Technol TIMES, 2026
https://www.timesjournalsuite.com/

Research Article

A Decentralized, Edge-AI Hydroponic System for Sustainable Urban Agriculture

Muhammad Saad Eldin and Amir Akramin Shafie*
Department of Mechatronics Engineering, Kulliyyah of Engineering,
International Islamic University Malaysia, Kuala Lumpur, Malaysia

Received: Feb 07, 2026
Accepted: Mar 12, 2026
Published: Mar 17, 2026

***Corresponding author:**
Amir Akramin Shafie
E-mail:
aashafie@iium.edu.my

ABSTRACT

Global pressures from urbanization, shrinking arable land, and climate volatility necessitate efficient food production methods. This paper presents an Intelligent Hydroponic System (IHS) designed for autonomy and resource efficiency, integrating edge-based TinyML for closed-loop environmental control with a computer-vision layer for periodic plant-health diagnostics. The system employs an ESP32 microcontroller to host a lightweight Multi-Layer Perceptron (MLP) model that processes real-time sensor data—pH, electrical conductivity/total dissolved solids (EC/TDS), temperature, and humidity—to automate nutrient dosing and LED lighting. A secondary diagnostic module uses an NVIDIA Jetson Nano running a YOLO-based object detection model for twice-daily visual health assessments. In an eight-week Nutrient Film Technique (NFT) trial with lettuce (Lactuca sativa) and mustard greens (Brassica juncea), the IHS maintained target nutrient parameters with 96.3% accuracy and achieved 18–21% higher fresh biomass compared to a conventional, manually managed hydroponic system. The system also demonstrated a 21% reduction in water consumption and improved nutrient-use efficiency from 72% to 87%. Automated LED and pump scheduling further reduced total energy consumption by 8%. With a core smart component cost of approximately RM 244.25, the IHS offers a scalable, low-cost, and connectivity-resilient solution for precision agriculture in urban and resource-limited settings. These results establish a compelling case for decentralized, edge-native intelligence as a practical alternative to cloud-dependent smart farming platforms.

KEYWORDS: Hydroponics; Precision Agriculture; Edge Computing; Internet of Things (IoT); Embedded Machine Learning; Computer Vision; Sustainable Farming; TinyML; Nutrient Film Technique

1. INTRODUCTION

The global population is projected to reach 9.7 billion by 2050, intensifying pressure on traditional soil-based agriculture amid constraints on arable land and freshwater resources (FAO, 2021). By 2100, more than two-thirds of the world's population is expected to reside in urban environments, fundamentally reshaping where and how food must be produced (UN DESA, 2018). Controlled-environment agriculture (CEA), and hydroponics in particular, presents a compelling pathway toward resilient, localized food production, offering water savings of up to 90% compared to conventional field farming (Hoagland & Arnon, 1950; Asao, 2012). Hydroponic systems achieve these efficiencies by delivering nutrients directly to plant roots in a controlled solution, eliminating soil-related constraints such as nutrient leaching and variable water retention.

Despite their inherent advantages, hydroponic systems have struggled to achieve widespread adoption in urban and small-scale settings. Two primary barriers persist: the labor-intensive nature of manual monitoring and dosing, and the latency, cost, and connectivity dependencies inherent in many cloud-centric 'smart farm' solutions. Manual systems require skilled operators to periodically measure and adjust pH, electrical conductivity (EC), and other nutrient parameters—a demanding routine that leaves wide windows for suboptimal conditions to develop and harm plant health. Cloud-based alternatives, while increasingly capable, introduce significant concerns around network reliability, data privacy, subscription costs, and real-time response latency, rendering them impractical for low-infrastructure urban deployments.

Recent advances in embedded machine learning—collectively referred to as TinyML—have opened a new design space for autonomous agricultural systems. The maturation of microcontrollers with sufficient computational resources (e.g., dual-core ESP32), alongside efficient model compression techniques such as quantization and pruning, now enables the deployment of neural network inference directly on low-cost hardware at the edge (Shareef et al., 2024). This paradigm eliminates dependence on cloud connectivity for real-time decisions, reduces latency to milliseconds, and substantially lowers hardware cost. Concurrently, compact single-board computers such as the NVIDIA Jetson Nano have made on-site computer vision inference feasible at accessible price points.

The convergence of these technologies motivates the system presented in this paper. The Intelligent Hydroponic System (IHS) is a decentralized, edge-native platform that integrates a TinyML-driven environmental control loop with a YOLO-based computer vision pipeline for plant health diagnostics. By co-locating intelligence with the physical system, the IHS achieves full operational autonomy with only intermittent internet connectivity—a critical feature for urban rooftops, community gardens, and resource-limited settings in developing economies. This design philosophy extends and synthesizes prior work in IoT-based hydroponic monitoring (Rathnayake et al., 2023), machine learning for precision agriculture (Adapa et al., 2024; Shareef et al., 2024), and smartphone-assisted plant diagnostics (Van et al., 2019).

The primary contributions of this work are fourfold: (1) the design and implementation of a fully integrated, edge-native hardware-software architecture for autonomous hydroponics; (2) the development and on-device deployment of a quantized MLP model for closed-loop environmental control; (3) the integration of a YOLO-based vision pipeline with complementary colorimetric sensing for redundant plant health monitoring; and (4) a validated comparative trial demonstrating significant gains in biomass yield, plant health, water use efficiency, and nutrient use efficiency relative to conventional manual management.

2. LITERATURE REVIEW

2.1 Hydroponic Systems and Controlled Environment Agriculture

Hydroponics encompasses a range of soilless cultivation techniques, of which the Nutrient Film Technique (NFT), deep water culture (DWC), and ebb-and-flow systems are among the most widely studied. NFT, first systematized by Cooper (1979), circulates a thin film of nutrient solution over plant roots, providing high oxygenation and efficient nutrient delivery. Asao (2012) provides a comprehensive review of hydroponic methodologies in controlled research settings, highlighting NFT's suitability for leafy greens due to its scalability and low water volume requirements. López-Chuken (2012) further documents the application of hydroponics to phytoremediation, underscoring the versatility of soilless systems beyond food production.

Within CEA, environmental parameters—particularly pH, EC/TDS, temperature, humidity, and light intensity—exert decisive influence over growth rates and nutritional quality. Optimal pH ranges for most leafy crops fall between 5.8 and 6.5, while EC targets vary by species and growth stage. Deviation from these setpoints, even for brief periods, can induce nutrient lockout, osmotic stress, or accelerated pathogen growth, resulting in measurable yield penalties. These sensitivities underscore the value of continuous, automated monitoring and rapid corrective actuation.

2.2 IoT and Automation in Hydroponic Systems

The application of Internet of Things (IoT) architectures to hydroponic monitoring has been extensively explored over the past decade. Rathnayake et al. (2023) presented ADVANCE, an automated indoor hydroponic unit that integrates environmental sensing with automated lighting and water management, reporting reduced user intervention without sacrificing yield. Van et al. (2019) introduced PlantTalk, a smartphone-based IoT platform for intelligent hydroponic management published in Sensors, establishing an early precedent for smartphone-driven diagnostics in small-scale hydroponic deployments.

A systematic literature review by Shareef et al. (2024) surveyed parameter optimization strategies for smart hydroponic systems across 84 studies, identifying pH and EC regulation as the most frequently targeted control objectives. The review identifies a prominent gap: the majority of intelligent hydroponic solutions rely on continuous cloud connectivity, leaving them vulnerable to network disruptions and imposing recurring bandwidth and service costs—constraints that are particularly prohibitive in low-income or rural-urban contexts.

2.3 TinyML and Edge Computing for Agricultural Applications

TinyML refers to the deployment of machine learning inference on microcontrollers and other severely resource-constrained hardware (Adapa et al., 2024). Techniques such as post-training quantization (reducing model weights from 32-bit floats to 8-bit integers), weight pruning, and knowledge distillation enable compact neural networks to fit within the kilobyte-scale memory budgets of modern microcontrollers. The ESP32, with its dual Xtensa LX6 cores, 520 KB of SRAM, and rich peripheral support, has emerged as a popular platform for TinyML agricultural applications. Deployed models can perform inference locally in under 50 ms, enabling closed-loop control cycles that are infeasible with cloud-roundtrip architectures.

In agricultural contexts, TinyML has been applied to soil moisture prediction, pest detection, and microclimate classification with promising results. However, its integration into a full-stack hydroponic control system—combining sensor-driven ML inference, actuator control, and vision-based diagnostics—has not been comprehensively reported in the peer-reviewed literature prior to this work.

2.4 Computer Vision for Plant Health Monitoring

Computer vision has been increasingly adopted for non-destructive plant health assessment, leveraging color analysis, morphological feature extraction, and deep learning classifiers to identify nutrient deficiencies, disease symptoms, and growth anomalies. YOLO (You Only Look Once) architectures, originally introduced for general object detection, have been successfully adapted for plant phenotyping tasks owing to their real-time inference speed and competitive accuracy on small training datasets when pre-trained weights are fine-tuned (Adapa et al., 2024). The integration of visual diagnostics with traditional physicochemical sensors creates a multi-modal feedback loop that is demonstrably more robust than either modality alone, as visual phenotypic changes—such as leaf chlorosis or wilting—often precede or accompany shifts in sensor chemistry.

3. MATERIALS AND METHODS

3.1 System Overview and Architecture

The IHS follows a modular architecture comprising four integrated components: (1) a sensor-actuator loop for real-time environmental data acquisition and actuation; (2) an ESP32-based control hub hosting the embedded MLP inference engine; (3) an NVIDIA Jetson Nano vision module for periodic plant health diagnostics; and (4) a lightweight MQTT-based IoT communication hub for telemetry logging and user alerts. The physical cultivation substrate utilizes the Nutrient Film Technique (NFT) with recirculating nutrient solution flowing through modular PVC channels. NFT was selected for its scalability, efficient water recirculation, and suitability for the target crop species.

The layered architecture ensures operational resilience: the ESP32 control loop functions fully offline, while the Jetson Nano communicates diagnostic results to a local MQTT broker. User-facing alerts are published to both a local dashboard and, when connectivity is available, to a remote monitoring endpoint. This design decouples critical control functions from network availability, a key distinction from cloud-dependent alternatives.

3.2 Hardware Implementation

3.2.1 Computational Core. The dual-processor hardware setup reflects a deliberate partitioning of computational roles. The ESP32-

WROOM-32 microcontroller was selected for its low cost (approximately RM 15), dual-core Xtensa LX6 processing at 240 MHz, 520 KB SRAM, and integrated Wi-Fi/Bluetooth. It interfaces with the sensor suite over I2C and UART buses, performs ML inference, and drives actuators through GPIO-controlled relays. The NVIDIA Jetson Nano (4 GB variant) provides the GPU compute required for real-time neural network inference on camera frames, running the YOLO-based vision pipeline on its 128-core Maxwell GPU.

3.2.2 Sensor Suite. Environmental sensing is accomplished by a SEN0161 analog pH probe (±0.1 pH accuracy, 0–14 pH range), a DFRobot TDS/EC probe (±2% full-scale accuracy), a DHT11 module for ambient temperature (±2°C) and relative humidity (±5% RH), and a GY-31 TCS3200 color sensor positioned beneath the canopy for chlorophyll proxy estimation via leaf reflectance. The GY-31 provides spectral data in red, green, and blue channels, enabling quantitative tracking of leaf color as an early proxy for nitrogen or iron deficiency (chlorosis) prior to observable changes in the nutrient solution chemistry.

3.2.3 Actuation System. Nutrient and pH buffer dosing is handled by a set of precision peristaltic pumps (dosing accuracy ±0.5 mL), selected to minimize cross-contamination and enable precise micro-dosing. Full-spectrum LED grow lights are controlled by solid-state relays (SSRs) driven by the ESP32, enabling PWM-based dimming and programmatic photoperiod management. A submersible pump recirculates the nutrient solution through the NFT channels on a configurable duty cycle.

3.2.4 Vision Module. The NVIDIA Jetson Nano is paired with a Logitech C922 Pro webcam (1080p at 30 fps) for image capture. To ensure color consistency and reduce background clutter, a controlled illumination enclosure with diffuse white LED panels was constructed around the camera field of view. Images are captured every 12 hours, with timestamps and inference results logged to the local MQTT broker.

Table 1. Bill of materials for the IHS smart electronic components

Component	Model / Specification	Key Parameter	Unit Cost (RM)
ESP32 MCU	ESP32-WROOM-32	Dual-core 240 MHz, 520 KB SRAM	~15.00
pH Sensor	DFRobot SEN0161	±0.1 pH accuracy, 0–14 pH	~45.00
TDS/EC Probe	DFRobot TDS Probe	±2% accuracy, 0–1000 ppm	~30.00
Temp/Humidity	DHT11	±2°C / ±5% RH	~5.00
Color Sensor	GY-31 TCS3200	RGB + clear channel	~8.00
Peristaltic Pumps	Mini Dosing Pump ×3	±0.5 mL accuracy	~36.00
SSR Relay Module	4-ch SSR	5 V control, 240 V AC load	~12.00
LED Grow Lights	Full-spectrum LED Panel	380–780 nm, 45 W	~55.00
NVIDIA Jetson Nano	Jetson Nano Dev Kit (4 GB)	128-core GPU, 4 GB LPDDR4	~N/A (lab)
Webcam	Logitech C922 Pro	1080p @ 30 fps	~N/A (lab)
Total (Smart Components)	—	—	~244.25

3.3 AI and Software Layers

3.3.1 Environmental Control Model (MLP on ESP32)

The first AI layer is a TinyML-based Multi-Layer Perceptron (MLP) deployed directly on the ESP32. The model was trained on 1,200 expert-labeled records collected during pilot runs, where an agronomist manually classified system states (sensor vector → corrective action). Each record consists of four input features: pH, TDS (ppm), ambient temperature (°C), and relative humidity (%). The output layer classifies among six discrete control actions: (1) dose nutrient solution A, (2) dose nutrient solution B, (3) dose pH-up buffer, (4) dose pH-down buffer, (5) adjust photoperiod (extend or reduce LED on-time), and (6) no action required.

The architecture consists of an input layer (4 neurons), two hidden layers (32 and 16 neurons, ReLU activation), and a 6-class softmax output. Post-training int8 quantization was applied using TensorFlow Lite, reducing model size to approximately 8 KB—well within the ESP32's 520 KB SRAM. The quantized model was converted to a C array and compiled directly into the ESP32 firmware using the TensorFlow Lite Micro runtime. The median on-device inference time was measured at less than 20 milliseconds, enabling a practical control loop frequency of up to 50 Hz.

3.3.2 Computer Vision Health Diagnostic Pipeline (YOLO on Jetson Nano)

The second AI layer is a YOLO-based object detection model running on the Jetson Nano. A training dataset of 1,033 labeled images was curated from the pilot trial, augmented through horizontal/vertical flipping, brightness jitter (±30%), and Gaussian blur to improve model generalization across lighting conditions and growth stages. Images were annotated with bounding boxes and one of two class labels: 'Healthy' or 'Stressed'. Stress indicators captured in the training data include chlorosis (yellowing), tip burn, wilting, and dark spotting associated with over-fertilization.

The model was fine-tuned from pre-trained YOLOv5s weights using transfer learning, achieving convergence in approximately 80 training epochs. Inference is triggered every 12 hours by a cron job on the Jetson Nano; results are published to the MQTT broker within 5 minutes of image capture. A stress detection event triggers an immediate alert and prompts the ESP32 to increase diagnostic sensor polling frequency from 15-minute to 5-minute intervals.

3.3.3 Colorimetric Redundancy Layer

The GY-31 color sensor provides a lightweight, continuous (5-minute polling interval) colorimetric proxy for early chlorosis detection. By tracking the ratio of red to green channel output from the crop canopy, the system generates a normalized Greenness Index (NGI) defined as G / (R + G + B). A decline in NGI below a calibrated threshold (species-specific, established during pilot runs) triggers an early warning flag even before the YOLO vision model detects visible phenotypic changes. This redundant, multi-modal feedback mechanism creates a critical early-warning capability that neither sensor data nor vision analysis could achieve independently.

3.4 Experimental Design

The system was evaluated over a continuous eight-week growth cycle under standard NFT conditions. The study employed a two-arm comparative design: the fully autonomous IHS and a Conventional Hydroponic System (CHS) managed by daily manual monitoring and adjustment by an experienced operator. Both systems used identical NFT channel configurations with the following environmental setpoints: pH 6.0–6.4, TDS 1000–1100 ppm, EC 1.4–1.6 mS/cm, temperature 21–23°C, and relative humidity 55–65%. Both trials were conducted simultaneously in the same controlled laboratory space to minimize environmental confounds.

Each system cultivated two crop species: lettuce (Lactuca sativa, 'Butterhead' variety) and mustard greens (Brassica juncea, 'Green Wave' variety), chosen for their relevance to urban food security, short growth cycles, and well-characterized hydroponic performance benchmarks. Each species was grown in three replicate NFT channels per system (six channels per arm). Primary endpoints measured were: (1) fresh biomass at harvest (g per plant); (2) the fraction of plants classified as healthy at trial end; (3) water consumption per growth cycle (L); (4) nutrient-use efficiency (NUE, proportion of supplied nutrients absorbed); and (5) total energy consumption (kWh per cycle).

4. RESULTS

4.1 AI Model Performance

The embedded MLP achieved a classification accuracy of 96.3% on the held-out test set (n = 240 records, stratified by class), with a median on-device inference time of 17.4 ms (range: 14–23 ms). Confusion matrix analysis revealed that the most frequent misclassifications occurred between the 'dose pH-up' and 'no action' classes at pH values near the upper setpoint boundary (pH 6.3–6.4), reflecting the inherent ambiguity of marginal states. No misclassifications occurred between chemically opposing actions (e.g., pH-up vs. pH-down), confirming the model's safety-critical reliability.

The YOLO-based vision pipeline attained a mean Average Precision at 50% IoU (mAP50) of 0.994 on the held-out test set, with precision and recall values of 0.991 and 0.988 for the 'Stressed' class, respectively. Colorimetric stress alerts from the GY-31 layer preceded YOLO-confirmed stress events by a median of 14 hours, validating the NGI as a meaningful early-warning signal. End-to-end time from stress onset to system alert was less than 5 minutes across all trials.

Table 2. Summary of AI model performance metrics

Model	Metric	Value
MLP (ESP32)	Test Accuracy	96.3%
MLP (ESP32)	Median Inference Time	17.4 ms
MLP (ESP32)	Model Size (quantized)	~8 KB
YOLOv5s (Jetson Nano)	mAP50 (Healthy / Stressed)	0.994
YOLOv5s (Jetson Nano)	Precision (Stressed)	0.991
YOLOv5s (Jetson Nano)	Recall (Stressed)	0.988
GY-31 Colorimetric	Median Lead Time before YOLO Alert	14 hours
System (end-to-end)	Alert Response Latency	< 5 min

4.2 Growth, Yield, and Resource Efficiency

The IHS consistently outperformed the CHS across all primary endpoints. Lettuce plants in the IHS reached a mean fresh biomass of 148 g per plant compared to 124 g in the CHS—a 19.4% increase. Mustard greens showed an even more pronounced improvement, reaching a mean of 172 g per plant in the IHS versus 131 g in the CHS (31.3% increase). Aggregated across both species, the IHS achieved an 18–21% higher fresh biomass at harvest. Plant health rates at trial end were 93% (IHS) versus 78% (CHS), with the 15-percentage-point difference primarily attributable to earlier detection and remediation of stress events in the IHS.

Resource efficiency improvements were equally substantial. Water consumption per growth cycle was reduced by approximately 21% in the IHS, attributable to tighter EC/TDS control reducing the frequency of solution replacement. Nutrient-use efficiency improved from 72% (CHS) to 87% (IHS), reflecting more precise dosing and reduced nutrient loss to solution dump events. Automated LED and pump scheduling resulted in an 8% reduction in total energy consumption per cycle.

Table 3. Comparative performance summary: IHS vs. CHS across key metrics

Metric	IHS (Intelligent)	CHS (Conventional)	Improvement
Lettuce fresh biomass (g/plant)	148 ± 12	124 ± 15	+19.4%
Mustard greens biomass (g/plant)	172 ± 14	131 ± 18	+31.3%
Healthy plants at harvest (%)	93%	78%	+15 pp
Water consumption (L/cycle)	Reduced by ~21%	Baseline	−21%
Nutrient-use efficiency (%)	87%	72%	+15 pp
Energy consumption (kWh/cycle)	Reduced by ~8%	Baseline	−8%
Smart component cost (RM)	~244.25	N/A (manual)	Low-cost

4.3 Environmental Parameter Stability

Throughout the eight-week trial, the IHS maintained pH within the target range (6.0–6.4) for 97.8% of monitored time intervals (15-minute sampling), compared to 84.3% for the CHS. TDS was maintained within 1000–1100 ppm for 96.1% of intervals in the IHS versus 79.5% in the CHS. These stability improvements reflect the cumulative advantage of continuous automated actuation over the discrete, daily correction intervals inherent in manual management. Notably, the IHS avoided all instances of acute pH excursion (>6.6 or <5.8), three of which were recorded in the CHS and correlated temporally with the stress events observed in that arm.

4.4 Cost Analysis

The total bill of materials for all smart electronic components—microcontrollers, sensors, actuators, relays, and grow lights—was RM 244.25 (approximately USD 52). This figure excludes the NVIDIA Jetson Nano and webcam, which were sourced from laboratory inventory; for a commercial deployment, the Jetson Nano (4 GB) retails at approximately USD 149, placing total smart component cost at approximately RM 944. This compares favorably with turnkey commercial hydroponic automation systems, which typically range from USD 500 to over USD 2,000 for comparable functionality. The open-source software stack and commodity hardware ensure that all components are replaceable and upgradable without vendor lock-in.

5. DISCUSSION

5.1 Edge Intelligence as an Enabler for Autonomous Agriculture

The results confirm that a decentralized, edge-native intelligence architecture can effectively overcome the principal limitations of cloud-dependent smart farming systems. By processing sensor data and executing control decisions locally on the ESP32, the IHS eliminates the latency and reliability risks associated with cloud roundtrips. The median inference time of 17.4 ms enables a control loop cadence that is orders of magnitude faster than the periodic manual adjustments in the CHS and significantly faster than most cloud-connected alternatives, which are bounded by network round-trip times of hundreds to thousands of milliseconds.

This speed advantage translates directly into tighter parameter stability. The 13.5-percentage-point improvement in pH maintenance (IHS: 97.8% vs. CHS: 84.3%) reflects the system's ability to detect and correct drift within minutes rather than hours. Given that even brief excursions outside the optimal pH range can induce nutrient lockout and measurable yield penalties, this improvement in stability is mechanistically consistent with the observed biomass and health gains. The absence of acute pH excursions in the IHS, contrasted with three such events in the CHS, provides additional direct evidence that autonomous edge control materially reduces risk of crop loss.

5.2 Multi-Modal Diagnostics and the Value of Sensor Fusion

A central innovation of the IHS is the fusion of three distinct sensing modalities: electrochemical sensors (pH, EC/TDS), ambient environment sensors (temperature, humidity), and visual/colorimetric sensing (camera + GY-31). Each modality captures a distinct dimension of plant and system state, and their combination creates a feedback loop that is substantially more informative than any single modality alone. The 14-hour median lead time of the colorimetric NGI signal ahead of YOLO-confirmed stress events is particularly significant: it demonstrates that the GY-31 layer provides actionable early warning that enables preventive—rather than reactive—intervention.

This multi-modal architecture also provides robustness against individual sensor failures. In scenarios where the pH probe drifts out of calibration (a documented failure mode for electrochemical sensors in hydroponic environments), the vision and colorimetric layers continue to provide independent indicators of plant stress, alerting operators to investigate the sensor rather than misattributing declining plant health to a non-chemical cause.

5.3 Generalizability and Scalability Considerations

The current validation is limited to two leafy green species (lettuce and mustard greens) in a single controlled laboratory environment. While these species are agronomically representative of the most common urban hydroponic crops, the trained MLP and YOLO models are species- and environment-specific in their current form. Extending the system to fruiting crops (e.g., tomatoes, strawberries) or herbs with substantially different nutrient requirements, morphologies, and stress phenotypes will require retraining or fine-tuning of both models with species-specific data.

The modular hardware architecture is well-suited to scaling. Additional NFT channels can be added without modifying the control software, subject to actuator capacity. The MQTT broker supports many-to-one telemetry from multiple ESP32 nodes, enabling a single Jetson Nano vision module to serve multiple cultivation zones through a scheduling rotation. These design features position the IHS as a scalable platform for small-to-medium commercial deployments, not merely a research prototype.

5.4 Comparison with Related Work

The IHS advances the state of the art in several respects relative to prior published systems. Compared to PlantTalk (Van et al., 2019), which requires a smartphone for vision inference and maintains cloud connectivity for control decisions, the IHS achieves comparable visual diagnostic capability with a fully offline control loop and no smartphone dependency. Relative to Rathnayake et al. (2023), who demonstrated automated hydroponic control but relied on basic sensor thresholds without ML inference, the IHS provides more adaptive closed-loop control, achieving 96.3% classification accuracy entirely on-device. The integration of redundant colorimetric sensing alongside vision diagnostics appears to be a novel contribution not reported in any directly comparable prior system.

5.5 Limitations and Future Work

Several limitations constrain the generalizability of the current findings. The training datasets for both the MLP (1,200 records) and YOLO (1,033 images) models are relatively small, and the MLP dataset was collected from a single laboratory environment by a single expert labeler. Larger, more diverse datasets—including records from multiple operators, environments, and seasonal conditions—would strengthen model robustness. The eight-week trial, while spanning a full growth cycle for the target species, represents a single temporal replicate; multi-season and multi-site validation is necessary to confirm the stability of the observed performance gains.

Future development will focus on four priorities: (1) expanding crop coverage to fruiting plants, herbs, and microgreens through few-shot learning and transfer learning techniques; (2)

deploying federated learning protocols to aggregate model updates from multiple IHS installations without centralizing raw data; (3) integrating a low-cost hyperspectral or multispectral imaging module to enhance the diagnostic precision of the vision layer for specific nutrient deficiency classification; and (4) conducting a full lifecycle cost-benefit analysis across diverse urban deployment scenarios to quantify return on investment for small-scale commercial growers.

6. CONCLUSION

This study successfully developed and validated a low-cost, decentralized Intelligent Hydroponic System combining embedded TinyML inference on an ESP32 microcontroller with a YOLO-based vision pipeline and colorimetric redundancy layer on an NVIDIA Jetson Nano. In an eight-week comparative trial with lettuce and mustard greens, the IHS demonstrated significant and consistent improvements over conventional manual management: 18–21% higher fresh biomass, 93% vs. 78% healthy plant rates, 21% lower water consumption, 15 percentage-point gain in nutrient-use efficiency, and 8% reduction in energy use—all at a smart component cost of approximately RM 244.25.

The system's decentralized, edge-native architecture provides full operational autonomy without internet connectivity, addressing the key practical barriers that have limited the adoption of smart hydroponic solutions in urban and resource-constrained settings. The multi-modal sensing architecture, combining electrochemical, ambient, colorimetric, and visual data streams, delivers a level of diagnostic redundancy and early-warning capability not achievable by sensor data or vision analysis alone. Together, these properties establish the IHS as a practical, accessible, and scalable platform for the advancement of sustainable precision agriculture.

By providing complete architectural documentation, model training procedures, and deployment instructions, this work aims to facilitate replication and adaptation by researchers and practitioners, contributing an open and reproducible engineering foundation for the next generation of intelligent urban food systems.

7. REFERENCES

Adapa R, Singh HS, Kaur P and Khan MI (2024) A comprehensive survey of machine learning techniques in smart agriculture. Journal of Agriculture and Food Research 26:100222. DOI: 10.1016/j.jafr.2024.100222.

Asao T (ed) (2012) Hydroponics – A standard methodology for plant biological researches. InTech. DOI: 10.5772/2215.

Cooper AJ (1979) The ABC of NFT: Nutrient film technique. Grower Books, London.

FAO (2021) The state of food and agriculture 2021: Making agrifood systems more resilient to shocks and stresses. FAO. DOI: 10.4060/cb4476en.

Hoagland DR and Arnon DI (1950) The water-culture method for growing plants without soil. California Agricultural Experiment Station Circulars 347:1–39.

López-Chuken UJ (2012) Hydroponics and environmental clean-up. In: T Asao (ed) Hydroponics – A standard methodology for plant biological researches, pp. 181–198. InTech.

Rathnayake RMUI, Liyanage DN, Ganege HC and Rathnathunga EUU (2023) Advance: An automated indoor hydroponic unit for plant growth detection with compatible for two different plant varieties. Indonesian Journal of Innovation and Applied Sciences 3(2):118–132. DOI: 10.47540/ijias.v3i2.807.

Shareef U, Rehman AU and Ahmad R (2024) A systematic literature review on parameters optimization for smart hydroponic systems. AI 5(3):1517–1533. DOI: 10.3390/ai5030073.

UN DESA (2018) World urbanization prospects: The 2018 revision. United Nations. Available at: https://population.un.org/wup/.

Van L-D, Lin Y-B, Wu T-H, Lin Y-W, Peng S-R, Kao L-H and Chang C-H (2019) PlantTalk: A smartphone-based intelligent hydroponic plant box. Sensors 19(8):1763. DOI: 10.3390/s19081763.

Author Contributions

All the authors conceived the concept, wrote and approved the manuscript.

Acknowledgements

Not applicable.

Funding

Not applicable.

Availability of data and materials

Not applicable.

Competing interest

The authors declare no competing interests.

Ethics approval

Not applicable.

Open Access

Citation: Muhammad Saad Eldin and Amir Akramin Shafie (2026) A Decentralized, Edge-AI Hydroponic System for Sustainable Urban Agriculture. Technol TIMES 1 (1): 23-29.