplanted into diverse agricultural landscapes, highlighting its relevance well beyond the cannabis niche. By incorporating a range of crop diseases, the HAMIS-DLE construct can be validated for a broader set of farming applications, thereby magnifying its practical value and societal footprint.

The pursuit of strategies to streamline HAMIS-DLE's computational appetite is equally critical. Approaches such as model pruning, quantization, and the engineering of more parsimonious neural architectures can lower the resource threshold, making the platform more attainable for smallholder operations and regions with sparse computational infrastructure. A thoroughgoing appraisal of the A-CLDC-S platform's operational characteristics—especially response latency, scaling behavior, and user interaction quality across a spectrum of network conditions—is indispensable. Prospective upgrades might incorporate edge-computing paradigms to reduce latency and improve system responsiveness.

Delivering richer documentation on the system's real-world fielding also ranks as a significant objective. This includes cataloguing performance indicators across varied environmental contexts, thereby shoring up the practical credibility of our contribution. Trials conducted in authentic settings that span climatic and cultivation regimes will yield instructive data on model behavior and performance, helping ensure the system can genuinely assist growers operating under diverse conditions.

Although the present investigation concentrates on cannabis, the underlying principles and techniques are transferable to other crops. Subsequent studies should examine the tailoring of HAMIS-DLE for disease sorting in additional economically important species, thereby broadening its impact on agricultural practice. Embedding a feedback loop that allows end users to contribute observations within the A-CLDC-S ecosystem will facilitate continual model refinement. Parallel to this, instructional initiatives designed to train farmers and cultivators to use the platform effectively can increase both uptake rates and real-world impact. By systematically addressing acknowledged limitations and advancing along the research avenues outlined here, the HAMIS-DLE model can be further honed, its applicability extended, and, in the long run, contribute to farming systems that are both more sustainable and more productive.

Conclusion

The impetus behind our study was the pressing requirement for faithful identification of cannabis leaf ailments, pursued to upgrade crop stewardship and product quality. Present-day classification tools fall short, particularly when operating across fluctuating environmental conditions, highlighting a significant gap in the current methodological landscape. To close this gap, we compiled a dedicated cannabis-centric dataset and unveiled a novel classification apparatus.

For our Automated Cannabis Leaf Disease Classification System (A-CLDC-S), we embraced a two-pronged ensemble design. The system wove together an array of image augmentation and segmentation pipelines, guided by a non-population-based decision-fusion logic, thereby capturing a fuller portrait of disease presentation. In a subsequent step, a population-based decision fusion scheme was called upon to unify the outputs of several Convolutional Neural Network (CNN) backbones—ResNet-50, Inception-ResNetV2, and EfficientNet-B3 among them. This dual-strategy methodology is crystallized in our Hybrid Adaptive Multiple Intelligence System (Hybrid AMIS) for Deep Learning Ensembles (HAMIS-DLE), an architecture equipped to steer both categories of fusion through heuristic search mechanisms.

HAMIS-DLE delivered standout results, achieving 99.58% accuracy on our cannabis corpus and demonstrating its competence across pathologies in other plant species. By way of illustration, when evaluated against external benchmarks for rice and groundnut leaf diseases, HAMIS-DLE achieved accuracies of 99.81% and 99.62%, respectively, outstripping established baselines such as VGG16 and DenseNet 169. These outcomes speak to the model's dependability and its latent value for a broad swath of agricultural applications.

Embedding HAMIS-DLE within A-CLDC-S markedly improved the system's capabilities, as evidenced by the chatbot interface's accessible design and its real-time diagnostic output. The deployed A-CLDC-S registered a strong mean accuracy of 98.70%. Moreover, results from a System Usability Scale (SUS) survey indicated high user satisfaction, reflecting the platform's ease of use and practical utility. Such feedback is instrumental in steering ongoing refinements and corroborating the system's tangible worth in farming contexts.

Seen as a whole, the fusion of HAMIS-DLE into A-CLDC-S yielded exceptional diagnostic precision, surpassing legacy approaches and confirming usability through favourable user appraisals. These findings highlight the potency of deep learning methodologies for disease classification, with encouraging implications for advancing crop health governance and agricultural output.

Looking ahead, subsequent investigations could productively focus on refining and operationalizing automated leaf disease classification platforms in real agricultural environments. This ambition might encompass widening the system's remit to cover a broader range of crops and ecological settings, integrating real-time surveillance tools such as unmanned aerial vehicles and IoT sensor networks, and reducing computational demands to expand accessibility. By tackling these dimensions, future endeavours aim to advance the on-the-ground deployment of automated disease management solutions, nurturing resilient, sustainable agricultural practices.

Acknowledgements

The authors would like to thank the National Science, Research and Innovation Fund and Maha Sarakham University for their valuable financial support of this research.

Conflict of interest

None

Financial support

This research was supported by the National Science, Research and Innovation Fund and Maha Sarakham University. The funding agencies had no role in the design, collection, analysis, or interpretation of the data, writing of the manuscript, or decision to submit the manuscript for publication.

Ethics statement

Dear editor-in-chief,

I am writing to submit our manuscript titled “Hybrid Adaptive Multiple Intelligence System (HybridAMIS) for Classifying Cannabis Leaf Diseases using Deep Learning Ensembles” for consideration for publication in *Smart Agricultural Technology*. This study investigates. The study addresses the critical need for precise identification of cannabis leaf diseases, aiming to enhance crop management and quality. Our developed system, the Hybrid AMIS for Deep Learning Ensembles (HAMIS-DLE), achieved an outstanding accuracy of 99.58 % on our cannabis dataset, demonstrating its effectiveness in disease classification. The innovative methodology combines various image augmentation and segmentation methods with multiple Convolutional Neural Network (CNN) architectures, including ResNet-50, Inception-ResNetV2, and EfficientNet-B3, using heuristic mechanisms to manage both strategy types comprehensively, which we believe will be of great interest to your journal's readership.

Received: 10 Dec 2023

Revised: 28 Jan 2024

Accepted: 27 Mar 2024

Published online: 15 September 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Mehmood N, Saeed M, Zafarullah S, Hyder S, Rizvi ZF, Gondal AS, et al. Multifaceted impacts of plant-beneficial *Pseudomonas* spp. in managing various plant diseases and crop yield improvement. *ACS Omega*. 2023;8(25):22296-315.
- Bakulova IV, Pluzhnikova IV, Kriushin NV. The Influence of agricultural practices on photosynthetic activity and productivity of a new variety of cannabis in the middle Volga region. *Agrarian Sci*. 2023;7:80-4.
- Luyckx M, Blanquet M, Isenborghs A, Guerriero G, Bidar G, Waterlot C, et al. Impact of silicon and heavy metals on hemp (*Cannabis sativa* L.) bast fibres properties: an industrial and agricultural perspective. *Int J Environ Res*. 2022;16(5):82.
- Elhag MM. Selected farming systems for improving crop production and rainwater productivity in Semi-Arid Zone, Sudan. *Open Access J Agric Res*. 2023;8(3):1-12.
- Aggarwal M, Khullar V, Goyal N, Singh A, Tolba A, Thompson EB, et al. Pre-trained deep neural network-based features selection supported machine learning for rice leaf disease classification. *Agriculture*. 2023;13(5):936.
- Amin H, Darwish A, Hassanien AE, Soliman M. End-to-end deep learning model for corn leaf disease classification. *IEEE Access*. 2022;10:31103-15.
- Dixit S, Kumar A, Haripriya A, Bohre K, Srinivasan K. Classification and recognition of soybean leaf diseases in Madhya Pradesh and Chhattisgarh using deep learning methods. In: *Proceedings of the 2nd International Conference on Paradigm Shifts in Communications Embedded Systems, Machine Learning and Signal Processing (PCEMS)*; 2023. p. 1-6.
- Jain S, Jaidka P. Mango leaf disease classification using deep learning hybrid model. In: *Proceedings of the International Conference on Power, Instrumentation, Energy and Control (PIECON)*; 2023. p. 1-6.
- Hall W, Stjepanović D, Caulkins J, Lynskey M, Leung J, Campbell G, et al. Public health implications of legalising the production and sale of cannabis for medicinal and recreational use. *Lancet*. 2019;394(10208):1580-90.
- Sack C, Simpson C, Pacheco K. The emerging spectrum of respiratory diseases in the U.S. cannabis industry. *Semin Respir Crit Care Med*. 2023;44(03):405-14.
- Short Gianotti AG, Harrower J, Baird G, Sepaniak S. The quasi-legal challenge: assessing and governing the environmental impacts of cannabis cultivation in the North Coastal Basin of California. *Land Use Policy*. 2017;61:126-34.
- Jerushalmi S, Maymon M, Dombrovsky A, Freeman S. Fungal pathogens affecting the production and quality of medical cannabis in Israel. *Plants*. 2020;9(7):882.
- Punja ZK. Emerging diseases of cannabis sativa and sustainable management. *Pest Manag Sci*. 2021;77(9):3857-70.
- Malabadi RB, Nethravathi T, Kolkar KP, Chalannavar RK, Mudigoudra BS, Abdi G, et al. Fungal infection diseases- nightmare for cannabis industries: artificial intelligence applications. *Int J Res Innov Appl Sci*. 2023;VIII(VIII):111-36.
- Punja ZK, Collyer D, Scott C, Lung S, Holmes J, Sutton D. Pathogens and molds affecting production and quality of cannabis sativa L. *Front Plant Sci*. 2019;10:1120.

Akanmu AO, Babalola OO, Venturi V, Ayilara MS, Adeleke BS, Amoo AE, et al. Plant disease management: leveraging on the plant-microbe-soil interface in the biorational use of organic amendments. *Front Plant Sci.* 2021;12:700507.

Tariq M, Khan A, Asif M, Khan F, Ansari T, Shariq M, et al. Biological control: a sustainable and practical approach for plant disease management. *Acta Agric Scand Sect B — Soil Plant Sci.* 2020;70(6):507-24.

Barber DM, Jackson VE, Ling KB. Innovations in crop protection towards sustainable agriculture. *Pest Manag Sci.* 2023;79(12):4754-6.

Mishra J, Prakash J, Arora NK. Role of beneficial soil microbes in sustainable agriculture and environmental management. *Clim Change Environ Sustain.* 2016;4(2):137.

Nanthasamroeng N, Khonjun S, Srichok T, Pitakaso R, Akkararungroengkul R, Jirasirilerd G, et al. Transborder logistics network design for agricultural product transportation in the Greater Mekong Subregion. *Asian J Shipping Logist.* 2022;38(4):245-62.

Sethanan K, Pitakaso R, Srichok T, Khonjun S, Thannipat P, Wanram S, et al. Double AMIS-ensemble deep learning for skin cancer classification. *Expert Syst Appl.* 2023;234:121047.

Asif S, Zhao M, Tang F, Zhu Y, Zhao B. Metaheuristics optimization-based ensemble of deep neural networks for mpox disease detection. *Neural Netw.* 2023;167:342-59.

Ritharson PI, Raimond K, Mary XA, Robert JE, J A. DeepRice: a deep learning and deep feature based classification of rice leaf disease subtypes. *Artif Intell Agric.* 2024;11:34-49.

Langa S, Magwaza LS, Mditshwa A, Tesfay SZ. Characterization of cannabis varieties and the intrinsic and extrinsic factors affecting cannabis germination and seedling establishment: a descriptive review. *Ind Crops Prod.* 2024;208:117861.

Sutton DB, Punja ZK, Hamarneh G. Characterization of trichome phenotypes to assess maturation and flower development in cannabis Sativa L. (Cannabis) by automatic trichome gland analysis. *Smart Agric Technol.* 2023;3:100111.

Sethy PK, Barpanda NK, Rath AK, Behera SK. Deep feature based rice leaf disease identification using support vector machine. *Comput Electron Agric.* 2020;175:105527.

MP A, Reddy P. Ensemble of CNN models for classification of groundnut plant leaf disease detection. *Smart Agric Technol.* 2023;6:100362.

Pitakaso R, Nanthasamroeng N, Srichok T, Khonjun S, Weerayuth N, Kotmongkol T, et al. A novel artificial multiple intelligence system (AMIS) for agricultural product transborder logistics network design in the Greater Mekong Subregion (GMS). *Computation.* 2022;10(7):126.

Wan X, Zuo X, Zhao X. A differential evolution algorithm combined with linear programming for solving a closed loop facility layout problem. *Appl Soft Comput.* 2022;121:108725.

He Z, Liu T, Liu H. Improved particle swarm optimization algorithms for aerodynamic shape optimization of high-speed train. *Adv Eng Softw.* 2022;173:103242.

Pitakaso R, Sethanan K, Jamrus T. Hybrid PSO and ALNS algorithm for software and mobile application for transportation in ice manufacturing industry 3.5. *Comput Ind Eng.* 2020;144:106461.

Bari BS, Islam MN, Rashid M, Hasan MJ, Razman MAM, Musa RM, et al. A real-time approach of diagnosing rice leaf disease using deep learning-based Faster R-CNN framework. *PeerJ Comput Sci.* 2021;7:e432.

Pandian JAKK, Rajalakshmi NR, Arulkumaran G. An improved deep residual convolutional neural network for plant leaf disease detection. *Comput Intell Neurosci.* 2022;2022:1-9.

Roy AM, Bhaduri J. A deep learning enabled multi-class plant disease detection model based on computer vision. *AI*. 2021;2(3):413-28.

Roy AM, Bhaduri J. Real-time growth stage detection model for high degree of occultation using DenseNet-Fused YOLOv4. *Comput Electron Agric*. 2022;193:106694.

Saberi Anari M. A hybrid model for leaf diseases classification based on the modified deep transfer learning and ensemble approach for agricultural AIoT-based monitoring. *Comput Intell Neurosci*. 2022;2022:1-15.