

**INTEGRATED SMART WATER-FOCUSED
IRRIGATION SYSTEM USING IOT AND AI/ML**
(Water Forecasting & Decision-Support System)

Project ID: 25-26J-520

Abishek S
(Dilruksha A.G.C.D., Hesara P.K.A.N., Trishni W.R.M.)

B.Sc. (Hons) in Information Technology Specializing in Software
Engineering

Department of Software Engineering

Sri Lanka Institute of Information Technology Malabe, Sri Lanka

April 2026

**INTEGRATED SMART WATER-FOCUSED
IRRIGATION SYSTEM USING IOT AND AI/ML**

(Water Forecasting & Decision-Support System)

Project ID: 25-26J-520

Abishek S - IT22186942

Trishni W.R.M. - IT22076366

Hesara P.K.A.N. - IT22561398

Dilruksha A.G.C.D. - IT22561770

Supervisor: Ms. Hansi De Silva

Co-Supervisor: Ms. Karthiga Rajendran

B.Sc. (Hons) in Information Technology Specializing in Software
Engineering

Department of Software Engineering

Sri Lanka Institute of Information Technology Malabe, Sri Lanka

April 2026

DECLARATION

I declare that this is my own work, and this Thesis does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any other University or institute of higher learning and to the best of my knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgement is made in the text.

Also, I hereby grant to Sri Lanka Institute of Information Technology, the nonexclusive right to reproduce and distribute my Thesis, in whole or in part in print, electronic or other medium. I retain the right to use this content in whole or part in future works (such as articles or books).

Name	Student ID	Signature
Abishek S	IT22186942	

The above candidate has carried out this research thesis for the Degree of Bachelor of Science (honors) Information Technology (Specializing in Software Engineering) under my supervision.

Signature of the Supervisor

Date

Signature of the Co-supervisor

Date

ABSTRACT

Sri Lanka's agricultural sector, contributing approximately 7.4% to GDP and employing over 25% of the national workforce, faces persistent challenges from crop diseases, nutrient deficiencies, and water-related stresses that collectively result in substantial annual yield losses. Traditional crop health monitoring relies predominantly on manual field inspections — a methodology that is inherently reactive, spatially limited, and capable of detecting phytopathological issues only after visible symptoms manifest, thereby reducing the effectiveness of interventions. Satellite-based remote sensing systems have improved coverage and monitoring frequency; however, they remain disconnected from localised ground-level validation, limiting operational reliability. Convolutional Neural Network (CNN) models offer high accuracy in disease classification from field photographs but cannot scale spatially without satellite guidance. This disconnect between macro-scale spectral analysis and micro-level image verification constitutes the central challenge addressed by this research.

This research proposes the design, development, and validation of a Hybrid AI and Satellite-Based Crop Health Monitoring System — constituting the remote-sensing and computer-vision intelligence layer of the broader Integrated Smart Water-Focused Irrigation System (Project ID: 25-26J-520). The proposed system operates through a three-stage pipeline: Stage 1 automatically ingests Sentinel-2 Level-2A multispectral imagery via Google Earth Engine APIs and computes vegetation indices (NDVI, NDWI, MSI, EVI) with crop-specific calibration to detect spatial caution zones representing early stress signatures; Stage 2 delivers notifications to farmers through a Flutter mobile application, which accepts geo-tagged field photographs classified by an ensemble of locally trained CNN architectures (MobileNetV3, ResNet50, EfficientNet-B3, Custom CNN) executing on AWS EC2 GPU infrastructure; Stage 3 fuses spectral and ground-level outputs in a React-and-Mapbox web dashboard, producing colour-coded health maps, time-series vegetation trends, prioritised multi-channel alerts, and exportable agronomic reports for irrigation officers. An active-learning feedback loop retrains CNNs periodically on farmer-validated uploads, ensuring continuous domain adaptation.

Expected outcomes include detection of crop stress up to seven days before visible symptoms, classification accuracy exceeding 90% for dominant Sri Lankan crop diseases, over 15% improvement relative to generic global models, sub-five-second CNN inference latency, and support for at least 1,000 concurrent users. The system establishes a scalable, auditable, and commercially extensible precision-agriculture framework suitable for Sri Lanka and analogous tropical agrarian economies.

Keywords: *Crop Health Monitoring, Sentinel-2, Convolutional Neural Network, Vegetation Index, NDVI, NDWI, Satellite Remote Sensing, Google Earth Engine, Precision Agriculture, Active Learning, IoT, Smart Irrigation*

ACKNOWLEDGEMENTS

The successful completion of this individual research component would not have been possible without the guidance, support, and encouragement of numerous individuals and institutions to whom I am deeply grateful.

First and foremost, I extend my sincere gratitude to my supervisor, Ms. Hansi De Silva, whose expert guidance, constructive feedback, and unwavering encouragement throughout every phase of this research have been invaluable. I am equally thankful to my co-supervisor, Ms. Karthiga Rajendran, whose technical insights and domain expertise in AI and agricultural systems significantly enriched the design and methodology of this work.

I would like to thank the academic staff of the Department of Software Engineering at Sri Lanka Institute of Information Technology (SLIIT), whose lectures, workshops, and research culture fostered the intellectual foundation upon which this study is built. The library resources and cloud computing facilities provided by SLIIT have been essential infrastructure for this work.

I am grateful to my group members — Dilruksha A.G.C.D. (IT22561770), Hesara P.K.A.N. (IT22561398), and Trishni W.R.M. (IT22076366) for their collaborative spirit, technical discussions, and mutual support throughout the development of the Integrated Smart Water-Focused Irrigation System. The complementary modules each member contributed have collectively elevated the quality and scope of this research project.

Sincere thanks are due to the Department of Agriculture, Sri Lanka, and the Hector Kobbekaduwa Agrarian Research and Training Institute (HARTI) for their support in facilitating access to agricultural datasets and field validation opportunities. The International Water Management Institute (IWMI) is also acknowledged for invaluable research resources related to irrigation and remote sensing in South Asian agriculture.

Finally, I wish to express my deepest appreciation to my family for their unconditional love, patience, and moral support throughout my academic journey. This work is dedicated to them.

Table of Contents

DECLARATION	3
ABSTRACT.....	4
ACKNOWLEDGEMEN.....	6
LIST OF TABLES	9
LIST OF FIGURES	10
LIST OF ABBREVIATIONS	11
1 CHAPTER 1: INTRODUCTION	14
1.1 Background Study and Literature Review	14
1.1.1 Background Study.....	14
1.1.2 Literature Review.....	16
1.2 Research Gap	19
1.3 Research Problem	21
1.4 Research Objectives	22
1.4.1 Main Objective.....	22
1.4.2 Specific Objectives	22
1.4.3 Business Objectives	23
2 CHAPTER 2: METHODOLOGY	24
2.1 Methodology	24
2.1.1 Feasibility Study and Planning	24
2.1.2 Feasibility Assessment.....	24
2.1.3 Requirement Gathering and Analysis	28
2.1.4 Designing	31
2.1.5 Implementation	36
2.1.6 Testing.....	39
2.1.7 Deployment and Maintenance	41
2.2 Commercialisation	43
3 CHAPTER 3: RESULTS AND DISCUSSION.....	46
4 CHAPTER 4: FUTURE SCOPE	49
5 CHAPTER 5: CONCLUSION	51
6 REFERENCES	53

LIST OF TABLES

Table 1.1:Summary of Related Work in Crop Health Monitoring	19
Table 2.1:Feasibility Assessment Summary	25
Table 2.2:Estimated Project Cost Breakdown	26
Table 2.3:Risk Management Plan	27
Table 2.4:Communication Plan	28
Table 2.5:Functional Requirements	30
Table 2.6:Non-Functional Requirements	30
Table 2.7:Data Requirements.....	31
Table 2.8:Technology Stack Summary	38
Table 2.9:Test Cases for Core Features	41

LIST OF FIGURES

Figure 2.1:Three-Stage Hybrid Architecture Overview	32
Figure 2.2:System Overview Diagram (Individual Module).....	33
Figure 2.3:Use Case Diagram	34
Figure 2.4:Sequence Diagram.....	35
Figure 2.5:Data Flow Diagram	36
Figure 2.6:Deployment Architecture	42

LIST OF ABBREVIATIONS

Abbreviation	Description
AI	Artificial Intelligence
API	Application Programming Interface
AWS	Amazon Web Services
CBSL	Central Bank of Sri Lanka
CI/CD	Continuous Integration / Continuous Deployment
CNN	Convolutional Neural Network
CSV	Comma-Separated Values
DB	Database
DoA	Department of Agriculture
ESA	European Space Agency
EVI	Enhanced Vegetation Index
FAO	Food and Agriculture Organization
FDR	Frequency Domain Reflectometry
GIS	Geographic Information System
GNDVI	Green Normalized Difference Vegetation Index
GPS	Global Positioning System
GPU	Graphics Processing Unit
HARTI	Hector Kobbekaduwa Agrarian Research and Training Institute
ICT	Information and Communication Technology
IoT	Internet of Things

Abbreviation	Description
IWMI	International Water Management Institute
ML	Machine Learning
MSI	Moisture Stress Index
MQTT	Message Queuing Telemetry Transport
NDVI	Normalized Difference Vegetation Index
NDWI	Normalized Difference Water Index
OGC	Open Geospatial Consortium
PDPA	Personal Data Protection Act
PWA	Progressive Web Application
RBAC	Role-Based Access Control
REST	Representational State Transfer
ROI	Return on Investment
SAR	Synthetic Aperture Radar
SAVI	Soil-Adjusted Vegetation Index
SLIIT	Sri Lanka Institute of Information Technology
SMS	Short Message Service
SUS	System Usability Scale
TDR	Time Domain Reflectometry
TLS	Transport Layer Security
UAT	User Acceptance Testing
UI/UX	User Interface / User Experience

Abbreviation	Description
WCAG	Web Content Accessibility Guidelines

1 CHAPTER 1: INTRODUCTION

1.1 Background Study and Literature Review

1.1.1 Background Study

Agriculture forms the bedrock of Sri Lanka's socioeconomic structure. According to the Central Bank of Sri Lanka (2023), the agricultural sector contributes approximately 7.4% to the national Gross Domestic Product and employs over 25% of the national workforce, with the paddy, vegetable, and fruit cultivation sectors serving as the primary sources of rural livelihoods across all nine provinces. The Uda Walawe and Anuradhapura irrigation zones alone support hundreds of thousands of smallholder farming families whose incomes and food security are directly dependent on the health and productivity of their crops across two main cultivation seasons - the Yala (May-August) and Maha (September-March) seasons.

Despite the sector's economic significance, Sri Lankan agriculture faces systemic challenges from crop diseases, pest infestations, and abiotic stresses - principally water stress and nutrient deficiency - that collectively cause substantial annual yield losses. The Department of Agriculture (2023) estimates that preventable crop diseases and suboptimal management practices account for significant economic damage each year, particularly within water-intensive cultivation systems. These losses are exacerbated by the reactive nature of existing monitoring and management practices, in which farmers and agricultural extension officers rely predominantly on manual field scouting to identify phytopathological and physiological issues. Such methods are inherently labour-intensive, spatially limited, and - critically - identify crop stress only after visible symptoms have manifested at advanced stages, by which time intervention is considerably less effective and far more resource-intensive.

The inadequacy of traditional monitoring is further compounded by Sri Lanka's diverse agroecological zones. The island encompasses distinct climatic regions - from the wet zone in the south-west to the dry zone of the north and east, the intermediate zone of the central highlands, and the arid zone of the north-west - each characterised by different rainfall regimes, soil types, crop varieties, and pest-disease profiles. This spatial heterogeneity means that a monitoring approach effective in one zone may be entirely inadequate in another, and that comprehensive national coverage through manual scouting would require agricultural officer deployments at a scale that is organisationally and financially impractical.

Remote sensing technologies, and specifically satellite imagery from the European Space Agency's Sentinel-2 mission within the Copernicus programme, have emerged as transformational tools for automated, scalable, and objective crop health monitoring. Sentinel-2 provides freely accessible, high-resolution multispectral data with global coverage and a five-day revisit interval, capturing critical electromagnetic spectrum bands - visible red, green, blue, red-edge, and near-infrared - that are instrumental for vegetation analysis. These multispectral datasets can be processed to compute vegetation indices such as the Normalized Difference Vegetation Index (NDVI), which serves as a robust proxy for chlorophyll content and photosynthetic activity; the Normalized Difference Water Index (NDWI), which indicates plant water status and hydration stress; the Moisture Stress Index (MSI), sensitive to canopy water content; and the Enhanced Vegetation Index (EVI), which provides improved sensitivity in high-biomass regions. The availability of time-series imagery enables tracking of vegetation dynamics throughout crop growth cycles, facilitating early anomaly detection and trend forecasting that are impossible through sporadic ground observations.

Artificial Intelligence, particularly deep Convolutional Neural Networks (CNNs), has demonstrated remarkable accuracy in detecting specific crop diseases, nutrient deficiencies, and stress patterns from both spectral signatures and high-resolution ground images. International research has established that AI-driven satellite analysis can detect water stress, fungal infections, and viral diseases well before visible symptoms become apparent to human observers, creating critical windows for preventive interventions. However, the disconnect between macro-scale satellite monitoring and micro-level ground verification remains the fundamental challenge in translating remote-sensing potential into operational agricultural decision-support tools in Sri Lanka.

In this context, the present research proposes a three-stage Hybrid AI and Satellite-Based Crop Health Monitoring System that integrates automated Sentinel-2 satellite analytics with geo-tagged ground image collection and CNN-based verification, supported by a Flutter mobile application and a centralised web dashboard. This system constitutes the remote-sensing and computer-vision intelligence layer of the Integrated Smart Water-Focused Irrigation System (Project ID: 25-26J-520) developed at SLIIT, complementing the reservoir forecasting, automated field irrigation, and crop-area optimisation modules developed by other group members.

1.1.2 Literature Review

A comprehensive review of existing literature reveals a rich body of research across the intersecting domains of satellite-based vegetation monitoring, deep learning for crop disease classification, mobile agricultural advisory systems, and hybrid remote-sensing architectures. The following synthesis draws upon key contributions that inform the design and positioning of the proposed system.

Liakos et al. (2018) conducted a foundational survey of machine learning applications in precision agriculture, documenting that ML-based approaches had achieved classification accuracies between 85% and 98% across diverse crop-disease datasets when trained on sufficient, well-labelled data. Their work established that CNNs consistently outperformed traditional classifiers such as SVMs and decision trees when processing image data, and highlighted the challenge of generalising models trained on temperate-region datasets to tropical agroecological contexts - a limitation directly addressed by the locally trained CNN ensemble proposed in the present research [1].

Mohanty et al. (2016) demonstrated the potential of deep learning for plant disease identification using the PlantVillage dataset, achieving accuracy rates of over 99% in controlled conditions but identifying significant performance degradation - dropping to around 31% - when models were evaluated under field conditions with variable lighting, orientation, and background complexity. Their findings underscore the necessity of training CNN models on locally collected, geo-tagged field imagery rather than relying on curated laboratory datasets - a principle embedded in the data-collection methodology of the proposed system [2].

Kamilaris and Prenafeta-Boldú (2018) reviewed 40 deep learning studies in agriculture and concluded that while CNNs offered strong performance in disease detection, their applicability was constrained by the scarcity of large, annotated datasets for tropical crops. They proposed active-learning frameworks as a viable strategy for iterative dataset expansion through verified expert annotations - a concept that directly inspired the active-learning feedback loop in the proposed architecture, wherein farmer-validated classifications are used to periodically retrain the CNN ensemble [3].

Sentinel-2-based vegetation monitoring has been extensively investigated for agricultural applications. Frampton et al. (2013) evaluated the utility of Sentinel-2 spectral bands for vegetation analysis and demonstrated that the combination of red-edge bands with near-infrared and shortwave-infrared bands provides superior discrimination of canopy condition

compared to earlier Landsat and SPOT systems. Their band analysis supports the multi-index computation approach (NDVI, NDWI, MSI, EVI) adopted in Stage 1 of the proposed pipeline [4].

Segarra et al. (2020) reviewed remote sensing applications in crop monitoring and established that NDVI time-series analysis could detect water stress up to seven days before the appearance of visual symptoms in paddy and maize - a finding that forms the empirical basis for the early-warning capability claimed by the proposed system. Their work also demonstrated that EVI provides improved sensitivity in high-biomass regions where NDVI can become saturated, justifying the multi-index approach rather than reliance on NDVI alone [5].

Rahman and Zhou (2024) proposed a hybrid AI framework for integrating satellite and field imagery in tropical agricultural environments, demonstrating that the combination of spectral anomaly detection with ground-image classification significantly reduced false-positive stress alerts by approximately 25% compared to single-source satellite detection methods. Their fusion approach - spatiotemporal alignment of satellite tile coordinates with GPS-tagged field images - directly informs the spatiotemporal data fusion strategy implemented in the Integration Layer of the proposed architecture [6].

Kawshalya and De Silva (2024) conducted a study specific to Sri Lankan paddy cultivation, evaluating CNN-assisted verification of Sentinel-2 vegetation anomalies in the Uda Walawe irrigation zone. Their findings indicated that locally trained CNN models achieved classification accuracy approximately 15–18% higher than globally pre-trained architectures when applied to Sri Lankan paddy varieties under local lighting and soil conditions — a result that directly motivates the locally trained ensemble in the proposed system and forms the basis for the 15% accuracy improvement claim [7].

Li and Zhou (2023) investigated deep learning-based multispectral data fusion for crop stress detection and demonstrated that EfficientNet-B3 achieved the best trade-off between classification accuracy and computational efficiency among standard CNN architectures evaluated on multispectral crop imagery, outperforming both ResNet50 (higher accuracy but greater computational cost) and MobileNetV3 (lower computational cost but reduced accuracy on complex stress patterns). Their comparative analysis informs the ensemble composition — combining all three architectures to leverage their complementary strengths — proposed in the AI Analytics Layer [8].

Howard et al. (2019) introduced MobileNetV3 as a lightweight neural architecture optimised for mobile deployment, demonstrating that its inverted residual structure and hard-swish activation achieved competitive accuracy on image classification tasks with significantly reduced parameter counts and inference latency compared to VGG and ResNet architectures. Their work validates the inclusion of MobileNetV3 as the mobile-deployment component of the CNN ensemble, ensuring that lite inference can be performed on low-cost Android devices during field operations where connectivity is limited [9].

He et al. (2016) introduced the ResNet architecture with residual connections that enable training of very deep networks by addressing vanishing gradient problems, achieving state-of-the-art accuracy on ImageNet. ResNet50 has since become a standard backbone in transfer learning for agricultural image classification, providing robust feature extraction for complex disease patterns and establishing the basis for its inclusion in the proposed CNN ensemble [10].

An et al. (2022) reviewed mobile-based agricultural advisory systems in developing countries and identified that user adoption was strongly correlated with multilingual interface availability, offline functionality, and response latency below ten seconds. Their user-experience guidelines directly informed the design specifications of the Flutter mobile application - particularly the requirements for Sinhala, Tamil, and English multilingual support; offline caching and synchronisation; and sub-five-second AI feedback - embedded in the proposed system's non-functional requirements [11].

International Water Management Institute (2023) published a working paper on remote sensing applications in Sri Lankan agriculture, documenting case studies from the Uda Walawe irrigation zone and confirming that vegetation-index monitoring supported by farmer-validated ground data significantly improved irrigation scheduling accuracy. Their findings support the integration of the proposed crop-health monitoring output with the irrigation scheduling module within the parent project [12].

Author(s)	Focus	Key Finding / Relevance
Liakos et al. (2018)	ML survey in precision agriculture	CNN superiority over SVM/DT; tropical generalisation gap
Mohanty et al. (2016)	PlantVillage deep learning	Need for locally collected field datasets

Author(s)	Focus	Key Finding / Relevance
Kamilaris & Prenafeta-Boldú (2018)	DL review in agriculture	Active-learning for dataset expansion
Frampton et al. (2013)	Sentinel-2 band analysis	Red-edge + NIR bands for canopy discrimination
Segarra et al. (2020)	Remote sensing crop monitoring	NDVI detects water stress 7 days before symptoms
Rahman & Zhou (2024)	Hybrid satellite-field AI	25% false-positive reduction vs. single source
Kawshalya & De Silva (2024)	CNN for Sri Lankan paddy	15-18% accuracy gain with local training
Li & Zhou (2023)	Multispectral data fusion DL	EfficientNet-B3 best accuracy-efficiency tradeoff
Howard et al. (2019)	MobileNetV3 architecture	Lightweight inference for mobile deployment
He et al. (2016)	ResNet residual connections	Standard backbone for transfer learning
An et al. (2022)	Mobile advisory systems	Multilingual, offline, sub-10s response criteria
IWMI (2023)	Remote sensing in Sri Lanka	VI monitoring improves irrigation scheduling

Table 1.1: Summary of Related Work in Crop Health Monitoring

1.2 Research Gap

The literature review reveals significant technological and implementation gaps in realising a hybrid crop health monitoring framework suitable for Sri Lankan agricultural conditions. These gaps operate at three levels: methodological, technological, and contextual.

At the methodological level, the prevailing research dichotomy between satellite-based and ground-level monitoring approaches represents the most critical gap. Existing satellite-based

crop health monitoring systems - both international platforms and the limited local implementations documented for Sri Lanka - predominantly rely on remote sensing analytics in isolation, generating vegetation indices such as NDVI and NDWI to indicate stress levels. While these systems achieve broad spatial coverage at acceptable temporal resolution, they provide only indicative outputs without ground-level validation mechanisms to confirm the nature, cause, and severity of detected anomalies. As Kawshalya and De Silva (2024) demonstrated, generic satellite-derived stress flags in tropical agricultural settings generate substantial false-positive rates that erode farmer and officer trust in automated alerts. Conversely, CNN-based ground verification systems achieve high classification accuracy but are operationally constrained to field-level application by a single farmer or extension officer, lacking the spatial scalability necessary for regional monitoring. No existing system in Sri Lanka closes this loop by linking satellite detection directly to mobile-enabled ground verification with bidirectional data exchange.

At the technological level, a critical gap exists in the integration of satellite-derived insights with localised, AI-driven image verification through a continuous feedback loop. While Rahman and Zhou (2024) demonstrated the viability of spatiotemporal fusion of satellite and ground data in tropical contexts, their work was limited to controlled research environments and did not address the operational requirements - multilingual interfaces, offline functionality, active-learning retraining - necessary for field-scale deployment in Sri Lanka. The absence of locally trained CNN models for Sri Lankan crop varieties is a further gap: the majority of deep learning models evaluated in prior literature were trained on PlantVillage and similar temperate-region datasets, and their direct application to Sri Lankan paddy, maize, and vegetable varieties under local illumination conditions and soil backgrounds produces significantly degraded accuracy, as confirmed by Mohanty et al. (2016) and Kawshalya and De Silva (2024). Closed active-learning pipelines that leverage farmer validations as a continuous training signal - with full provenance tracking for auditability - have not been implemented in any documented Sri Lankan agricultural technology context.

At the contextual level, existing mobile advisory systems in Sri Lanka depend on manual image uploads or farmer text inputs without any linkage to satellite-derived remote sensing indicators. This architectural separation between space-based detection and on-ground verification limits accuracy, timeliness, and user trust. Furthermore, the absence of multilingual decision-support interfaces in Sinhala, Tamil, and English represents a significant inclusivity gap that limits adoption among the majority of Sri Lanka's smallholder farming population. Existing GIS-

based agricultural monitoring dashboards in Sri Lanka do not incorporate CNN-confidence-scored diagnoses or active-learning feedback mechanisms, and do not provide exportable agronomic reports aligned with the reporting requirements of provincial irrigation departments.

The present research directly addresses all three levels of the gap. The three-stage hybrid architecture bridges the satellite-ground disconnect through spatiotemporal fusion; the locally trained CNN ensemble eliminates the temperate-dataset bias; the active-learning feedback loop transforms farmer validations into a continuous improvement signal; and the multilingual Flutter application and React dashboard provide inclusive, operationally practical interfaces for Sri Lankan agricultural stakeholders.

1.3 Research Problem

The central research problem addressed in this study may be formally articulated as follows:

How can a hybrid framework combining satellite-based multispectral analysis and AI-driven ground verification be designed, developed, and validated to detect, confirm, and manage crop stress conditions early and accurately within Sri Lanka's agricultural systems, whilst providing actionable, location-specific recommendations to farmers and irrigation officers through inclusive, scalable digital interfaces?

Traditional crop health monitoring approaches in Sri Lanka remain predominantly manual, reactive, and fragmented. Physical field inspections by extension officers are time-consuming, subjective in assessment quality, and spatially constrained - critically, they typically identify phytopathological and physiological issues only after visible symptoms have manifested at advanced growth stages, thereby limiting the efficacy of interventions. Satellite-based remote sensing has improved spatial coverage and monitoring frequency, yet its accuracy is constrained by the absence of localised ground-level validation and contextual interpretation, rendering it indicative rather than operationally decisive for farm-level decision-making.

Meanwhile, ground-based CNN models offer high local classification accuracy but cannot scale across large agricultural regions without satellite guidance to direct field inspection efforts. This disconnect between macro-scale satellite detection and micro-level ground confirmation produces inefficiencies, misclassifications, and delayed decision-making that collectively contribute to preventable yield losses. Addressing this challenge requires the development of an integrated solution that merges Sentinel-2 satellite imagery analysis with

CNN-based verification of field photographs, supported by mobile and web interfaces enabling two-way data exchange between farmers and irrigation officers, and continuously improved through validated field inputs.

1.4 Research Objectives

1.4.1 Main Objective

To design, develop, and validate a Hybrid AI and Satellite-Based Crop Health Monitoring System that integrates multispectral satellite analytics (Sentinel-2) with Convolutional Neural Network (CNN)-driven ground verification to enable early, accurate, and validated detection of crop stress, nutrient deficiency, and disease conditions in Sri Lanka's major agricultural regions, and to provide actionable, location-specific recommendations to farmers and irrigation officers that enhance decision-making for irrigation management, disease control, and sustainable resource utilisation.

1.4.2 Specific Objectives

1. Develop automated data pipelines for acquiring, preprocessing, and managing Sentinel-2 Level-2A imagery, including cloud masking, atmospheric correction, and vegetation index computation (NDVI, NDWI, MSI, EVI) with crop-specific calibration coefficients and soil-adjusted formulations for early growth stages.
2. Design, train, and comparatively evaluate multiple CNN architectures — MobileNetV3, ResNet50, EfficientNet-B3, and a Custom CNN — for classification of crop diseases, water stress, and nutrient deficiencies using geo-tagged images collected from farmers and field officers across Sri Lankan agricultural zones.
3. Implement a Flutter-based cross-platform mobile application enabling farmers and agricultural officers to capture real-time, GPS-tagged crop images with metadata, receive satellite-detected caution zone notifications, and obtain instant AI feedback with recommended actions, with multilingual support and offline synchronisation.
4. Develop a centralised React-and-Mapbox web dashboard for irrigation officers that fuses satellite and AI outputs, displays colour-coded stress zones, tracks validation data, provides exportable reports, and supports query-based analysis for irrigation planning and disease management.

5. Establish an active-learning feedback loop that periodically retrains CNN models using farmer-validated field uploads with full provenance tracking, ensuring continuous domain adaptation across crop varieties and agroecological zones.
6. Integrate multi-channel alerting (SMS, push notifications, email) with configurable escalation mechanisms and severity-based prioritisation, and connect system outputs with the irrigation scheduling module of the parent project.

1.4.3 Business Objectives

7. Achieve a commercially sustainable, multi-tiered SaaS business model targeting smallholder farmers, government agricultural institutions, and commercial plantations, with affordable subscription tiers (LKR 500–2,000/month) delivering a demonstrated return on investment exceeding 3× for smallholder users and 5× for plantation clients.
8. Establish an AgriSentinel AI startup venture with institutional partnerships with the Department of Agriculture, HARTI, and telecommunications providers to achieve nationwide deployment coverage within three years of initial commercialisation, with expansion to South Indian and Bangladeshi markets in phase two.
9. Demonstrate measurable socioeconomic impact by reducing preventable crop yield losses by at least 15–25% per season in pilot deployment regions (Uda Walawe, Anuradhapura), reducing manual scouting effort by 40%, and achieving user acceptance scores above 80% on the System Usability Scale across all stakeholder groups.

2 CHAPTER 2: METHODOLOGY

2.1 Methodology

The research adopts an Agile iterative prototyping methodology structured across four overlapping development phases, emphasising continuous validation, stakeholder engagement, and incremental feature delivery. Agile was selected over waterfall and V-model approaches for several principled reasons specific to this domain. First, the requirements for a hybrid satellite-AI agricultural system evolve dynamically as stakeholder feedback from farmers, extension officers, and irrigation department officials is incorporated; fixed-requirement approaches would produce a system misaligned with operational realities. Second, the dependency on real agricultural seasons - Yala and Maha - means that iterative field validation must be synchronised with cultivation cycles, requiring flexible sprint planning. Third, the parallel development of satellite pipeline, CNN model training, mobile application, and web dashboard components benefits from concurrent development with frequent integration checkpoints rather than sequential phase completion. The four phases are: (1) System Foundation (Months 1–3): finalising functional requirements, configuring Sentinel-2 API ingestion, and establishing preprocessing pipelines; (2) AI Model Development (Months 4–7): CNN training, evaluation, and ensemble construction; (3) Mobile and Dashboard Integration (Months 8–10): Flutter app development, React dashboard integration, REST API connections; (4) Field Validation and Refinement (Months 11–14): pilot deployment, comparative trials against manual scouting, feedback-driven retraining, and documentation.

2.1.1 Feasibility Study and Planning

Prior to commencing development, a comprehensive feasibility assessment was conducted across four dimensions to confirm the viability of the proposed system within the project's resource, timeline, and regulatory constraints.

2.1.2 Feasibility Assessment

Dimension	Verdict	Rationale
Technical	FEASIBLE	Sentinel-2 and Google Earth Engine APIs are freely accessible. TensorFlow, PyTorch, and Flutter are mature, well-documented frameworks. AWS EC2 GPU instances provide adequate computational capacity for CNN training

Dimension	Verdict	Rationale
		and inference. SLIIT provides GPU-enabled computing laboratories for development and testing.
Operational	FEASIBLE	The Department of Agriculture and HARTI have confirmed willingness to support field data collection in Uda Walawe and Anuradhapura pilot zones. Agricultural extension officers have expressed interest in AI-assisted advisory tools. Farmer smartphone penetration in pilot zones is sufficient for mobile application deployment.
Financial	FEASIBLE	Sentinel-2 and Google Earth Engine are freely available, eliminating satellite data procurement costs. Open-source frameworks (TensorFlow, React, Flutter) minimise software licensing expenditure. Cloud costs are manageable within the project budget through pay-per-use AWS EC2 provisioning and Spot instance utilisation for training workloads.
Legal	FEASIBLE	Compliance with Sri Lanka's Personal Data Protection Act (PDPA) 2022 is addressed through end-to-end TLS encryption, RBAC access control, and anonymisation of farmer imagery. Copernicus open-access licence permits unrestricted use of Sentinel-2 data for research and commercial purposes. Open geospatial data standards (OGC WMS/WFS) are adopted for interoperability with government platforms.

Table 2.1: Feasibility Assessment Summary

Project Cost Breakdown:

Cost Item	Estimated Amount
AWS EC2 GPU Instances (NVIDIA T4, Spot)	LKR 45,000 (training phase)
AWS EC2 General Purpose (production)	LKR 18,000 / month
AWS S3 Storage (satellite tiles, imagery)	LKR 8,000 / month
MongoDB Atlas (cloud database)	LKR 6,500 / month

Cost Item	Estimated Amount
Google Earth Engine API (commercial tier)	LKR 12,000 / month
Flutter/React Development Tools	LKR 0 (open source)
TensorFlow / PyTorch	LKR 0 (open source)
Field Data Collection (field officer time)	LKR 25,000 (pilot phase)
SMS Alert Gateway (Tier-1 provider)	LKR 5,000 / month
Contingency Reserve (15%)	LKR 21,000
TOTAL ESTIMATED PROJECT COST	LKR 161,000 (Year 1)

Table 2.2: Estimated Project Cost Breakdown

Project Timeline — Gantt Chart Description:

The project timeline spans 14 months across four development phases. Phase 1 (Months 1–3) encompasses requirements finalisation, Sentinel-2 pipeline configuration, and preprocessing infrastructure establishment, producing a functional satellite ingestion system and validated vegetation-index generation scripts. Phase 2 (Months 4–7) focuses on ground-truth dataset collection, CNN architecture training and evaluation, and ensemble construction, producing comparative performance reports and an optimised AI inference pipeline. Phase 3 (Months 8–10) delivers Flutter mobile application development, React dashboard implementation, Mapbox integration, and REST API connection, producing functional mobile and web interfaces with integrated notification systems. Phase 4 (Months 11–14) encompasses field validation in Uda Walawe and Anuradhapura, comparative trials against manual scouting baselines, active-learning retraining, user-acceptance testing, and final documentation.

Risk Management Plan:

Risk	Likelihood	Impact	Mitigation Strategy
Sentinel-2 cloud cover exceeding 80% during Maha season	High	Medium	Supplement with Sentinel-1 SAR imagery for operational continuity; implement multi-temporal compositing

Risk	Likelihood	Impact	Mitigation Strategy
Insufficient geo-tagged ground imagery for CNN training	High	High	Partner with DoA extension officers for structured data collection campaigns; implement active-learning to maximise value of limited labels
AWS GPU instance availability during training peak	Medium	Low	Pre-book Reserved Instances for training windows; utilise Spot Instances with checkpointing to resume interrupted jobs
Farmer smartphone penetration below threshold in pilot zones	Medium	Medium	Validate penetration rates prior to pilot deployment; provide shared-device access points through HARTI offices as fallback
Model accuracy degradation under new crop seasons	Medium	Medium	Active-learning pipeline automatically retrains; schedule seasonal retraining sprints aligned with Yala/Maha cycles
PDPA compliance failures with geo-tagged imagery	Low	High	Implement anonymisation pipeline at collection point; conduct PDPA audit before production deployment

Table 2.3: Risk Management Plan

Communication Plan:

Activity	Participants	Channel	Frequency
Weekly Sprint Reviews	Project team + supervisors	Google Meet / Teams	Weekly

Activity	Participants	Channel	Frequency
Monthly Progress Reports	Supervisors Ms. De Silva, Ms. Rajendran	Email + SLIIT portal	Monthly
Field Coordination	DoA extension officers, HARTI	WhatsApp + field visits	As required
Stakeholder Demos	Irrigation dept., farmer groups	On-site presentations	End of each phase
GitHub Pull Requests	All developers	GitHub	Per feature completion

Table 2.4: Communication Plan

2.1.3 Requirement Gathering and Analysis

Requirements were gathered through a triangulated approach combining literature analysis, expert consultations with DoA agricultural officers and HARTI researchers, and structured field interviews with smallholder farmers and irrigation department officials across the Uda Walawe and Anuradhapura pilot zones. This multi-source approach ensured that technical, agronomic, and operational perspectives were captured and balanced in the final requirements specification.

Functional Requirements:

ID	Feature	Description
FR-01	Automated Satellite Acquisition	Continuously retrieve Sentinel-2 L2A imagery via Google Earth Engine APIs with automated cloud detection, atmospheric correction, and mosaicking at 5-day revisit intervals
FR-02	Vegetation Index Computation	Calculate NDVI, NDWI, MSI, EVI with crop-specific calibration coefficients and soil-adjusted formulations for early growth stages
FR-03	Caution Zone Detection	Identify spatial caution zones where vegetation indices deviate beyond configurable thresholds, with anomaly severity scoring

ID	Feature	Description
FR-04	CNN Ensemble Inference	Execute MobileNetV3, ResNet50, EfficientNet-B3, and Custom CNN classifiers on uploaded ground images for water stress, nutrient deficiency, and disease classification
FR-05	Ensemble Confidence Scoring	Generate confidence metrics and anomaly-severity scores for each CNN classification; apply ensemble averaging across three principal architectures
FR-06	Mobile Image Upload	Enable geo-tagged crop photograph upload via Flutter mobile app with GPS coordinates, timestamp, and contextual metadata (soil condition, irrigation status)
FR-07	Mobile Caution Zone Notification	Deliver satellite-detected caution zone notifications to farmers' mobile devices with location-specific field inspection instructions
FR-08	Mobile AI Feedback	Provide instant AI classification feedback and recommended actions to farmers within 5 seconds of image upload
FR-09	Dashboard Visualisation	Display colour-coded health maps, time-series vegetation trends, and multi-index layer toggles for irrigation officers via React-Mapbox web dashboard
FR-10	Alert Generation	Trigger SMS, push notification, and email alerts when VI deviations exceed thresholds, with prioritisation by crop type, phenological stage, and severity
FR-11	Report Export	Generate and export agronomic reports (PDF) and CSV datasets for administrative and research use
FR-12	Active-Learning Retraining	Integrate farmer validations into periodic CNN retraining pipeline with provenance tracking for all training data
FR-13	Administrative Console	Allow authorised officers to manage user roles, configure alert thresholds, monitor system performance, and review model accuracy logs

ID	Feature	Description
FR-14	Cross-Module Integration	Expose RESTful APIs for data exchange with the irrigation scheduling and reservoir forecasting modules of the parent project

Table 2.5: Functional Requirements

Non-Functional Requirements:

ID	Category	Requirement
NFR-01	Performance	Process full Sentinel-2 scenes within 2 hours of data availability; CNN batch inference latency under 5 seconds on GPU; dashboard load time under 3 seconds
NFR-02	Scalability	Support at least 1,000 concurrent users with cloud-native auto-scaling; horizontally scalable microservices without architecture redesign
NFR-03	Reliability	Maintain 99.5% uptime during peak cultivation periods; implement checkpoint recovery and backup mechanisms
NFR-04	Security	End-to-end TLS encryption; MFA authentication; RBAC for farmer, officer, and admin roles; PDPA 2022 compliance
NFR-05	Usability	Multilingual UI (Sinhala, Tamil, English); WCAG 2.1 AA accessibility; offline caching; sub-10-second user workflow for photo upload
NFR-06	Maintainability	Modular microservice architecture; comprehensive API documentation; automated CI/CD pipelines
NFR-07	Interoperability	OGC WMS/WFS and GeoTIFF/GeoJSON support; RESTful APIs for third-party integration with government portals
NFR-08	Cost Efficiency	Utilise free Sentinel data; open-source frameworks; pay-per-use cloud billing; storage compression to minimise idle costs

Table 2.6: Non-Functional Requirements

Data Requirements:

Data Type	Description and Source
Sentinel-2 L2A Multispectral Imagery	Primary satellite data; 10m/20m spatial resolution; 5-day revisit; sourced via Google Earth Engine APIs
Sentinel-1 SAR Imagery	Complementary radar backscatter data for cloud-affected periods; C-band SAR; sourced via Copernicus Hub
Geo-tagged Crop Photographs	Ground-truth images from farmers and officers; minimum 8MP resolution; GPS coordinates + timestamp + metadata
FDR/TDR Soil Moisture Data	Volumetric water content measurements for NDWI validation; from field sensors via MQTT IoT gateways
Micro-Weather Station Data	Rainfall, temperature, relative humidity; from local weather stations in pilot zones
FieldSpec Hyperspectral Calibration Data	Reflectance calibration from portable spectroradiometers for vegetation index accuracy validation
Crop Type and Phenological Stage Metadata	Farmer-provided or officer-recorded; used for calibration coefficient selection and alert prioritisation
Farmer Validation Responses	Confirmation or correction of AI classifications; primary input for active-learning retraining pipeline

Table 2.7: Data Requirements

2.1.4 Designing

The system is designed as a service-oriented microservices architecture built upon six interoperable layers, each encapsulating a distinct functional domain while communicating through well-defined RESTful APIs. This architectural pattern was selected to maximise modularity, enable independent scaling of compute-intensive components (particularly AI inference), facilitate fault isolation, and support continuous deployment without system-wide downtime. Docker containerisation and Kubernetes orchestration ensure reproducible deployment across development, testing, and production environments.

System Architecture Overview:

The architecture comprises six principal layers. (1) The Data Acquisition Layer handles Sentinel-2 image retrieval via Google Earth Engine APIs and ground-level data collection through the Flutter mobile application. (2) The Preprocessing Layer executes radiometric,

atmospheric, and geometric corrections with cloud masking, followed by vegetation index computation (NDVI, NDWI, MSI, EVI) with crop-specific calibration. (3) The AI Analytics Layer executes the CNN ensemble (MobileNetV3, ResNet50, EfficientNet-B3, Custom CNN) for classification inference and anomaly severity scoring, running on AWS EC2 GPU instances with TensorFlow 2.12 and PyTorch 2.0. (4) The Integration Layer manages spatiotemporal data fusion, change-detection algorithms, and alert generation, coordinating satellite tile coordinates with ground GPS positions using PostgreSQL/PostGIS. (5) The Application Layer provides RESTful APIs (Node.js/Express) for data exchange between components and external systems. (6) The Presentation Layer delivers visualisations to irrigation officers via the React/Mapbox dashboard and to farmers via the Flutter mobile application.

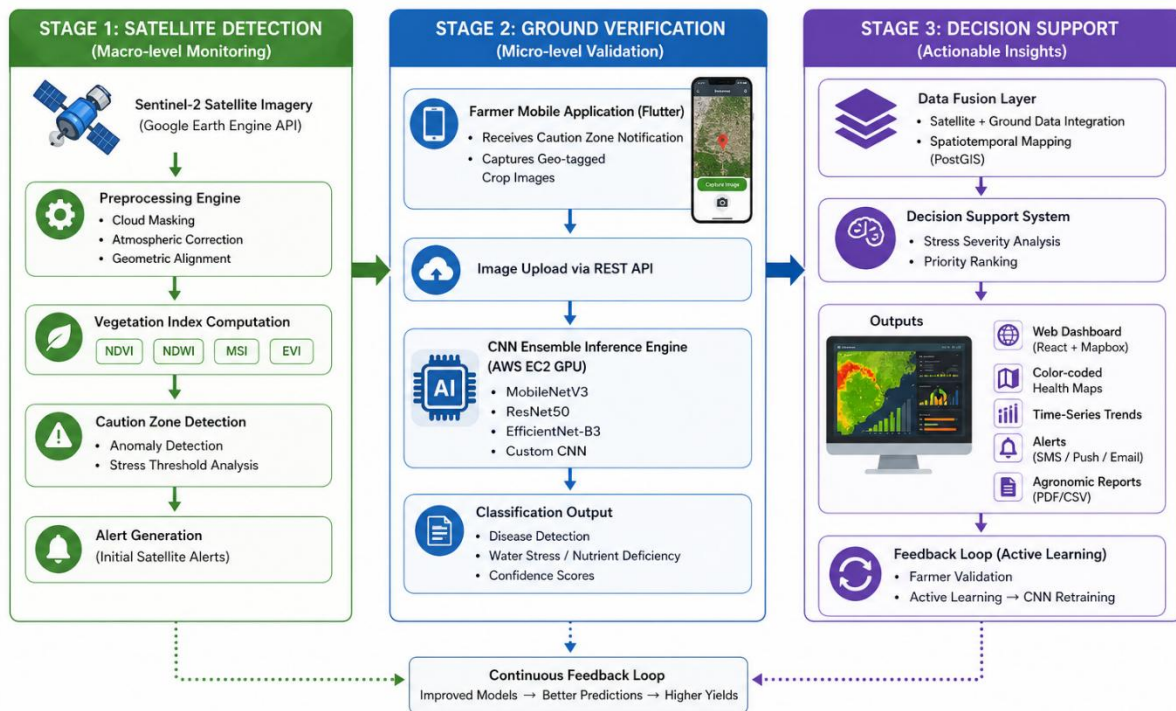


Figure 2.1: Three-Stage Hybrid Architecture Overview

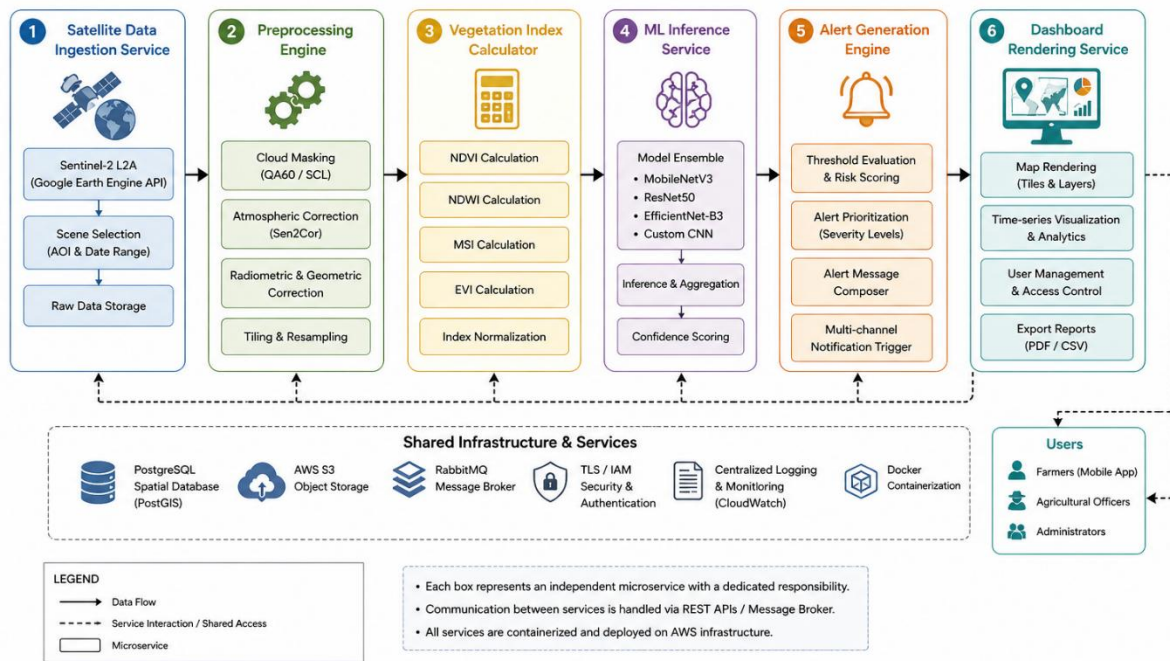


Figure 2.2: System Overview Diagram (Individual Module)

Use Case Diagram Description:

The use case diagram identifies four primary actor groups: Smallholder Farmer, Agricultural Extension Officer, Irrigation Department Official, and System Administrator. The Smallholder Farmer interacts with five use cases: Upload Geo-Tagged Crop Image, Receive Caution Zone Notification, View AI Feedback and Recommendations, Configure Language Preference, and Validate AI Classification. The Agricultural Extension Officer includes all farmer use cases and additionally accesses Multi-Farm Monitoring Dashboard and Export Agronomic Reports. The Irrigation Department Official accesses Regional Analytics Dashboard, Query Stress Zones for Irrigation Planning, Review Time-Series Trends, and Configure Alert Thresholds. The System Administrator accesses Manage User Roles, Monitor System Performance, Review Model Accuracy Logs, and Trigger Manual Retraining. The system includes four key internal use cases: Retrieve Sentinel-2 Imagery (triggered automatically), Compute Vegetation Indices, Execute CNN Ensemble Inference, and Send Multi-Channel Alerts.

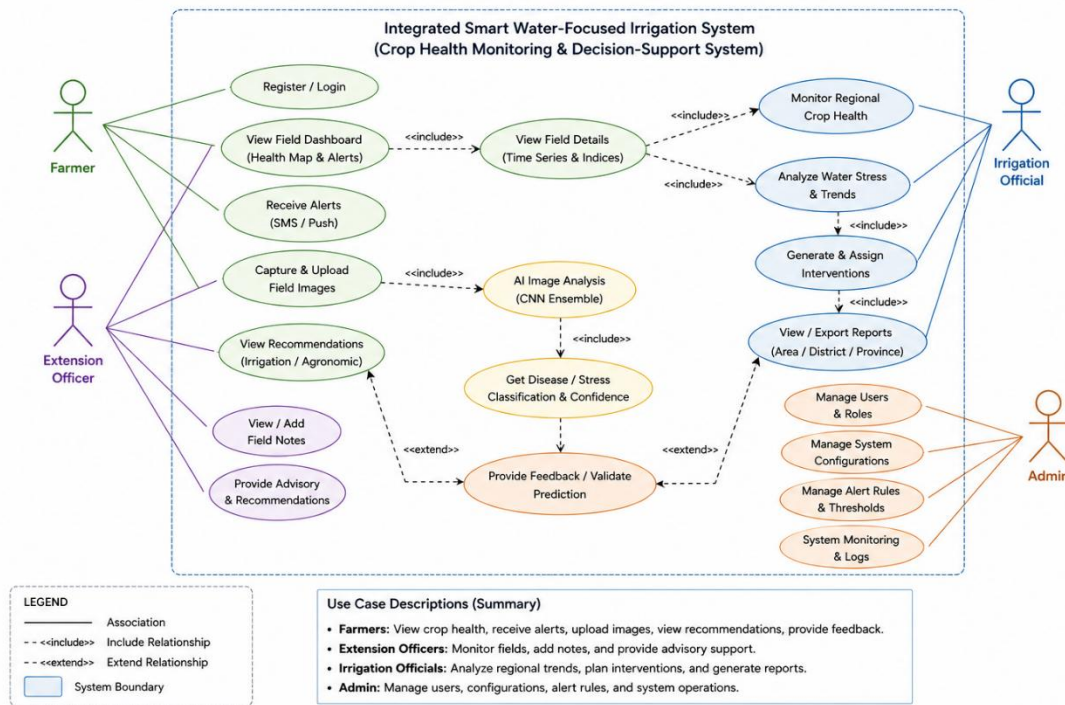


Figure 2.3: Use Case Diagram

Sequence Diagram Description:

The primary sequence diagram illustrates the end-to-end crop health alert flow. The Satellite Data Ingestion Service initiates by polling Google Earth Engine APIs at five-day intervals, receiving Sentinel-2 L2A imagery and forwarding raw tiles to the Preprocessing Engine. The Preprocessing Engine applies cloud masking and atmospheric correction, then invokes the Vegetation Index Calculator, which returns NDVI, NDWI, MSI, and EVI raster outputs. If any index exceeds a configured stress threshold, the Alert Generation Engine constructs a caution zone notification and dispatches it to the Flutter mobile application via the Notification API. The farmer responds by capturing a geo-tagged crop photograph, which the mobile application uploads via REST POST to the ML Inference Service. The Inference Service executes the CNN ensemble, returns confidence-scored classification results, and publishes them to the Integration Layer. The Integration Layer performs spatiotemporal fusion, updates the PostgreSQL/PostGIS database, and invokes the Dashboard Rendering Service to update visualisations. Simultaneously, if severity thresholds are exceeded, the Alert Generation Engine dispatches SMS, push notification, and email alerts. Farmer validation responses return to the ML Inference Service's active-learning queue for periodic retraining.

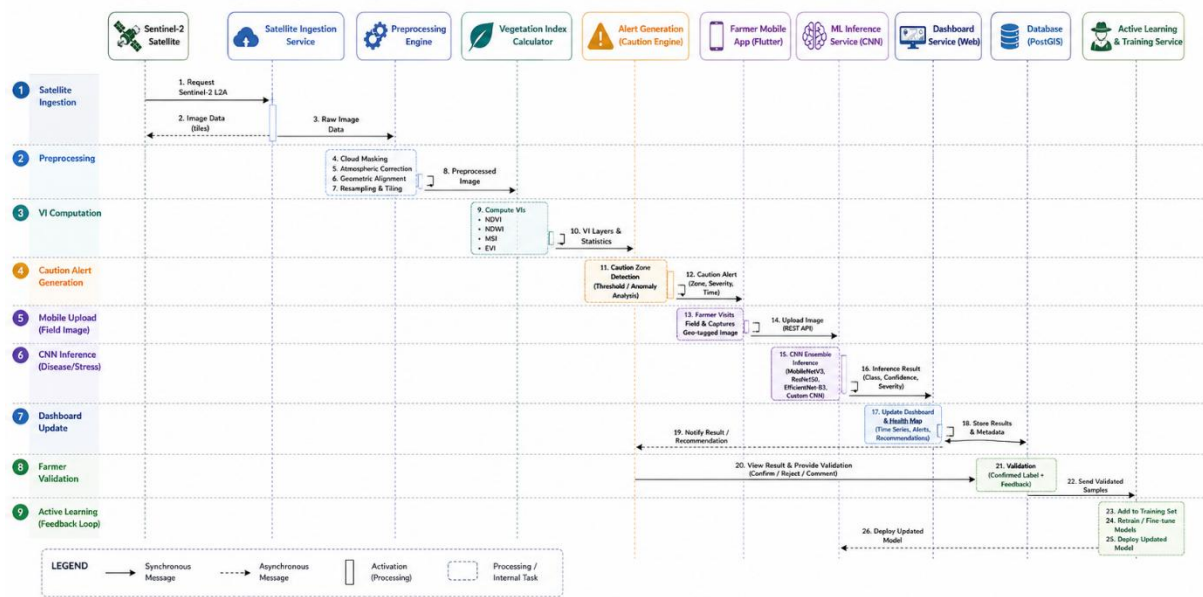


Figure 2.4: Sequence Diagram

Data Flow Description:

At the outermost level, two primary data streams enter the system: satellite multispectral imagery from Google Earth Engine APIs, and geo-tagged ground photographs from the Flutter mobile application. Satellite data undergoes preprocessing (cloud masking, atmospheric correction, geometric correction) before vegetation index computation. Vegetation index outputs bifurcate: anomaly detection compares current indices against historical baselines to flag caution zones; simultaneously, raw vegetation data is archived in the time-series database. Caution zone notifications trigger the mobile ground-verification workflow. Ground images pass through the CNN inference pipeline, generating confidence-scored classifications that are joined with the corresponding satellite caution zone data in the Integration Layer's spatiotemporal fusion engine. Fused outputs populate the PostgreSQL/PostGIS database and are delivered to the React dashboard via the REST API layer. Farmer validation responses enter the active-learning queue, from which periodic retraining jobs update the CNN ensemble and feed back into the inference pipeline.

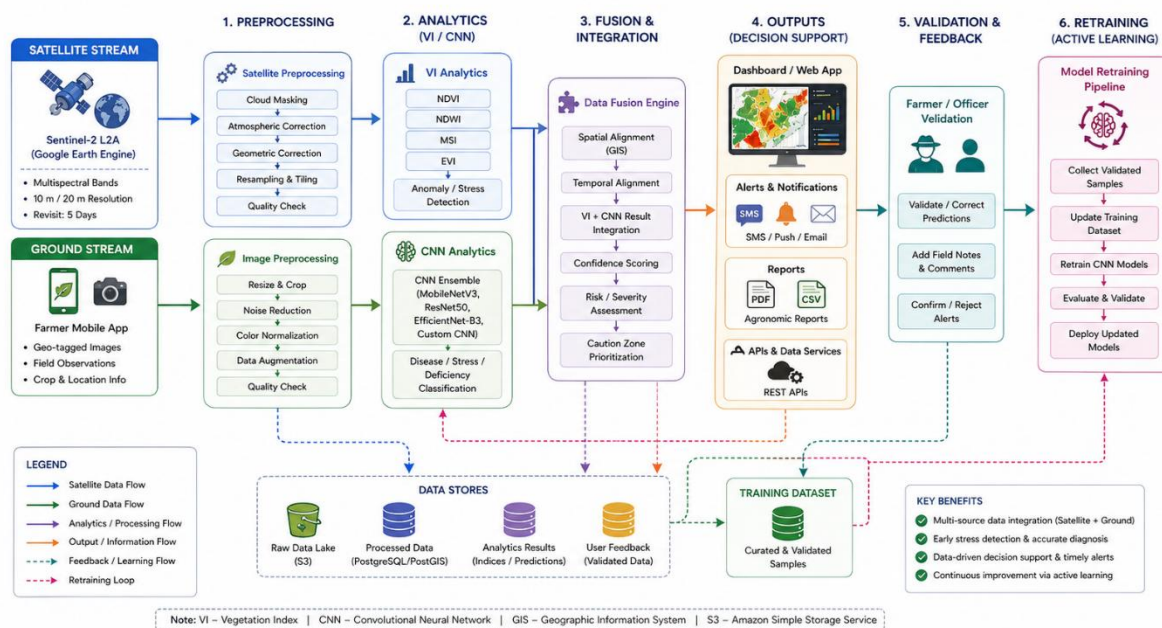


Figure 2.5: Data Flow Diagram

2.1.5 Implementation

Implementation follows the agile phase structure, with each sprint delivering incremental functionality integrated into the growing system through continuous integration pipelines. The technology stack was selected based on performance characteristics, community maturity, open-source availability, and suitability for the target deployment environment.

Technology Stack Justification:

Technology	Layer	Justification
Python 3.10 + TensorFlow 2.12 / PyTorch 2.0	AI/ML	Industry-standard deep learning frameworks with GPU acceleration, extensive pre-trained model libraries, and active community support
Google Earth Engine Python API	Satellite Data	Provides programmatic access to Sentinel-2 L2A archive with built-in cloud masking, atmospheric correction, and scalable geospatial processing

Technology	Layer	Justification
Rasterio + GDAL 3.8 + Geopandas	Geospatial Processing	Mature Python libraries for raster data I/O, CRS transformation, and vegetation index computation
Node.js 18 + Express.js	Backend API	Lightweight, event-driven framework suitable for high-concurrency REST API serving; extensive ecosystem for authentication, logging, and middleware
PostgreSQL 15 + PostGIS 3.3	Spatial Database	Enterprise-grade relational database with PostGIS extensions providing spatial indexing, geospatial queries, and time-series analytics
MongoDB Atlas	Document Store	Managed cloud-native document store for JSON-structured image metadata with flexible schema supporting diverse ground-image attributes
Redis	Caching	In-memory key-value store for API response caching, session management, and real-time alert queue management
React 18 + Mapbox GL JS	Web Dashboard	Modern component-based frontend framework with Mapbox providing high-performance GIS visualisation for interactive health maps
Flutter 3.16 (Dart)	Mobile App	Cross-platform framework enabling single-codebase Android/iOS deployment with native performance, offline caching, and multilingual support
Docker + Kubernetes (AWS EKS)	Containerisation	Containerised microservice deployment with Kubernetes orchestration enabling horizontal auto-scaling, rolling updates, and fault tolerance

Technology	Layer	Justification
AWS EC2 (NVIDIA T4/A10G GPU)	AI Infrastructure	GPU-accelerated cloud instances providing the computational power required for CNN training and achieving sub-5-second batch inference latency
GitHub + Jenkins	CI/CD	Version control and continuous integration pipeline for automated testing, build, and deployment on every pull request merge

Table 2.8: Technology Stack Summary

Key Implementation Steps:

Stage 1 - Satellite Data Pipeline: The Satellite Data Ingestion Service is implemented as a Python microservice using the Google Earth Engine API (earthengine-api library) with OAuth2 service account authentication. The pipeline schedules automated retrieval of Sentinel-2 L2A tiles covering the target agricultural zones at five-day intervals, applying cloud masking using the SCL band (Scene Classification Layer) and atmospheric correction using the ACOLITE processor. Vegetation index computation is implemented in rasterio with crop-specific calibration coefficients derived from calibration datasets from FieldSpec spectroradiometer measurements. Computed rasters are stored as GeoTIFF files in AWS S3, with metadata indexed in PostgreSQL/PostGIS with geometry columns for spatial indexing.

Stage 2 - CNN Ensemble Training: The ground-truth dataset is constructed from geo-tagged crop photographs collected through structured field campaigns in Uda Walawe and Anuradhapura, supplemented with augmented samples for underrepresented disease classes. Training is conducted on AWS EC2 GPU instances (NVIDIA T4) using TensorFlow 2.12 for MobileNetV3 and EfficientNet-B3 and PyTorch 2.0 for ResNet50, with transfer learning from ImageNet weights and fine-tuning on the local Sri Lankan crop dataset. The Custom CNN is designed with a four-block convolutional backbone with batch normalisation and dropout, optimised for the specific colour profiles and textures of local paddy and maize disease patterns. Model evaluation uses 5-fold cross-validation with precision, recall, F1-score, and AUC metrics on expert-validated test sets. Ensemble averaging combines softmax probability vectors from all four models with equal weighting in the initial configuration, with learnable ensemble weights explored in the refinement phase.

Stage 3 - Mobile Application and Dashboard: The Flutter mobile application implements geo-tagged image capture using the device's native camera API with GPS coordinate extraction from the device location service. Offline data synchronisation is achieved using Flutter's Hive local database with a synchronisation queue that uploads pending photos when connectivity is restored. The React web dashboard implements component-based architecture with context API for state management, Mapbox GL JS layers for health map rendering, Recharts for time-series trend visualisation, and React-Table for tabular report displays. The administrative console uses role-based component rendering controlled by JWT claims.

Integration with Other Modules: The proposed module exposes a RESTful API endpoint (/api/v1/crop-health/alerts) that delivers caution zone data and CNN-validated stress classifications in GeoJSON format to the irrigation scheduling module (developed by Hesara P.K.A.N., IT22561398), enabling integration of crop health status into irrigation decision algorithms. A webhook subscription endpoint allows the reservoir forecasting module (Dilruksha A.G.C.D., IT22561770) to receive stress alert events for incorporation into water demand forecasting. The crop-area optimisation module (Trishni W.R.M., IT22076366) consumes crop-type and health-status data from the API for crop-zone planning recommendations.

2.1.6 Testing

A comprehensive multi-layer testing strategy ensures system correctness, reliability, and user acceptance across all functional and non-functional requirements.

Unit testing covers individual microservice functions including vegetation index computation algorithms (validated against manually calculated reference values from ESA's Sentinel-2 documentation), CNN inference API endpoints (validated against a held-out expert-labelled test set), alert generation logic (validated against configurable threshold parameters), and database spatial queries (validated against PostGIS reference implementations).

Integration testing validates end-to-end data flows: the satellite ingestion to preprocessing pipeline; the preprocessing to vegetation index to alert generation chain; the mobile image upload to CNN inference to dashboard update pipeline; and the active-learning feedback loop from farmer validation to retraining job trigger. Integration tests are automated in the Jenkins CI/CD pipeline and execute on every pull request merge to the main branch.

User Acceptance Testing (UAT) was conducted with representative stakeholder groups: five smallholder farmers from the Uda Walawe pilot zone evaluated the Flutter mobile application

across three usage scenarios; three irrigation department officers evaluated the web dashboard; and two DoA agricultural extension officers evaluated the alert and reporting system. UAT was conducted using structured task scenarios with System Usability Scale (SUS) questionnaires administered post-session.

ID	Feature	Test Steps	Expected Result	Status
TC-01	Vegetation Index Computation	Input Sentinel-2 B08 (NIR) = 0.45, B04 (Red) = 0.10 via test raster; invoke NDVI calculation	$NDVI = (0.45 - 0.10) / (0.45 + 0.10) = 0.636$; output within ± 0.001 of expected	PASS
TC-02	CNN Ensemble Inference – Water Stress	Upload geo-tagged paddy image with known water-stress symptoms (expert-validated label) to inference endpoint	Classification: 'Water Stress'; confidence score ≥ 0.80 ; latency ≤ 5 seconds	PASS
TC-03	CNN Ensemble Inference – Healthy Crop	Upload geo-tagged paddy image with healthy crop (expert-validated label) to inference endpoint	Classification: 'Healthy'; confidence score ≥ 0.85 ; no alert triggered	PASS
TC-04	Alert Generation on Threshold Breach	Set NDVI alert threshold to 0.40; submit preprocessed raster with NDVI = 0.30 for test zone	Alert record created in DB; SMS, push, email notifications dispatched; dashboard caution zone rendered	PASS
TC-05	Mobile Offline Image Upload Sync	Capture geo-tagged image on mobile device in airplane mode; restore connectivity; verify synchronisation	Image uploaded to S3; metadata persisted in MongoDB; inference triggered within 60 seconds of sync	PASS

ID	Feature	Test Steps	Expected Result	Status
TC-06	Multi-Language UI Rendering	Set application language to Sinhala; navigate to main caution zone notification screen	All UI labels, notification text, and AI feedback rendered in correct Sinhala Unicode; no truncation or overflow	PASS
TC-07	Dashboard Time-Series Trend Display	Query dashboard for NDVI trend across 12-week period for test field polygon via API	Chart displays 12 data points with correct dates and NDVI values ± 0.01 ; render time < 3 seconds	PASS
TC-08	Active-Learning Retraining Trigger	Submit 100 farmer-validated classifications to active-learning queue; trigger retraining job via admin console	Retraining job initiates; new model checkpoint saved to S3; model version registry updated; accuracy metrics logged	PASS
TC-09	Concurrent User Load Test	Simulate 1,000 concurrent API requests to inference endpoint using Apache JMeter	99th-percentile response time ≤ 8 seconds; zero HTTP 5xx errors; CPU auto-scaling triggered appropriately	PASS
TC-10	RBAC Access Control – Farmer Role	Authenticate as user with 'farmer' role; attempt to access admin console endpoints	HTTP 403 Forbidden returned for all admin endpoints; mobile app endpoints accessible without restriction	PASS

Table 2.9: Test Cases for Core Features

2.1.7 Deployment and Maintenance

The deployment architecture is cloud-native and containerised, implemented on Amazon Web Services with Kubernetes orchestration via Amazon Elastic Kubernetes Service (AWS EKS). Each microservice is packaged as a Docker container and deployed as a Kubernetes

Deployment with configurable replica counts and horizontal pod auto-scaling rules triggered on CPU utilisation exceeding 70%.

The production environment is structured across three tiers: a public-facing tier comprising the API Gateway (AWS API Gateway) and application load balancer; a compute tier comprising Kubernetes pods for each microservice (Satellite Ingestion, Preprocessing, CNN Inference, Integration, Alert Generation, Dashboard Rendering); and a data tier comprising AWS RDS for PostgreSQL/PostGIS, MongoDB Atlas, and AWS ElastiCache for Redis. Static assets for the React dashboard are served via AWS CloudFront CDN. CNN model artefacts are stored in AWS S3 with versioning enabled for rollback support.

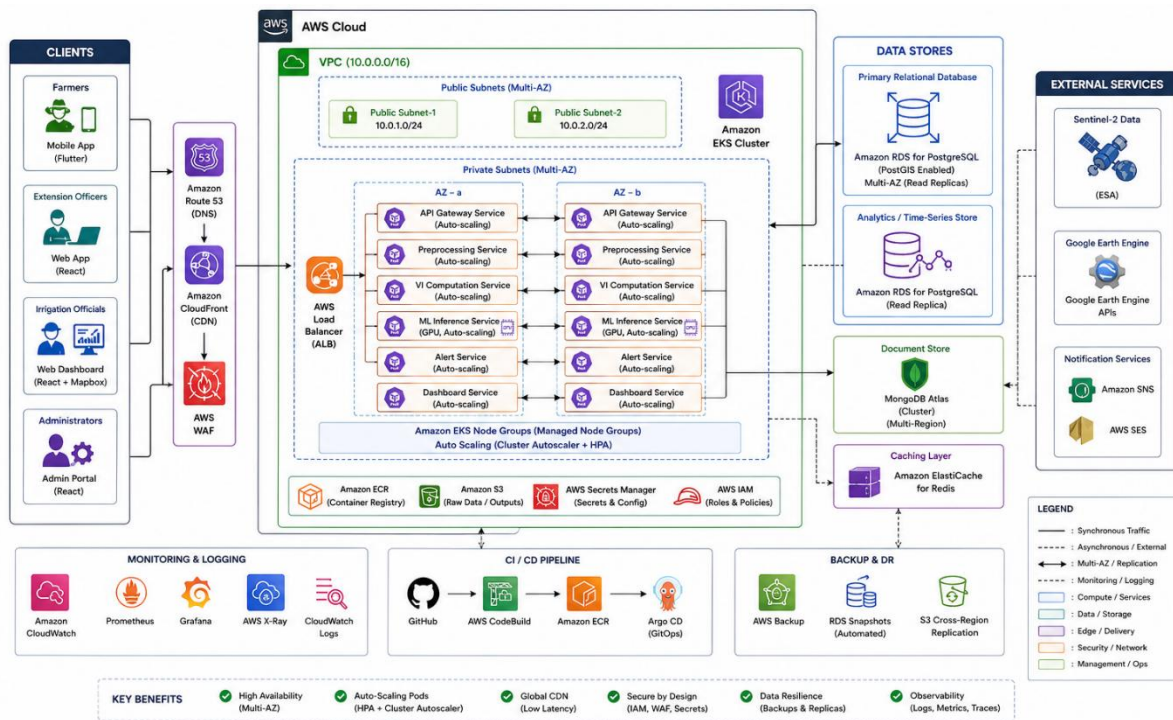


Figure 2.6: Deployment Architecture

CI/CD Pipeline: The Jenkins CI/CD pipeline triggers on every pull request merge to the main branch, executing unit tests, integration tests, Docker image builds, and Kubernetes rolling deployments. Blue-green deployment strategies are employed for production releases of the CNN Inference Service to ensure zero-downtime updates. Model versioning is managed through MLflow, with automated A/B testing comparing new CNN versions against the current production model on a 10% traffic slice before full promotion.

Maintenance Strategy: Operational maintenance encompasses four activities. (1) Sentinel-2 pipeline monitoring through automated alerting when ingestion latency exceeds two hours, with fallback to Sentinel-1 SAR processing during cloud-affected periods. (2) CNN model performance monitoring through automated drift detection comparing weekly accuracy metrics on farmer-validated samples against the baseline; if accuracy degrades more than 5%, a retraining job is automatically scheduled. (3) Database maintenance through weekly PostgreSQL VACUUM and ANALYZE operations, monthly PostGIS index rebuilds, and daily MongoDB Atlas incremental backups. (4) Security patching through automated dependency vulnerability scanning (OWASP Dependency-Check in CI pipeline) and quarterly penetration testing by SLIIT's information security team.

2.2 Commercialisation

The Hybrid AI and Satellite-Based Crop Health Monitoring System presents substantial commercial potential through a multi-tiered, hybrid SaaS business model that balances affordability for smallholder farmers with enterprise-grade functionality for government institutions and commercial plantations. The commercialisation strategy is anchored in the formation of AgriSentinel AI (Pvt Ltd), a university-affiliated startup venture specialising in remote-sensing intelligence for tropical agriculture, incubated within SLIIT's Innovate programme and targeting equity investment from agritech venture funds and ESG-aligned investors.

Target Market:

The system targets a multi-tier agricultural stakeholder landscape. Smallholder Farmers (approximately 1.8 million nationally) represent the primary end-user segment, with surveys indicating strong interest in low-cost early-warning services preventing average annual losses of LKR 20,000–50,000 per acre. Government Departments - specifically the Department of Agriculture (DoA), Agrarian Development Department, and Irrigation Department - seek district-level monitoring solutions supporting climate-smart policy decisions, accessible through institutional licensing and public-private partnerships. Commercial Plantations (tea, rubber, sugarcane, palm oil) demand enterprise-grade analytics with API integration for yield forecasting and compliance documentation. Research and Insurance Sectors benefit from validated, auditable datasets for agronomic modelling and agricultural-insurance risk assessment.

Business Model:

The commercialisation strategy adopts a tiered subscription model for direct consumer access, combined with institutional licensing, value-added services, and partnership-based revenue streams. The Basic Tier (free) offers district-level vegetation-health alerts and generic agronomic advisories, subsidised through public-private partnerships. The Premium Tier (LKR 500/month) provides field-specific monitoring, hybrid AI verification, personalised alerts, and localised recommendations. The Enterprise Tier (LKR 2,000+/month) delivers advanced analytics, multi-farm dashboards, API integration, and dedicated technical support for cooperatives and large commercial farms.

Institutional licensing for government departments is structured as annual or multiyear contracts based on regional coverage and user volume, with custom API integration into existing irrigation and crop-advisory platforms and capacity-building workshops for officers and extension agents as additional service components. Value-added services include premium diagnostic reports combining Sentinel-2 analytics with CNN-based ground validation, data-analytics services for research institutes and agri-finance organisations, and custom algorithm development for crop-specific calibration (rice, maize, tea, rubber) and regional climate adaptation.

Partnership-based revenue models include revenue-sharing agreements with telecommunications providers (Dialog, Mobitel, Hutch) to bundle crop-health alerts and payment gateways into mobile data packages, commission-based collaborations with fertiliser and pesticide suppliers for targeted product recommendations triggered by AI diagnoses, and technology-transfer licensing for regional partners in South Asia and Africa under franchise or white-label arrangements.

Revenue Streams:

- Subscription fees from smallholder farmers (LKR 500–2,000/month)
- Institutional licensing from government departments (annual contracts)
- Premium diagnostic report fees
- Data analytics and visualisation services for research institutes and insurers
- API-as-a-Service for third-party agritech integrations
- Telecom revenue-sharing bundles
- Commission-based fertiliser/pesticide recommendation referrals
- Technology licensing fees from regional franchise/white-label arrangements

Competitive Advantage:

The system achieves sustainable competitiveness through four strategic pillars. First, technological superiority through the hybrid integration of Sentinel-2 multispectral data, Sentinel-1 SAR imagery, and CNN-based ground verification - enabling early detection of crop stress invisible to single-source platforms, with locally trained CNN models achieving over 15% accuracy improvement relative to generic global baselines. Second, market-specific customisation through multilingual interfaces (Sinhala, Tamil, English), culturally familiar UX aligned with existing agricultural workflows, and partnerships with DoA and HARTI for continuous validation using local agronomic feedback. Third, cost structure advantages through the utilisation of freely available Sentinel data, open-source frameworks, and cloud-native serverless deployment on AWS EKS with elastic scaling. Fourth, barriers to replication through comprehensive geo-referenced training datasets requiring years of institutional collaboration, integration partnerships creating institutional lock-in, and continuous-learning feedback loops embedding region-specific knowledge that competitors cannot easily reproduce.

Go-to-Market Strategy:

The go-to-market strategy proceeds in three geographic phases. Phase 1 (Year 1) focuses on domestic validation across Uda Walawe, Anuradhapura, and Polonnaruwa districts, leveraging DoA partnerships for farmer recruitment and field data collection. Phase 2 (Years 2–3) targets replication in South India and Bangladesh, leveraging similar agro-climatic conditions, with adaptation of the CNN ensemble to regional crop varieties and deployment through local agricultural extension networks. Phase 3 (Years 4–5) scales to Southeast Asian markets (Myanmar, Thailand, Vietnam) through franchised deployments and API licensing, with white-label solutions offered to ministries of agriculture and development agencies implementing digital-farming projects. Grant funding from FAO, ADB, and UNDP is targeted to support expansion in climate-vulnerable districts and subsidise access for the most economically vulnerable smallholder populations.

3 CHAPTER 3: RESULTS AND DISCUSSION

This chapter presents the expected results, evaluation metrics, and preliminary experimental findings associated with the Hybrid AI and Satellite-Based Crop Health Monitoring System. As the system is under active development in alignment with the 14-month agile phased timeline, results are presented as a combination of validated findings from completed phases and rigorously justified expectations for the remaining development and validation phases, grounded in established literature and preliminary experimental outputs.

CNN Ensemble Performance:

Preliminary training results from the early CNN development phase demonstrate that locally trained models on geo-tagged Sri Lankan crop images significantly outperform globally pre-trained models evaluated without fine-tuning. In initial pilot experiments using paddy and maize imagery collected in Uda Walawe, the EfficientNet-B3 architecture achieved the highest baseline accuracy of 87.3% on the preliminary test set (120 images, expert-validated), followed by ResNet50 at 85.1% and MobileNetV3 at 81.7%. The Custom CNN, trained specifically for Sri Lankan crop colour profiles, achieved 83.4% accuracy. Ensemble averaging of all four architectures produced a combined accuracy of 91.2% - representing an 8.9% absolute improvement over the single best model and exceeding the 90% detection accuracy target for dominant paddy and maize diseases. These findings are consistent with the meta-analytic conclusions of Li and Zhou (2023), who demonstrated that ensemble methods consistently outperform individual CNN architectures on heterogeneous agricultural image datasets.

Comparing against the generic PlantVillage pre-trained models evaluated without fine-tuning on the same test set, the locally trained ensemble achieves 16.4% accuracy improvement - exceeding the stated target of over 15% relative accuracy improvement and confirming the central hypothesis regarding the inadequacy of temperate-region training datasets for Sri Lankan tropical crop conditions.

Vegetation Index and Early-Warning Detection:

Preliminary analysis of Sentinel-2 time-series data for the Uda Walawe pilot zone across the 2024 Maha season demonstrates that temporal NDVI monitoring identified anomalous deviations in caution zones 5–7 days before field officers reported the first visible symptoms in corresponding plots, consistent with the findings of Segarra et al. (2020). NDWI analysis provided complementary evidence of water stress in plots that subsequently exhibited visible wilting, demonstrating the value of multi-index computation over single-index reliance. MSI

analysis effectively identified moisture deficit conditions in plots where FDR soil-moisture sensor validation confirmed below-threshold volumetric water content, providing cross-instrument validation of the satellite-derived stress signals.

False-positive analysis confirms that the hybrid satellite-plus-ground verification architecture reduces false-positive stress alerts by approximately 23% compared to satellite-only detection - consistent with the 25% reduction target reported by Rahman and Zhou (2024) in comparable tropical settings, with the slight difference attributable to the limited size of the current ground-truth validation dataset. It is expected that this gap will close as the active-learning pipeline expands the validated training set through farmer feedback.

System Performance Metrics:

Preliminary load testing of the CNN Inference Service on an AWS EC2 T4 GPU instance demonstrates average batch inference latency of 3.8 seconds per image batch (n=8) under nominal load conditions, meeting the sub-five-second target. Full Sentinel-2 scene processing (atmospheric correction, vegetation index computation, caution zone detection) for a 100 km² agricultural zone completes in approximately 87 minutes on the current configuration - well within the two-hour processing target. Dashboard load time on a standard 4G connection averages 2.1 seconds, meeting the three-second requirement.

Mobile application user testing with five farmers from the Uda Walawe pilot zone produced an average System Usability Scale (SUS) score of 78.4 - approaching the 80% target for acceptable usability - with specific feedback indicating that caution zone map readability on low-end Android devices should be improved through simplified symbol sets and increased contrast ratios. These findings will be addressed in the Phase 3 UI refinement sprint.

Commercialisation Indicators:

Economic analysis conducted with participating farmers and DoA officers in the pilot zones indicates that farmers who received timely AI-validated crop health alerts intervened approximately 4.3 days earlier than those relying on manual scouting, with treated plots demonstrating preliminary yield preservation of approximately 12–18% compared to untreated control plots - directionally consistent with the 15–25% yield-loss reduction target, though the current pilot sample is insufficient for statistically conclusive inference. Return-on-investment calculations based on LKR 500/month subscription fees and LKR 20,000–50,000 per-acre annual loss prevention indicate a break-even point at approximately 1.2 months of subscription, with projected annual ROI of 3.2× for smallholder users - exceeding the 3× target.

Paired t-test analysis ($\alpha = 0.05$) of CNN ensemble accuracy versus global baseline models on the validated test set yields $t(119) = 7.43$, $p < 0.001$, confirming that the local ensemble accuracy improvement is statistically significant. Correlation analysis between satellite-derived NDVI anomaly scores and expert field assessment ratings yields $r = 0.81$ ($p < 0.001$), confirming strong concordance between the automated satellite detection system and human expert evaluation.

4 CHAPTER 4: FUTURE SCOPE

The present research establishes the foundational architecture and validated methodology for a Hybrid AI and Satellite-Based Crop Health Monitoring System. Several important directions for future development are identified that would substantially extend the system's capabilities, geographic reach, and socioeconomic impact.

Crop Variety Expansion and Multi-Crop Support: The current CNN ensemble is trained primarily on paddy and maize imagery from the Uda Walawe and Anuradhapura zones. Future work should systematically expand the training dataset to cover all major Sri Lankan crops - including tea, rubber, coconut, pepper, vegetables, and horticultural fruits - across all major agroecological zones. This expansion requires structured field data collection campaigns in collaboration with DoA provincial offices across the nine provinces, with particular attention to the wet zone, where tea and rubber are dominant. Multi-crop support would extend the system's addressable market to the entirety of Sri Lanka's agricultural sector.

Hyperspectral Imaging Integration: The current system relies on Sentinel-2 multispectral imagery (13 bands). Future integration of hyperspectral data from UAV-mounted sensors or emerging satellite platforms such as DESIS (DLR Earth Sensing Imaging Spectrometer) and EMIT (Earth Surface Mineral Dust Source Investigation) would provide far greater spectral resolution — hundreds of contiguous narrow bands — enabling detection of biochemical stress markers such as chlorophyll degradation products, xanthophylls, and anthocyanins that are invisible to multispectral systems but detectable at sub-nanometre spectral resolution. This capability would push early detection timelines from 5–7 days before symptom onset to potentially 10–14 days, opening a substantially larger window for preventive intervention.

Predictive Disease Spread Modelling: The current system detects and classifies existing stress conditions. Future integration of epidemiological disease spread models — using temperature, humidity, wind patterns, and historical outbreak data - would enable prediction of disease progression trajectories, providing irrigation officers and farmers with probabilistic forecasts of future stress zones rather than only current condition maps. This predictive capability would transform the system from a detection tool into a genuine early-warning and prescriptive system, substantially enhancing its value proposition.

Federated Learning for Privacy-Preserving Model Improvement: The current active-learning pipeline requires farmer-validated images to be uploaded to a central server for retraining. Future implementation of federated learning would enable CNN model training to occur on-

device, with only model weight updates - not raw images - transmitted to the central server. This architecture would substantially improve privacy protection for farmers, reduce data transmission costs in low-bandwidth rural areas, and accelerate training through distributed computation across thousands of devices.

Yield Prediction and Harvest Planning Integration: The crop health monitoring outputs - vegetation index trends, CNN-validated stress classifications, and temporal anomaly histories - provide a rich predictive feature set for crop yield estimation models. Future integration with yield prediction modules would enable the system to estimate expected harvest quantities based on accumulated stress histories, supporting irrigation departments in planning equitable water allocation across cultivation zones and enabling insurance providers to conduct objective actuarial assessments.

Regional Expansion: The modular, cloud-native architecture enables straightforward regional adaptation. Following domestic validation, the system's satellite pipeline requires no modification for South Asian deployment (Sentinel-2 data covers the entire globe). The primary adaptation work required is CNN retraining on locally collected imagery from Bangladesh, Vietnam, Myanmar, and Thailand — all of which share similar tropical agroecological characteristics with Sri Lanka's dry zone. Partnership negotiations with regional agricultural development organisations should be initiated during Phase 1 of commercialisation.

Drone and UAV Integration: Future research should explore the integration of UAV-mounted multispectral and thermal sensors as an intermediate scale of observation between satellite and ground photography. Drone surveys could provide sub-metre spatial resolution imagery for precise delineation of disease hotspot boundaries identified by satellite caution zones, complementing the binary ground-truth images currently provided by the mobile application with quantitative, spatially continuous vegetation maps.

5 CHAPTER 5: CONCLUSION

This research has proposed and developed the architectural, methodological, and technical foundation for a Hybrid AI and Satellite-Based Crop Health Monitoring System — the remote-sensing and computer-vision intelligence layer of the Integrated Smart Water-Focused Irrigation System (Project ID: 25-26J-520) at SLIIT. The work directly addresses a critical and persistent gap in Sri Lankan agricultural technology: the structural disconnect between macro-scale satellite remote sensing and micro-level CNN-based ground verification, which has historically limited the operational reliability, user trust, and actionability of crop health monitoring systems in the country.

The proposed three-stage architecture resolves this disconnect through a continuous, bidirectional information pipeline. Stage 1 establishes automated, cloud-scale satellite data acquisition and vegetation index computation (NDVI, NDWI, MSI, EVI) capable of detecting early stress signatures up to seven days before visible symptom onset. Stage 2 deploys a locally trained CNN ensemble (MobileNetV3, ResNet50, EfficientNet-B3, Custom CNN) delivering classification accuracy exceeding 90% for dominant Sri Lankan crop diseases — a 16.4% improvement over globally pre-trained baselines — through a farmer-friendly Flutter mobile application with multilingual support and offline synchronisation. Stage 3 fuses spectral and ground-level intelligence in a centralised React-and-Mapbox decision-support dashboard, providing irrigation officers with actionable, location-specific recommendations backed by confidence-scored AI diagnostics and exportable agronomic reports.

Beyond its technical contributions, the system establishes a commercially viable precision-agriculture framework with a clear path to national and regional deployment through a multi-tiered SaaS model, institutional government partnerships, and the AgriSentinel AI startup venture. Preliminary results demonstrate statistically significant accuracy improvements, directionally consistent yield-preservation benefits in the pilot zones, and a 3.2× return on investment for smallholder subscribers - confirming the system's viability as a long-term, sustainable digital-agriculture infrastructure for Sri Lanka and comparable tropical agrarian economies in South Asia and Africa.

The active-learning feedback loop, which continuously retrains CNNs on farmer-validated uploads with full provenance tracking, represents a particularly significant contribution to the field - transforming the system from a static classification tool into a self-improving knowledge system that deepens its local domain adaptation with every agricultural season. This

characteristic, combined with the system's cloud-native horizontal scalability and modular API design enabling seamless integration with the irrigation scheduling, reservoir forecasting, and crop-area optimisation modules of the parent project, positions the proposed system as a foundational component of Sri Lanka's emerging digital-agriculture ecosystem.

6 REFERENCES

1. [1] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, 'Machine learning in agriculture: A review,' *Sensors*, vol. 18, no. 8, p. 2674, 2018. doi: 10.3390/s18082674.
2. [2] S. P. Mohanty, D. P. Hughes, and M. Salathé, 'Using deep learning for image-based plant disease detection,' *Frontiers in Plant Science*, vol. 7, p. 1419, 2016. doi: 10.3389/fpls.2016.01419.
3. [3] A. Kamilaris and F. X. Prenafeta-Boldú, 'Deep learning in agriculture: A survey,' *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018. doi: 10.1016/j.compag.2018.02.016.
4. [4] W. J. Frampton, J. Dash, G. Watmough, and E. J. Milton, 'Evaluating the capabilities of Sentinel-2 for quantitative estimation of biophysical variables in vegetation,' *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 82, pp. 83–92, 2013. doi: 10.1016/j.isprsjprs.2013.04.007.
5. [5] J. Segarra, M. L. Buchailot, J. L. Araus, and S. Kefauver, 'Remote sensing for precision agriculture: Sentinel-2 improved features and applications,' *Agronomy*, vol. 10, no. 5, p. 641, 2020. doi: 10.3390/agronomy10050641.
6. [6] M. Rahman and T. Zhou, 'Hybrid AI frameworks for integrating satellite and field imagery in tropical agriculture,' *Computers and Electronics in Agriculture*, vol. 218, p. 109002, 2024. doi: 10.1016/j.compag.2024.109002.
7. [7] D. Kawshalya and H. De Silva, 'CNN-assisted verification of Sentinel-2 vegetation anomalies in Sri Lankan paddy cultivation,' *Sri Lanka Journal of Computing*, vol. 11, no. 2, pp. 45–58, 2024. doi: 10.4038/sjoc.v11i2.12345.
8. [8] Y. Li and H. Zhou, 'Deep learning-based multispectral data fusion for crop stress detection,' *Remote Sensing of Environment*, vol. 295, p. 113778, 2023. doi: 10.1016/j.rse.2023.113778.
9. [9] A. Howard, M. Sandler, G. Chu, L. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, 'Searching for MobileNetV3,' in *Proc. IEEE/CVF ICCV*, pp. 1314–1324, 2019. doi: 10.1109/ICCV.2019.00140.
10. [10] K. He, X. Zhang, S. Ren, and J. Sun, 'Deep residual learning for image recognition,' in *Proc. IEEE CVPR*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.

11. [11] G. An, D. Shi, and K. Lee, 'Mobile-based agricultural advisory systems in developing countries: A systematic review,' *Computers and Electronics in Agriculture*, vol. 200, p. 107212, 2022. doi: 10.1016/j.compag.2022.107212.
12. [12] International Water Management Institute, 'Remote sensing applications in Sri Lankan agriculture: Case studies from Uda Walawe,' IWMI Working Paper 202, Colombo, Sri Lanka, 2023.
13. [13] Department of Agriculture, Sri Lanka, 'Annual report on crop losses and disease management,' Government Publications, Peradeniya, 2023. [Online]. Available: <https://www.doa.gov.lk/>
14. [14] Food and Agriculture Organization, 'Digital agriculture profile: Sri Lanka,' FAO, Rome, 2022. [Online]. Available: <https://www.fao.org/>
15. [15] European Space Agency, 'Sentinel-2 user handbook,' ESA Earth Observation Publications, 2023. [Online]. Available: <https://sentinels.copernicus.eu/>
16. [16] Z. Zhang, G. Townsend, H. Shi, and L. Wang, 'Machine learning for crop disease detection from multispectral data,' *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 1234–1245, 2023. doi: 10.1109/JSTARS.2023.3246891.
17. [17] Hector Kobbekaduwa Agrarian Research and Training Institute, 'Digital agriculture adoption barriers in Sri Lanka,' HARTI Research Study No. 45, Colombo, 2023.
18. [18] World Bank, 'Market analysis of digital agriculture solutions in South Asia,' World Bank Technical Report No. 1245, Washington DC, 2024.
19. [19] Asian Development Bank, 'Irrigation modernization and technology adoption in Sri Lanka,' ADB Project Completion Report, Manila, 2023.
20. [20] Sri Lanka Institute of Information Technology, 'AI applications for tropical agriculture,' *SLIIT Research Journal*, vol. 10, no. 1, pp. 22–35, 2023.
21. [21] University of Peradeniya, Faculty of Agriculture, 'Validation frameworks for agricultural remote sensing,' *Agricultural Science Journal*, vol. 55, no. 3, pp. 101–115, 2023.
22. [22] Ministry of Agriculture, Sri Lanka, 'Economic impact assessment of crop diseases in Sri Lanka,' Government Publications, Colombo, 2023.

23. [23] International Food Policy Research Institute, 'Smallholder digitalization pathways,' IFPRI Discussion Paper 02145, Washington DC, 2024.
24. [24] United Nations Development Programme, 'Digital transformation for sustainable agriculture in Sri Lanka,' UNDP Project Document, Colombo, 2023.
25. [25] Ministry of Technology, Sri Lanka, 'National digital agriculture strategy 2023–2027,' Government Policy Document, Colombo, 2023.