الجمهورية الجزائرية الديمقراطية الشعبية
PEOPLE'S DEMOCRATIC REPUBLIC OF ALGERIA

وزارة التعليم العالي والبحث العلمي
MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH

جامعة عمار ثليجي بالأغواط
UNIVERSITY OF LAGHOUAT

FACULTY OF SCIENCES
DEPARTMENT OF COMPUTER SCIENCE

Master thesis

Domain : Mathematics and Computer Science
Sector : Computer Sciences
Option : Data Science and Artificial intelligence

Submitted by

Farouk Brachemi & Youcef Boudia

Real-Time Crop Health Monitoring in Precision Agriculture Using Geospatial Data and Deep Learning

Defended Publicly in Front of the Jury Composed of:

Mr. Lakhdar OULAD DJEDID	President
Mr. Laradj CHELLAMA	Examiner
Mr. Lakhdar KACHNA	Examiner
Mr. Mohamed El Habib MAICHA	Supervisor

Academic Year 2024 / 2025

بِسْمِ اللَّهِ الرَّحْمَٰنِ الرَّحِيمِ

Acknowledgments

First, all thanks and praise be to Allah, who gave us the strength and patience to complete this work.
I would like to thank our supervisor, Prof. Mohamed El Habib Maicha, for his help, guidance, and valuable advice during this project. His support made a big difference.
Many thanks to our families for their love, support, and encouragement throughout this journey.
I also thank all our professors who taught us and shared their knowledge during these years.
To my classmates and friends, it was a great honor to study with you and learn together.
Finally, thank you to everyone who helped or supported us in any way, whether directly or indirectly.

Dedications

This thesis is dedicated to:
- Our beloved families, whose constant support, prayers, and encouragement made this journey possible.
- Our friends, for their motivation and companionship throughout this work.
- All our teachers and mentors who shared their knowledge with us.
- The professors and students of Laghouat University, for their guidance and collaboration.
- My friend Slimane Hocini, whose influence and input greatly impacted this project.
- Finally, to everyone—near or far—who contributed in any way to the success of this thesis.

Abstract

Precision agriculture increasingly relies on advanced technologies to monitor crop health and optimize farming operations. This thesis presents a real-time crop monitoring system that leverages UAV-acquired multispectral imagery, deep learning, and a scalable data streaming infrastructure. High-resolution drone images are processed to generate vegetation indices such as NDVI, which serve as input for U-Net models. These models perform semantic segmentation and classification to detect various crop stress types, including nutrient deficiency, drydown, planter skips, water stress, and weed clusters.

To ensure real-time processing and system scalability, the architecture integrates Apache Kafka for streaming UAV imagery and prediction outputs, supporting parallel operation across multiple producers and consumers. A React-based dashboard displays live predictions and metadata, enabling field operators to make timely and informed decisions.

The system was validated using the Agriculture-Vision 2021 dataset and tested under simulated UAV deployment. The best segmentation performance was achieved using a U-Net model with RGB+NIR input at 512×512 resolution, reaching a Dice score of 0.70 and classification accuracy of 95%. In contrast, a NDVI-based model offered faster inference (3–5s) with slightly lower accuracy, making it suitable for resource-constrained environments. Kafka demonstrated low-latency image transmission (<2.5s) even with up to 15 parallel producers.

Keywords: Precision Agriculture, NDVI, UAV, U-Net, Real-time Inference, Kafka, Vegetation Indices, Machine Learning, Deep Learning.

الملخص

تعتمد الزراعة الدقيقة بشكل متزايد على التقنيات المتقدمة لمراقبة صحة المحاصيل وتحسين العمليات الزراعية. تقدم هذه الأطروحة نظامًا لمراقبة المحاصيل في الوقت الحقيقي، يستفيد من الصور متعددة الأطياف التي يتم التقاطها بواسطة الطائرات بدون طيار، إلى جانب تقنيات التعلم العميق وبنية تحتية قابلة للتوسّع لبث البيانات.

تُعالَج صور الطائرات عالية الدقة لاستخراج مؤشرات الغطاء النباتي مثل مؤشر NDVI، والتي تُستخدم كمدخلات لنماذج U-Net التي تقوم بعمليات التجزئة الدلالية والتصنيف لاكتشاف أنواع متعددة من إجهاد المحاصيل، بما في ذلك نقص المغذيات، الجفاف، الفجوات الناتجة عن تخطي الغراس، الإجهاد المائي، وتكتلات الأعشاب الضارة.

ولضمان المعالجة في الوقت الفعلي وقابلية النظام للتوسّع، تم دمج Kafka Apache لبث صور الطائرات ونتائج التنبؤات، مما يدعم التشغيل المتوازي عبر عدد من المنتجين والمستهلكين. كما تعرض لوحة تحكم مبنية باستخدام React التنبؤات والبيانات الوصفية بشكل مباشر، مما يمكّن العاملين في الحقل من اتخاذ قرارات سريعة ومدروسة.

تم التحقق من فعالية النظام باستخدام مجموعة بيانات Agriculture-Vision 2021، واختباره في بيئة محاكاة لنشر طائرات بدون طيار. وقد حقق أفضل أداء في التجزئة باستخدام نموذج U-Net مع مدخلات مكونة من RGB+NIR وبدقة 512×512، حيث وصلت درجة Dice إلى 70.0 ودقة التصنيف إلى 95%. في المقابل، قدم النموذج المعتمد على مؤشر NDVI استدلالًا أسرع (بين 3 إلى 5 ثوانٍ) مع دقة أقل قليلاً، مما يجعله مناسباً للبيئات ذات الموارد المحدودة. وأظهرت Kafka قدرة على نقل الصور بزمن استجابة منخفض (أقل من 5.2 ثانية) حتى مع وجود ما يصل إلى 15 منتجًا متوازيًا.

الكلمات المفتاحية: الزراعة الدقيقة، NDVI، الطائرات بدون طيار، U-Net، الاستدلال في الوقت الحقيقي، Kafka، مؤشرات الغطاء النباتي، التعلم الآلي، التعلم العميق.

الكلمات المفتاحية: الزراعة الدقيقة، NDVI، الطائرات بدون طيار، U-Net، الاستدلال في الوقت الحقيقي، Kafka، مؤشرات الغطاء النباتي، التعلم الآلي، التعلم العميق.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1 GENERAL INTRODUCTION

1.1 Context and Motivation

Modern agriculture faces increasing pressure to produce more food while contending with climate variability, resource constraints, and the need for sustainability. Precision agriculture has emerged as a key approach, using technology to optimize crop monitoring, resource allocation, and decision-making.

Unmanned Aerial Vehicles (UAVs), equipped with multispectral sensors, offer high-resolution imagery capable of capturing key vegetation indicators such as NDVI (Normalized Difference Vegetation Index). When combined with machine learning models, this imagery can reveal early signs of crop stress—such as nutrient deficiency, drydown, or weed infestations—before they become visually apparent. However, much of the current analysis is conducted offline, limiting its effectiveness for real-time decision support.

This thesis addresses this gap by proposing a real-time crop health monitoring pipeline that integrates UAV-based data collection, deep learning segmentation models, and scalable data streaming via Apache Kafka.

1.2 Problem Statement

Current crop monitoring methods often rely on delayed or manual analysis, which can be inefficient and reactive. While UAVs provide valuable high-resolution imagery, turning this raw data into timely insights remains a challenge due to:

- The computational demands of deep learning models like U-Net when applied to large datasets.
- The lack of real-time infrastructure to support fast data ingestion, inference, and delivery.
- Fragmented systems that do not integrate image processing, classification, and visualization in a unified pipeline.

There is a clear need for a lightweight, scalable, and real-time solution that leverages vegetation indices like NDVI and modern ML techniques to enable proactive crop stress detection.

1.3 Research Objectives

This research aims to develop a real-time UAV-based crop monitoring system using deep learning and streaming infrastructure. The specific objectives include:

1. Designing a dual-output lightweight U-Net model optimized for NDVI-based anomaly detection and classification.
2. Preprocessing UAV imagery to simulate real-world conditions using sequences of high-resolution field images.
3. Implementing a Kafka-based pipeline for real-time data ingestion, model inference, and visualization.
4. Evaluating the system's performance under varying conditions, including latency and prediction accuracy across multiple UAV streams.

1.4 Document Organization

This thesis is structured to reflect the full development cycle of the proposed system:

- Chapter 2 provides a literature review of machine learning, geospatial analysis, and real-time streaming systems in agriculture.
- Chapter 3 identifies strengths and limitations of previous works and presents key research gaps this thesis addresses.
- Chapter 4 outlines the system methodology, including model design, dataset preparation, and pipeline implementation.
- Chapter 5 presents experimental results, including inference latency, prediction accuracy, system scalability, and discusses findings, limitations, and implications for real-time precision agriculture..
- Chapter 6 concludes the thesis and suggests future directions for extending this work.

CHAPTER 2

BACKGROUND AND THEORETICAL FRAMEWORK

This chapter provides the theoretical foundation for our study by reviewing key concepts and technologies in precision agriculture. We begin with an overview of crop health monitoring techniques, geospatial data usage, and remote sensing tools. We then explore how machine learning is applied to agricultural tasks such as disease detection and yield prediction. The chapter also introduces Apache Kafka as a real-time streaming platform and examines how these technologies are being integrated. Finally, we highlight the current gaps in fully combining machine learning, geospatial data, and streaming systems—pointing to the need for unified, intelligent solutions in modern agriculture.

2.1 Precision Agriculture and Crop Health Monitoring

2.1.1 Overview of Precision Agriculture

Farming has always been about working with the land. But today, it's also about working with data. Precision agriculture (PA) flips traditional farming on its head by treating fields not as uniform plots, but as mosaics of micro-environments, each with its own needs.

The goal is to use resources like water, fertilizers, and pesticides more efficiently, so crops grow better and farming is more sustainable. Tools like GPS, sensors, and UAVs help farmers understand the condition of their fields in detail. [2].

At its core, PA relies on three pillars:

1. Data Collection: GPS-guided drones, soil sensors, and satellites gather real-time info on moisture, nutrients, and plant health.

2. Analysis: Machine learning models crunch this data to spot patterns (e.g., "This corner of the field always dries out first").

3. Action: Farmers use these insights to apply water, fertilizer, or pesticides only where needed cutting costs and boosting yields.

2.1.2 Crop Health Monitoring Techniques

Monitoring crop health is a critical component of precision agriculture, involving the use of advanced technologies to assess plant conditions and detect stress, disease, or nutrient deficiencies at early stages. Techniques such as computer vision analyze images to track growth and identify abnormalities like leaf discoloration or stunted development. Thermal imaging helps detect water stress by measuring temperature variations, while multispectral imaging captures reflectance across different wavelengths, including near-infrared, which reveal subtle signs of nutrient shortages or disease before they become visible to the naked eye. [3] [4]

Drones equipped with multispectral cameras have become indispensable tools in this process. They can rapidly survey large fields, capturing high-resolution, geo-tagged images that provide detailed spatial data on crop vigor, water stress, and pest infestations. These drones combine a high-resolution RGB camera with multiple multispectral sensors (Green, Red, Red Edge, Near-Infrared), delivering precise crop health insights with centimeter-level accuracy. These insights enable farmers to implement targeted interventions such as variable-rate fertilizer application or early pest control, which improve yields, reduce input costs, and minimize environmental impact. By integrating crop health monitoring with drone-based multispectral imaging, precision agriculture empowers farmers to make data-driven decisions that enhance productivity and sustainability. [5] [6]

2.1.3 Geospatial Data in Agriculture

Geospatial data plays a crucial role in modern farming by providing detailed information about crops and fields. This data comes from satellites, drones, soil sensors, and GPS-enabled equipment, helping farmers see how conditions change across their land. When combined with Geographic Information Systems (GIS), it allows for better visualization and more informed decisions.

Multispectral data refers to information captured across multiple, discrete bands of the electromagnetic spectrum, extending beyond the visible range to include ultraviolet (UV), near-infrared (NIR), and sometimes shortwave infrared (SWIR) wavelengths. Unlike standard color images, which record only red, green, and blue (RGB) channels, multispectral imaging systems can acquire data in several specific spectral bands, each chosen for its relevance to the application or material under investigation.

Satellites capture multispectral images that show how plants reflect different types of light, including visible and near-infrared (NIR). Drones complement this by taking high-resolution photos closer to the crops. GPS technology tags all data points with exact locations, so farmers know precisely where to focus their attention. GIS then layers this information to create maps that highlight specific zones needing care [7, 8].

A key part of analyzing this data involves vegetation indices. These are simple formulas that use light reflectance to estimate plant health. The most common index is the Normalized Difference Vegetation Index (NDVI), calculated as

$$\text{NDVI} = \frac{\text{NIR} - \text{Red}}{\text{NIR} + \text{Red}}$$

NDVI values range from -1 to 1, where higher values generally mean healthier, denser vegetation. However, NDVI can sometimes lose sensitivity in very dense crops. To address this, the Enhanced Vegetation Index (EVI) is used:

$$\text{EVI} = 2.5 \times \frac{\text{NIR} - \text{Red}}{\text{NIR} + 6 \times \text{Red} - 7.5 \times \text{Blue} + 1}$$

EVI is better at handling dense vegetation and atmospheric effects, giving a clearer picture in those situations [9, 8].

Farmers use these indices to detect stress early, manage irrigation and fertilization more precisely, predict yields, and monitor pests. GIS platforms help integrate other data like soil type and terrain to optimize how fields are managed.

While working with geospatial data can be complex due to different formats and large datasets, cloud computing and machine learning are making these tools easier to use and more accessible. Beyond individual farms, geospatial data also supports regional and global food security by identifying areas with yield gaps and helping prioritize resources [10].

2.1.4 Remote Sensing

Remote sensing in precision agriculture involves the use of various sensors and platforms to gather detailed information about soil and crop conditions, enabling site-specific management to optimize agricultural inputs and improve crop yield and health.

Remote Sensing Applications

- Monitoring crop health by detecting crop vigor, damage, and stress through optical sensors that capture visible and infrared wavelengths.
- Observing soil conditions such as organic matter, texture, pH, and moisture content to understand their impact on crop performance.
- Mapping soil properties, classifying crop species, detecting crop water stress, monitoring weeds and diseases, and estimating crop yield.
- Supporting variable-rate technology (VRT) for precise application of fertilizers, pesticides, and irrigation based on spatial variability within fields.
- Detecting abiotic and biotic stresses, including pests and diseases, and estimating crop growth parameters like leaf area index, chlorophyll content, biomass, and canopy temperature.
- Enabling real-time monitoring and decision-making for improved farm management. [11]

Types of sensors used

- Satellite Sensors: Provide wide-area coverage with varying spatial, spectral, and temporal resolutions. Examples include the Sentinel-2 A+B constellation, which offers improved temporal, spatial, and spectral resolution for agricultural monitoring.

- Unmanned Aerial Vehicles (UAVs): Offer high spatial resolution and flexible sensor payloads for detailed field-scale monitoring.

- Aircraft-based Sensors: Used for intermediate-scale observations with customizable sensor suites.

- Ground-based Sensors: Include handheld or tractor-mounted multispectral and hyperspectral sensors for very high-resolution data collection.

- Multispectral Sensors: Capture data in several broad spectral bands (e.g., visible, near-infrared).

- Hyperspectral Sensors: Capture data in many narrow spectral bands, allowing detailed analysis of crop biochemical and biophysical traits.

- Thermal Sensors: Measure canopy or soil temperature to detect water stress and other physiological conditions.

- LiDAR and Fluorescence Spectroscopy: Advanced sensors for structural and physiological crop assessments. [12]

2.2 Machine Learning for Crop Health Assessment

Machine learning (ML) is a branch of Artificial Intelligence (AI) that enables computers to learn from data without being explicitly programmed for specific tasks [1]. ML can be categorized into four primary learning paradigms, each with distinct characteristics and applications in agriculture, as shown in Figure 2.1.

2.2.1 Supervised Learning

Supervised learning involves training models on labeled datasets, where each input is paired with the correct output. This paradigm is particularly useful for tasks like disease detection in crops, where the model learns to identify diseased plants based on annotated examples. Techniques such as Support Vector Machines (SVMs) and Random Forests are commonly employed in this context [13].

2.2.2 Unsupervised Learning

Unsupervised learning deals with unlabeled data, allowing models to identify patterns or groupings without predefined labels. In agriculture, this approach can assist in segmenting fields based on soil properties or detecting anomalies in crop growth, facilitating early intervention strategies [13].

2.2.3 Reinforcement Learning

Reinforcement learning involves training models to make sequences of decisions by rewarding desired behaviors. In farming, this paradigm can optimize irrigation schedules or pesticide application by learning the most effective strategies over time through trial and error [13].

2.2.4 Semi-Supervised and Self-Supervised Learning

Semi-supervised learning combines a small amount of labeled data with a large amount of unlabeled data, making it beneficial when labeled data is scarce but unlabeled data is abundant. Self-supervised learning, on the other hand, generates labels from the data itself, enabling models to learn useful representations without manual l abeling. These approaches are particularly advantageous in agricultural settings where obtaining labeled data can be challenging [14].

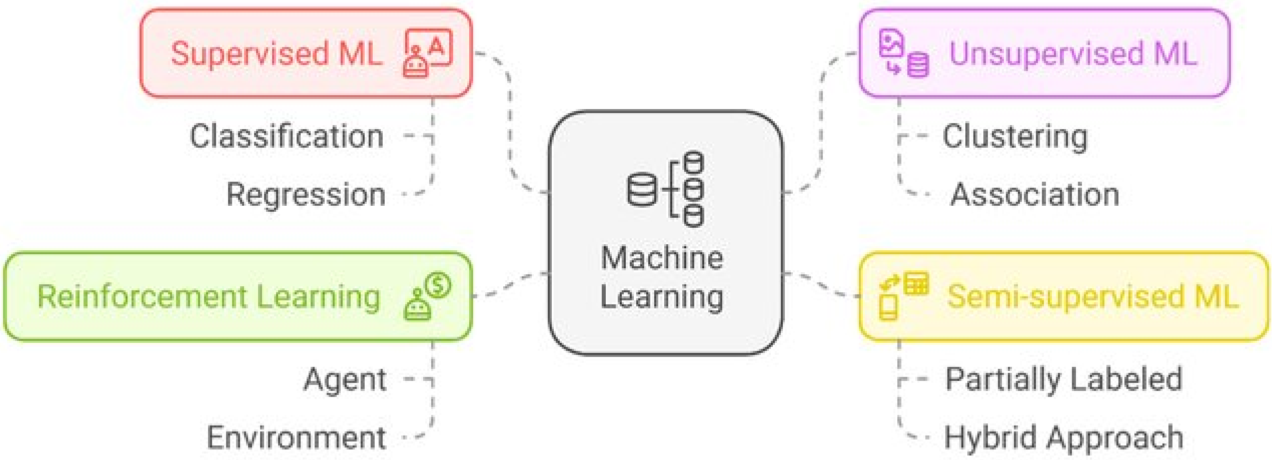

Figure 2.1: Learning Paradigms [1]

2.2.5 Applications in Crop Health Assessment

Machine learning (ML) techniques have been increasingly applied in agriculture to enhance crop health assessment. These applications leverage various data sources and algorithms to monitor, predict, and manage crop conditions effectively.

Disease Detection

ML models, particularly deep learning algorithms, have been utilized to detect and classify plant diseases from leaf images. Convolutional Neural Networks (CNNs) can identify patterns associated with specific diseases, enabling early intervention and reducing crop losses [15].

Stress Monitoring

By analyzing data from sensors and remote sensing technologies, ML algorithms can monitor plant stress factors such as drought, nutrient deficiencies, and pest infestations. These models help in timely decision-making to mitigate stress impacts on crops [16].

Yield Prediction

ML approaches, including Random Forests and Support Vector Machines, have been employed to predict crop yields based on historical data, weather patterns, soil conditions, and farming practices. Accurate yield predictions assist in resource planning and market strategies [17].

Precision Agriculture

Integrating ML with precision agriculture enables site-specific crop management. ML models analyze spatial and temporal data to optimize inputs like water, fertilizers, and pesticides, enhancing productivity and sustainability [18].

2.3 Real-Time Data Processing with Apache Kafka

Apache Kafka is a distributed streaming platform designed for high-throughput, fault-tolerant real-time data processing. Its architecture and core concepts enable scalable ingestion, storage, and analysis of streaming data, making it indispensable for precision agriculture applications requiring immediate insights from diverse data sources.

2.3.1 Kafka Architecture and Key Concepts

Kafka architecture is built around several core components that work together to ensure reliability, scalability, and fault tolerance, as can be seen in Figure 2.2.

- Broker: A Kafka broker is a server that receives data from producers, stores it on disk, and serves it to consumers. Brokers form a cluster where one broker acts as the cluster controller, managing partition assignments and monitoring broker health. Each broker can handle thousands of partitions and millions of messages per second depending on hardware.

- Topic: Topics are logical channels that organize messages. Each topic represents a stream of events, such as user notifications or transaction logs. Topics are append-only and immutable; messages once written cannot be changed or deleted by consumers.

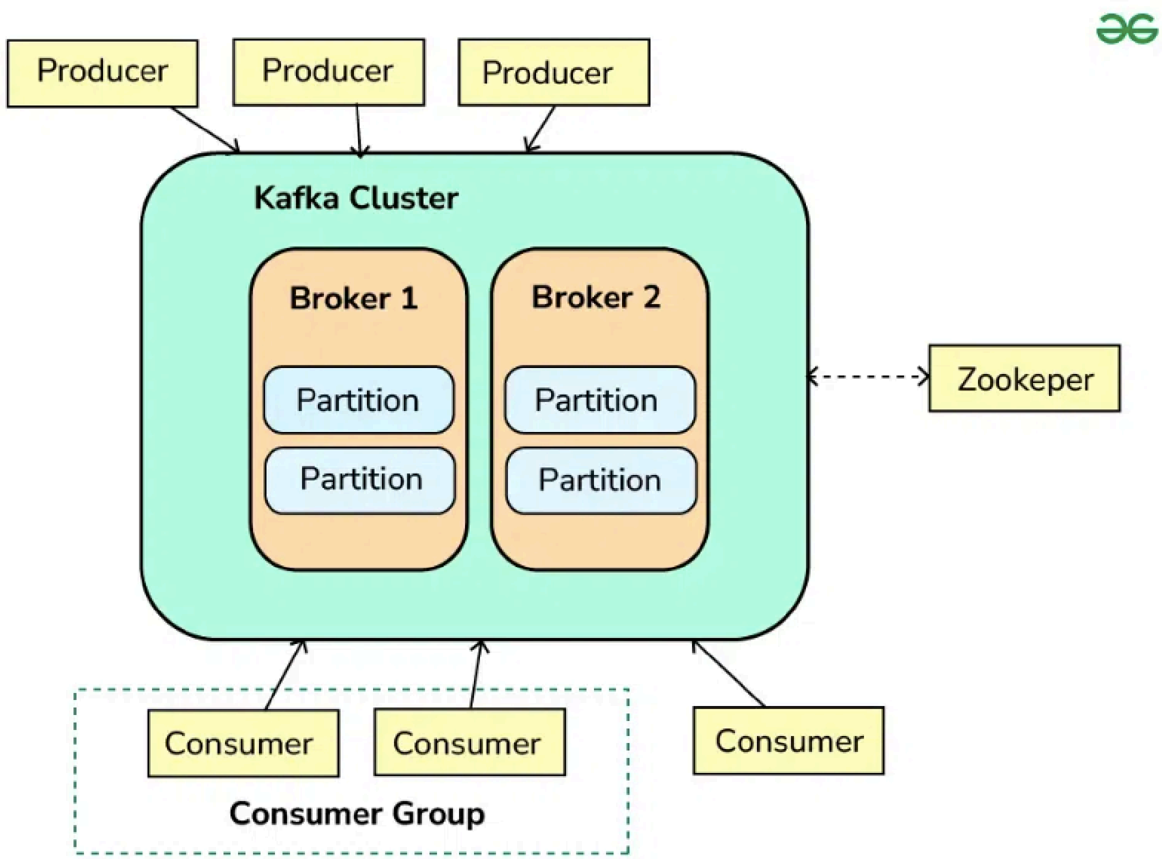

Figure 2.2: Kafka Architecture Overview [1]

- Partition: Topics are split into partitions, which are ordered, append-only logs. Each partition is hosted on a single broker and assigned a unique offset to every message, ensuring strict ordering within that partition. Partitions enable horizontal scalability and redundancy by distributing data across multiple brokers.
- Producer: Producers are client applications that publish messages to Kafka topics. They serialize data into key-value pairs and decide which partition to send each message to, often using the message key to ensure ordering guarantees within partitions.
- Consumer: Consumers subscribe to topics to read messages. They track their position using offsets, allowing them to resume reading from where they left off. Consumers can operate individually or as part of consumer groups to share the workload of processing partitions.
- Consumer Group: A consumer group is a collection of consumers that coordinate to consume messages from partitions of a topic. Each partition is assigned to only one consumer within the group, enabling parallel processing and fault tolerance through automatic partition reassignment if a consumer fails. [1]

[1] https://www.redpanda.com/guides/kafka-architecture

- Zookeeper: A centralized system that supports distributed applications by handling configuration, naming, synchronization, and group services. Kafka originally relied on ZooKeeper to manage its cluster by tracking active brokers, partition leaders, topic configurations (like partitions and replicas), consumer group membership, and access permissions for reading and writing to topics. [2]

In Kafka, data flows through several key stages to ensure reliable and scalable event streaming. First, producers serialize and transmit messages to the leader broker responsible for the target partition within a topic. The broker then appends these messages to the partition log on disk and replicates them to follower brokers for durability and fault tolerance. Consumers retrieve messages by polling brokers for data from their assigned partitions, processing each message in the order determined by its offset. To maintain reliable processing and support message replay, consumers also commit their offsets, keeping track of which messages have already been consumed. This mechanism ensures both fault-tolerant data delivery and the ability to resume processing from the correct point in the event of failures. [3]

2.3.2 Kafka Use Cases in IoT and Agriculture

Apache Kafka plays a vital role in precision agriculture by serving as a scalable, fault-tolerant platform for real-time data streaming and processing. It efficiently handles high-throughput, heterogeneous data from IoT sensors, drones, and satellites, enabling continuous ingestion and distribution of key variables such as soil moisture, temperature, and environmental conditions. This real-time capability supports AI-driven crop monitoring and decision-making, making Kafka a foundational component in sustainable precision agriculture systems.

A practical example of this approach involves deploying Kafka to stream data from multiple IoT sensors monitoring vineyard conditions. This real-time data pipeline feeds AI analytics platforms that automatically trigger precise irrigation adjustments based on soil moisture thresholds. By leveraging Kafka's robust streaming capabilities, the system optimizes water usage, reduces waste, and maintains crop health, demonstrating Kafka's effectiveness in supporting sustainable farming practices.

This example highlights Kafka's ability to integrate diverse data sources and deliver continuous, reliable data streams to AI models, making it a backbone technology for modern precision agriculture systems aiming to enhance resource efficiency and crop productivity. [19]

2.4 Integration of Machine Learning, Geospatial Data, and Streaming Platforms

Despite the increasing importance of real-time intelligent systems in precision agriculture, the full integration of Machine Learning (ML), Geospatial Data (GD), and Streaming Platforms

[2] https://www.geeksforgeeks.org/kafka-architecture
[3] https://www.redpanda.com/guides/kafka-architecture

(SP) remains underexplored. However, there are notable advancements in the pairwise combinations of these technologies, which demonstrate the potential for a future holistic framework. Below, we discuss each of these pairwise integrations with supporting literature.

2.4.1 Machine Learning and Geospatial Data

The integration of machine learning and geospatial data has been extensively explored, particularly for tasks such as land cover classification, crop yield prediction, and disaster assessment. In the context of precision agriculture, this pairing enables models to extract patterns from spatially referenced data to inform decisions on crop health and resource allocation. Linaza et al. exemplify this approach by employing deep learning models such as U-Net on high-resolution satellite imagery. Their system supports agricultural decision-making by performing automatic segmentation and disease classification, leveraging geospatial data to provide a spatial understanding of field conditions.[19]

2.4.2 Machine Learning and Streaming Platforms

ML and streaming platforms are increasingly integrated to support real-time inference and anomaly detection from continuous data sources, such as IoT sensors or mobile devices.

In this pair, ML models operate on continuous data streams from farm sensors or imaging. Bonacci et al. demonstrate such integration: they built an on-board push-broom hyperspectral scanner whose spectral data are classified by an embedded neural network, and classification results are streamed in real time via Apache Kafka.

In practice, the system continuously acquires hyperspectral crop images, applies an MLP classifier to identify plant type, and publishes the results over Kafka. This real-time ML+streaming pipeline enables live crop monitoring and adaptive decision support (e.g. weed/crop discrimination) in precision agriculture. [20]

2.4.3 Geospatial Data and Streaming Platforms

Here the focus is on streaming spatial data feeds into analytics platforms. Theofilou et al. present an agricultural big-data warehouse where Apache Flume/Kafka ingest spatiotemporal feeds (IoT sensor logs, satellite metadata) and Spark Streaming processes them. In their architecture, weather, soil and satellite sensor data (all georeferenced) are treated as continuous streams: Flume/Kafka handle high-velocity ingestion and Spark Streaming performs real-time processing. This design explicitly supports geospatial data (e.g. partitioning by region, time) in a streaming context, enabling up-to-date spatial analytics (like live mapping of soil moisture or climate patterns) to inform farm management decisions. [21]

2.5 Conclusion

This chapter has outlined the key technological components underpinning modern precision agriculture: crop health monitoring, geospatial data, machine learning, and real-time streaming platforms. We examined how each of these elements contributes to smarter, data-driven farming practices, and reviewed existing efforts to integrate them in pairs. However, the full convergence of machine learning, geospatial analytics, and real-time data processing remains an open challenge. This gap highlights the need for integrated systems capable of delivering timely, spatially-aware, and intelligent insights—an objective that frames the motivation for the work presented in the following chapters.

CHAPTER 3

RELATED WORK

This chapter reviews recent research efforts at the intersection of machine learning, geospatial data, and real-time processing in the context of precision agriculture. We first explore how various machine learning models—ranging from classical classifiers to deep neural networks—have been applied for crop monitoring, disease detection, and yield prediction. Next, we examine the role of real-time data processing frameworks, such as Apache Kafka, in handling large-scale agricultural data streams. We then highlight the use of geospatial data, particularly satellite and UAV imagery, in crop health assessment through vegetation indices like NDVI and NDWI. Finally, we identify key research gaps that motivate the need for more generalizable models, efficient infrastructure, and integrated solutions in smart agriculture.

3.1 Machine learning in Precision Agriculture

The integration of machine learning (ML) with remote sensing has revolutionized precision agriculture by enabling advanced data analysis for crop monitoring, yield prediction, and resource management. Below is a detailed overview of the key ML models and their applications, as evidenced by recent research:

Support Vector Machines (SVM) and Random Forest (RF) are the most widely adopted models. SVM accounts for over 20% of applications, excelling in crop classification and disease detection due to its robustness in handling high-dimensional remote sensing data. RF follows closely at 18%, particularly effective for yield prediction and land-use classification by leveraging ensemble learning to reduce overfitting. [17]

For soil moisture estimation, Decision Tree and Extra Tree Regression models have demonstrated exceptional performance, achieving approximately 91% accuracy when analyzing time-series satellite imagery. These models are favored for their interpretability in identifying soil health indicators critical for olive cultivation and other crops. [18]

In UAV-based applications, RF classifiers have shown remarkable success in detecting crop lodging (e.g., rice varieties), achieving 96.1% accuracy by integrating multispectral UAV data

(444–842 nm). SVM is also frequently paired with hyperspectral remote sensing (>30% adoption) for detailed vegetation stress analysis. [22]

These models address challenges such as data heterogeneity from diverse sensors (e.g., satellites, UAVs) while balancing accuracy with computational efficiency. Future advancements aim to enhance model generalizability across regions and crop types, particularly in resource-constrained agricultural systems.

Deep convolutional networks, especially U-Net and its variants, have been widely applied to UAV imagery for vegetation segmentation and health classification. For instance, Zhang et al in 2021 they proposed Ir-UNet, an improved U-shaped CNN, to segment irregular wheat yellow rust lesions in multispectral UAV data, reporting an overall accuracy of 97.1% for identifying diseased vs. healthy regions, outperforming a standard U-Net [15]. In Segmentation / Classification architectures. For example, Silva et al in 2024 they compared Mask R-CNN, U-Net, and YOLO models on high-resolution UAV images of soybean and bean fields for weed detection, they found that modern single-shot detectors example YOLOv8 reached the highest detection and segmentation accuracy, comparing Mask R-CNN and U-Net [23]. A study by El Hoummadi et al. used UAVs with RGB and NIR imagery to map agricultural areas in Dubai, achieving 89.7% precision with NDVI as input compared to 72. 8% with RGB alone, showing that vegetation index like NDVI can affect more in the performance of model [24].

For the past five years, the Agriculture Vision Teams [1] have been hosting challenges focused on analyzing agricultural patterns. In their 2020 challenge, participants evaluated models like DeepLab V3+ and Feature Pyramid Network (FPN) architectures for semantic segmentation of crop anomalies. The FPN-based model achieved superior performance (62.3% mIoU) compared to DeepLab variants, demonstrating better capability to handle irregular agricultural patterns.

These studies highlight that machine learning, especially deep networks, is able to perform fine-grained vegetation segmentation or disease classification from multispectral UAV data, often with precision in the 80s to 90% range.

3.2 Real-Time Data Processing in Agricultural Applications

The integration of real-time data processing frameworks into agricultural systems has become increasingly important for enabling precision farming, early detection of plant stress, and efficient resource management. With the advent of UAVs, remote sensors, and high-resolution imagery, agriculture now generates vast amounts of data that require immediate analysis to support timely and informed decisions. Technologies such as Apache Kafka and Spark Structured Streaming are being leveraged to build scalable pipelines capable of ingesting and processing streaming data with minimal delay.

In one relevant study, a real-time processing pipeline was developed using a producer-consumer architecture powered by Apache Kafka and Spark Structured Streaming. The producer component reads images from a local directory, converts them into byte streams, and

[1] https://www.agriculture-vision.com/

publishes them to a Kafka topic, with each image tagged using a unique identifier derived from its filename to preserve ordering.

On the consumer side, Spark Structured Streaming subscribes to the Kafka topic and ingests image byte streams in real time. These images are normalized and passed through a pre-trained deep learning segmentation model to generate corresponding flood masks.

To ensure scalable and efficient storage, the system uses MongoDB's GridFS, which handles large image files by chunking them into smaller parts. This integration supports continuous, low-latency image processing, making it well-suited for time-sensitive applications like disaster monitoring and response [25].

3.3 Crop Health Monitoring using Geospatial Data

High-resolution satellite and UAV imagery are now routinely used to assess crop health. Multispectral satellites (e.g. Sentinel-2, Landsat, PlanetScope) deliver frequent global coverage in visible, red-edge, NIR and SWIR bands, while drones can capture fine-scale multispectral data on demand [26, 27]. These data enable calculation of vegetation indices: for example, NDVI (Normalized Difference Vegetation Index, using NIR and red bands) is a well-established proxy for green biomass and vigor [27, 28], and NDWI (using NIR and SWIR) highlights canopy water content and drought stress [27]. In practice, mapping NDVI/NDWI over fields lets farmers spot stressed areas (low NDVI/NDWI) before visual symptoms appear. For instance, Jeevan et al. (2024) note that drone multispectral cameras (visible + NIR) produce NDVI maps that correlate closely with plant health [28].

Studies show the complementary benefits of satellite and UAV data. Sentinel-2 imagery offers wide-area, regular revisit (10m – 20m pixels every 2–5 days) with bands optimized for vegetation indices [26], whereas UAV flights yield centimeter-scale detail at key growth stages. Li et al. (2022) compared Sentinel-2 vs UAV multispectral data on wheat fields: Sentinel-2 provided smooth, broad NDVI trends, while UAV-based NDVI captured fine spatial variability in the canopy [26]. In essence, drones reveal small-scale issues (e.g. patchy nutrient stress or localized disease), and satellites supply context over larger areas. By fusing these data and tracking indices like NDVI/NDWI through the season, practitioners can diagnose nutrient deficiencies, irrigation problems or emerging diseases with better precision, informing targeted interventions and resource management [28, 27]

3.4 Research Gaps and Opportunities

Despite notable progress in applying machine learning to precision agriculture, several challenges and limitations still hinder broader adoption and effectiveness. Key research gaps and opportunities include:

- Lack of Generalizability: Many studies focus on specific crop types and limited geographies, leading to models that do not generalize well across diverse agro-ecological zones and underrepresented crops.

- Computational Complexity of Models: Deep learning models like U-Net and YOLO offer high accuracy but have complex architectures that are computationally intensive and unsuitable for real-time or edge deployment.
- Need for Lightweight Models: There is a need to develop simplified or lightweight versions—such as a lightweight U-Net—that can perform efficiently in real-world farming conditions.
- Limitations of Vegetation Indices: Relying on certain vegetation indices like NDWI can be restrictive, as they may not always be computable due to missing spectral bands.
- Neglect of Data Infrastructure: While model development is well-studied, the supporting infrastructure for real-time agricultural intelligence (e.g., streaming frameworks like Apache Kafka) is largely overlooked.

3.5 Conclusion

This chapter has reviewed the evolving intersection of machine learning, geospatial analytics, and real-time data streaming within the context of precision agriculture. Traditional classifiers like SVM and Random Forest have demonstrated strong performance on structured tasks such as crop type classification and disease detection, while deep networks such as U-Net and YOLO variants have enabled fine-grained vegetation segmentation and anomaly localization in high-resolution UAV imagery.

Despite these advancements, current solutions face persistent challenges: limited model generalizability, high computational costs, and underutilized streaming infrastructures. While vegetation indices such as NDVI and NDWI remain central to crop health monitoring, their integration into scalable, real-time systems remains underdeveloped.

These identified gaps directly motivate the development of our lightweight, dual-output U-Net model and Kafka-based streaming pipeline. In the following chapter, we detail the complete system architecture, data preprocessing pipeline, model design, and real-time evaluation strategies used to address these challenges and enable efficient crop health monitoring at scale.

CHAPTER 4 METHODOLOGY

This chapter presents the design and implementation of a real-time crop monitoring system based on UAV imagery, deep learning models, and scalable streaming infrastructure. The methodology includes the system architecture, data sources, model design, and the end-to-end data pipeline leveraging Apache Kafka.[1]

4.1 System Architecture Overview

Our system is designed to process high-resolution imagery from unmanned aerial vehicles (UAVs) equipped with multispectral cameras capturing both RGB and near-infrared (NIR) data. We use Apache Kafka to stream this raw data in real time, simulating multiple UAVs acting as producers that send images to a single consumer, as illustrated in Figure 4.1.

The raw images, which can be as large as 200 MB, are processed and downsized to a more manageable 10 MB for efficiency. Historical data is stored in a MongoDB database, while the processed images are saved using a MinIO bucket-style storage system.

The system outputs (visualization) three types of prepared images:

- RGB images for visual inspection,
- NDVI (Normalized Difference Vegetation Index) images to assess and analyze crop health.
- Prediction masks from our machine learning model, overlaid on RGB images, highlighting stress types (e.g. nutrient deficiency, drydown, or water stress) if detected. Each image comes with metadata, including band details and resolution, to ensure traceability and usability.

[1]The full implementation is available at: https://github.com/ucef-b/RT-streaming

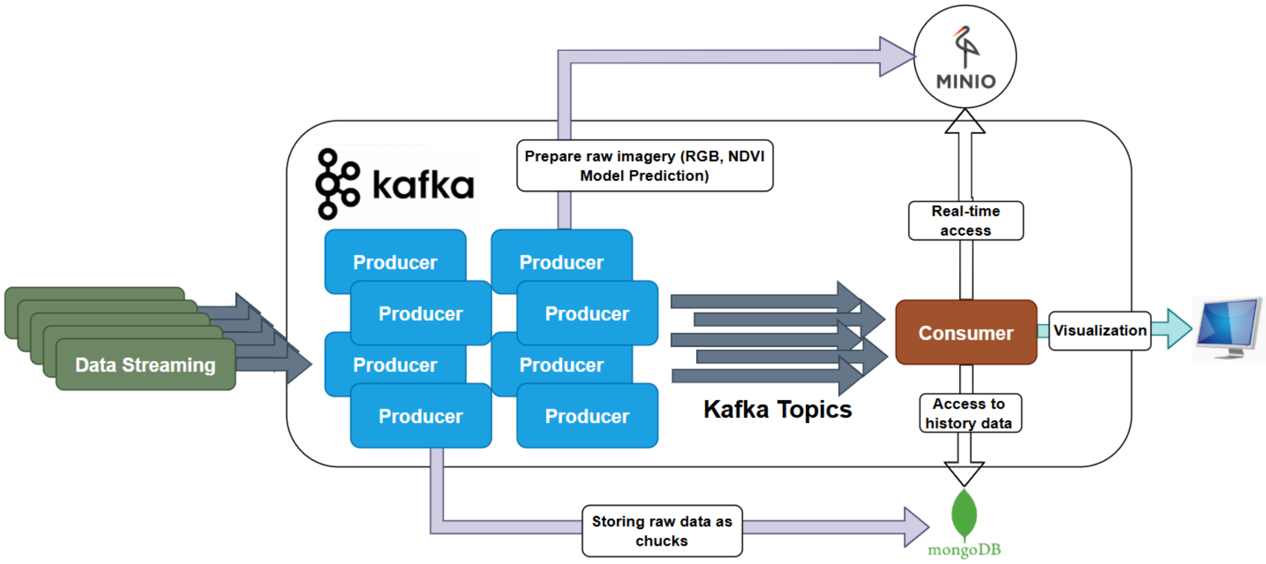

Figure 4.1: System Architecture

4.2 Data Preprocessing

For this work, we relied on the Agriculture-Vision 2021 dataset, a comprehensive collection of aerial imagery tailored for precision agriculture. This dataset is split into two main components: a supervised dataset and a raw image dataset. Below, we break down each part and explain how we preprocessed the data to suit our needs.

4.2.1 Supervised Dataset

Description

The supervised dataset contains cropped versions of the original captured image, divided into patches of 512x512 pixels, each with a filename in the format (field_id)_(x1)−(y1)−(x2)−(y2) · (jpg/png). The field_id uniquely identifies the farmland, while the coordinates (x1, y1, x2, y2) indicate the specific region of the field the image covers. Each image comes with corresponding masks labeling various agricultural conditions or anomalies.The dataset originally includes nine mask classes: double plant, endrow, planter skip, waterway, weed cluster, drydown, nutrient deficiency, storm damage, and water.

Prepossessing

Since our focus is on crop health, we narrowed down the classes to five that are most relevant: nutrient deficiency, drydown, water, planter skip and weed cluster. Figure 4.2 shows examples of two of the selected classes. The remaining classes (double plant, endrow, storm damage, and waterway) were grouped into a sixth ”none-selected” class to represent non-health-related conditions. This selection ensured our model targets issues directly impacting crop vitality.

To prepare the supervised dataset, all images and their corresponding masks were resized

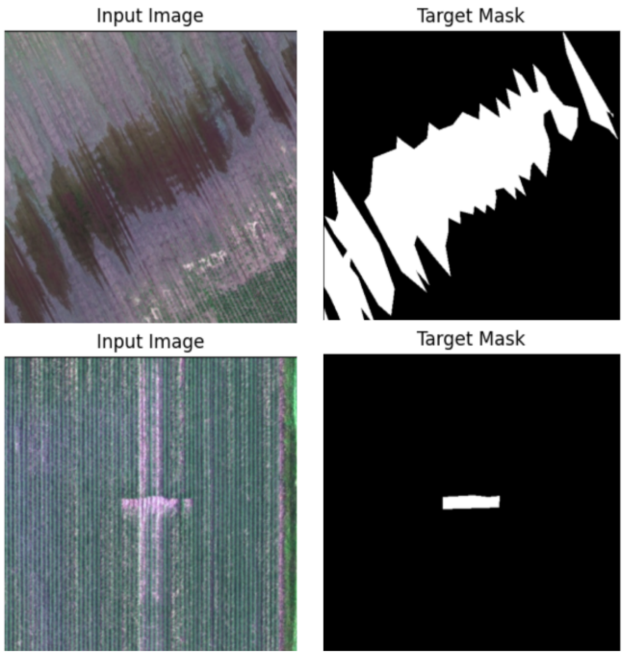

Figure 4.2: Example Stress type of planter skip and Water bodies

to 256×256 pixels. For each of the five selected classes, as well as a "none-selected" class (comprising waterway, endrow, storm damage, and double plant), we included almost 5,824 samples across classes to ensure class balance. This resulted in a total of 6 labeled samples.

Dataset Integration & Split

We consolidated our supervised dataset into training, validation, and test subsets, maintaining class balance across all six categories. The split ratios were 75% for training, 5% for validation, and 20% for testing.

Table 4.1: Supervised dataset split (6 balanced classes)

Subset	Count	Percentage
Training	4,369	75%
Validation	291	5%
Test	1,164	20%
Total	5,824	100%

4.2.2 Raw Image Dataset

Description

The raw image dataset comprises high-resolution imagery, capturing 261 full-field images at 10 cm resolution across 54 fields from 2017 to 2020. These images originate from 2–7 flights per field, forming a time series of data. Each image is identified by a path in the format raw\(field_id)\(flight_id)_(color_channel) · (tif), where the color channel includes RGB and NIR bands. Some fields also include Sentinel satellite imagery at 10 m resolution.

Unlike the supervised dataset, the raw images remain unprocessed, making them suitable for tasks such as generating NDVI maps or assessing the model's robustness on real-world data. We primarily used these images to simulate the UAV data stream in our Kafka pipeline, reducing their size from approximately 200 MB to around 10 MB for real-time processing.

Preprocessing

To evaluate our data processing pipeline under UAV-like conditions, we generated simulated image sequences. For each raw image, we created 120-frame sequences by applying subtle, pixel-level rotations. This approach simulates gradual UAV drift and provides realistic streaming data for performance testing, An example 4.3 is shown below, with the first frame and the 90th frame of the sequence for a selected field.

Figure 4.3: Example of simulated UAV sequence: the first and 90th frames of the same field after applying pixel-level drift.

4.3 Machine Learning Model Design and Training

Our crop health analysis system employs a dual-output U-Net architecture optimized for both computational efficiency and predictive accuracy. The model accepts 256×256 pixel NDVI (Normalized Difference Vegetation Index) inputs a critical choice for enhancing vegetation stress detection, and simultaneously produces segmentation masks and multiclass predictions through dedicated output branches. Figure 4.4 illustrates the complete architecture with its key components.

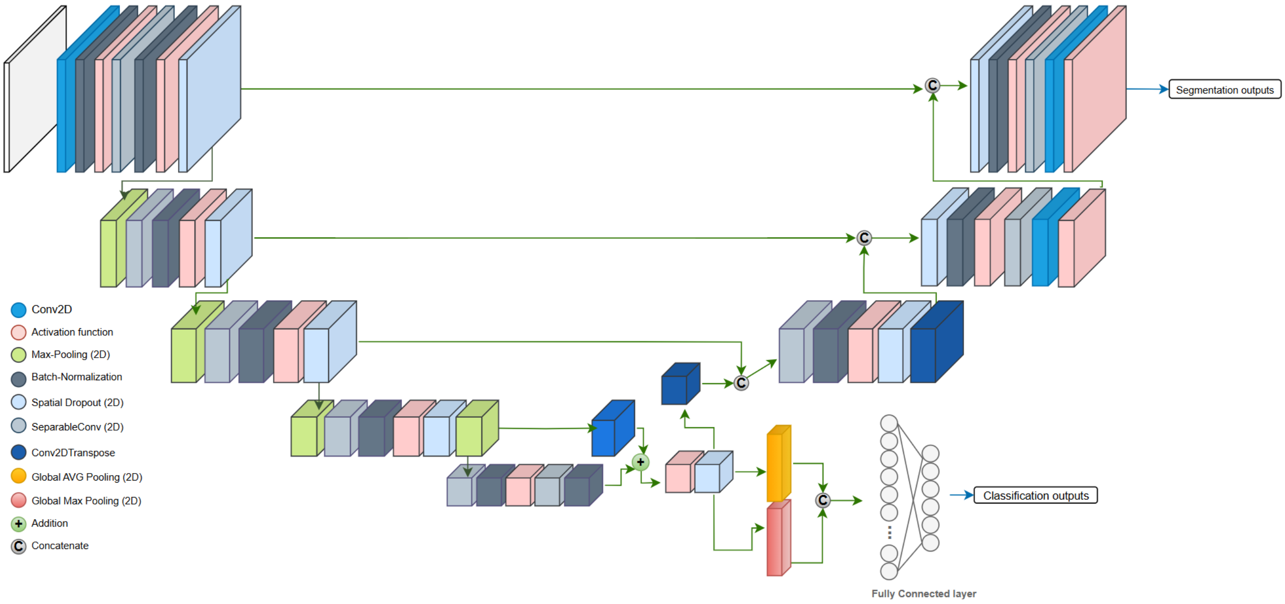

Figure 4.4: Proposed Multi-Task U-Net Architecture .

4.3.1 Architecture Overview

The network processes NDVI inputs through four progressive stages:

- Encoder Blocks (Downsampling): Four sequential blocks with separable convolutions (24, 48, 96, 192 filters) and 2×2 max pooling
- Bottleneck: Residual block with 384 depthwise separable filters and spatial dropout (20%)
- Decoder Blocks (Upsampling): Four transposed convolution blocks (192, 96, 48, 24 filters) with skip connections

4.3.2 Key Architectural Features

- Multi-Scale Feature Fusion: Skip connections between encoder and decoder at matching resolution levels
- Depthwise Separable Convolutions: 75% parameter reduction vs standard convolutions
- Dual Output Heads:
  - Segmentation Head: 1×1 convolution with sigmoid activation.
  - Classification Head: Global pooling 128-unit dense layer 6 class output with softmax activation.

4.3.3 Training Procedure

The model was trained with several specialized components:

- Loss Functions:
  - Segmentation: Hybrid loss (70% Dice Coefficient + 30% weighted BCE)
  - Classification: Binary cross-entropy with class weights
- Optimization: Adam optimizer (learning rate = 0.001)
- Regularization: Spatial dropout (10-20%) + 25% dropout rate in classification part.

4.4 Kafka-based Data Pipeline Implementation

To support real-time ingestion, preprocessing, model inference, and archival, we built an end-to-end streaming pipeline using Apache Kafka:

- Kafka Producers: Simulated UAVs send RGB, NDVI and Model Predictions frames via dedicated Kafka topics.
- Kafka Broker: Handles message queuing and delivery.
- Kafka Consumer: Receives image data, performs preprocessing and inference, then publishes prediction results.
- Storage: MongoDB stores historical original captured data.
- Frontend Dashboard: Displays streamed predictions using React and WebSockets.

Kafka's scalability allows multiple UAV streams to be handled concurrently without major delays, ensuring real-time throughput even as producer count increases.

4.5 Evaluation Metrics and Validation Approach

To effectively evaluate our multi-task model, we adopted a dual-metric strategy aligned with the model's architecture: one metric for the segmentation output and another for the classification output. Each component plays a critical role in identifying and characterizing crop anomalies in real-time.

4.5.1 Segmentation Evaluation: Dice Coefficient

For the segmentation head, which produces a binary mask indicating the location of anomalies (such as drydown or planter skips), we use the Dice Coefficient as our primary evaluation metric. This metric is particularly well-suited for image segmentation tasks because it emphasizes the overlap between the predicted and ground truth masks, making it more informative than simple pixel accuracy—especially when dealing with imbalanced datasets.

The Dice coefficient is calculated as follows:

$$\text{Dice} = \frac{2 \cdot |P \cap G|}{|P| + |G|} = \frac{2TP}{2TP + FP + FN} \tag{4.1}$$

Where:

- P is the set of predicted pixels,
- G is the set of ground truth pixels,
- TP = True Positives,
- FP = False Positives,
- FN = False Negatives.

This coefficient ranges from 0 to 1, where a value of 1 indicates perfect overlap.

4.5.2 Classification Evaluation: Accuracy, Precision, Recall, and Confusion Matrix

The classification head is responsible for predicting the type of anomaly (e.g., nutrient deficiency, drydown). For this multi-label classification task, we use a combination of the following metrics:

- Accuracy — measures the overall proportion of correct predictions.
- Precision — indicates how many of the predicted positive cases were actually correct.
- Recall — measures how many actual positive cases were correctly identified.
- Confusion Matrix — provides a class-by-class breakdown of correct and incorrect predictions.

These metrics provide a comprehensive view of the classifier's behavior, especially in cases where some classes might be harder to detect than others.

4.5.3 Real-time and Scalable Validation

In our real-time setup, each incoming UAV image (containing RGB and NDVI bands) is processed on the fly. The model produces predictions, and evaluation metrics are computed per image and aggregated over time. As the number of simulated UAVs scales from 5 to 30, we monitor both predictive accuracy and system performance to ensure scalability under fixed computational resources.

This streaming validation approach enables continuous performance monitoring and ensures that the model remains effective even as the volume of incoming data increases.

4.6 Conclusion

This chapter detailed the methodology for our real-time crop monitoring system, combining UAV imagery with deep learning and Kafka-based streaming. Key components included:

- A curated dataset of 5,824 samples across 6 stress types
- An efficient dual-output U-Net processing 256×256 NDVI inputs
- A scalable Kafka pipeline handling multiple UAV streams

The architecture balances computational efficiency with detection accuracy through depth-wise separable convolutions and multi-task learning. Our preprocessing pipeline transformed raw field images while preserving critical agricultural features. The following chapter evaluates this system's performance using Dice coefficient and precision/recall metrics.

CHAPTER 5

EXPERIMENTAL SETUP AND RESULTS

This chapter presents the experimental design, system configuration, and evaluation results for the proposed real-time crop monitoring pipeline. The primary objective is to assess the system's performance under realistic UAV deployment conditions, both in terms of image streaming and model inference. A modular architecture was implemented using containerized services (Docker) to ensure scalability, reproducibility, and efficient resource usage across GPU-enabled hardware.

We evaluated multiple deep learning models, including U-Net variants, across various input types (NDVI, RGB, RGB+NIR) and resolutions. Key metrics such as Dice coefficient, classification accuracy, streaming latency, and throughput were measured to provide a comprehensive analysis of segmentation quality and system responsiveness. Special attention was also given to the impact of image resolution and spectral input channels on per-class segmentation performance.

5.1 System Overview and Configuration

To ensure scalable, reproducible, and efficient real-time crop monitoring, the system was deployed using containerized services on a high-performance server. The following subsections describe the hardware environment, software stack, and message streaming architecture.

1. Deployment Hardware

- Operating System: Ubuntu 24.04 LTS (64-bit)
- CPU: Intel Xeon Gold
- Memory: 128 GB RAM
- GPU: 2 × NVIDIA RTX A5000 (24 GB VRAM each, total 48 GB)

2. Containerized Architecture

The system's components were deployed via Docker and orchestrated using docker-compose. Each service runs in an isolated container to enhance modularity and fault tolerance.

- Inference Server: FastAPI-based service for prediction
- Messaging: Apache Kafka for image streaming and coordination
- Frontend: React-based dashboard for live image and mask visualization
- Database: MongoDB for storing historical data

3. Real-Time Streaming via Kafka

Each UAV stream is assigned individual Kafka topics for image types:

- uav.{ID}.images.rgb
- uav.{ID}.images.ndvi
- uav.{ID}.images.predicted

This design enables easy scaling across multiple UAVs and data types. The dashboard consumes these topics in real time and updates via WebSocket endpoints.

5.2 Results

We evaluated the system under multiple configurations, testing inference latency, message throughput, and segmentation/classification quality across models and image types.

1. Kafka Throughput and Streaming Latency

The transmission of image frames via Kafka was tested under realistic UAV sequence simulations. Despite large image sizes, the system achieved reliable delivery:

Table 5.1: Average Kafka image transmission time per frame.

Image Type	Kafka Transmission Time
RGB	< 1.0 s
NDVI	0.5 – 2.5 s

MongoDB writes for historical raw data exhibited intermittent delays, especially under high throughput. Although it did not affect inference directly, it introduced a slight delay of 5s.

2. Inference Latency Across Producers

We tested the system with varying numbers of Kafka producers. Results showed minimal degradation in inference speed up to 15 producers, suggesting sufficient parallelism and system resilience.

Table 5.2: Model inference time under varying producer counts. Higher producer counts lead to increased inference latency, especially for U-Net (RGB+NIR).

Producers	Model Type	Average Inference Time
1–5	Light U-Net (NDVI)	3.1 – 4.2 s
15	Light U-Net (NDVI)	3.9 – 4.8 s
15–30	Light U-Net (NDVI)	7.0 – 8.0 s
1–5	U-Net (RGB+NIR)	9 – 11.1 s
15	U-Net (RGB+NIR)	15.2 – 17 s
15–30	U-Net (RGB+NIR)	22 – 25 s

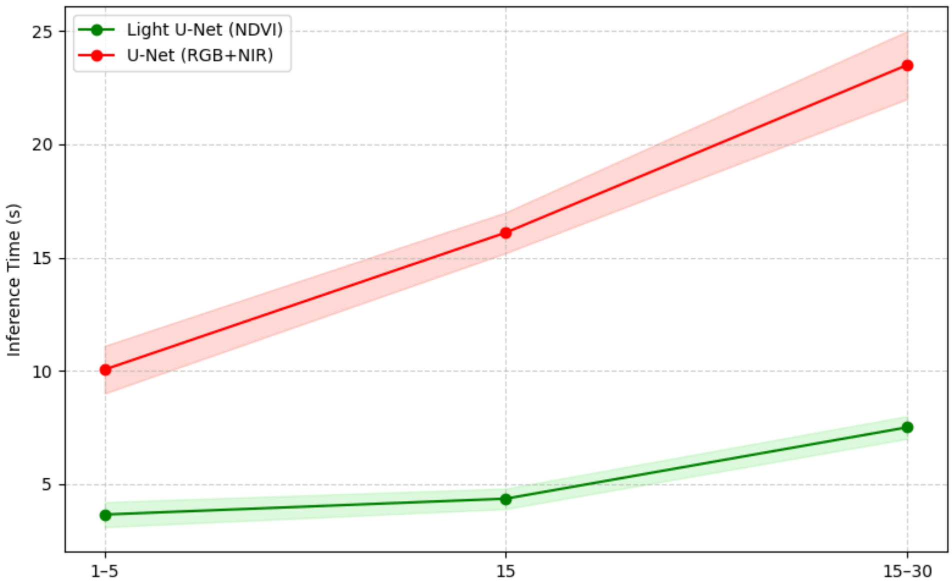

Figure 5.1: Inference latency ranges for Light U-Net and U-Net models across different Kafka producer counts. Shaded regions represent min-max latency.

Figure 5.1 illustrates the impact of Kafka producer count on inference latency for two different models. As the number of producers increases, the latency for U-Net (RGB+NIR) rises significantly due to the added computational load and input complexity. In contrast, Light U-Net (NDVI) exhibits only a modest increase, suggesting it is more suitable for deployments requiring fast response under high-throughput conditions.

The Dashboard for testing simulation that can views RGB and NDVI and view metadata per UAV stream, including image resolution, spectral bands, processing times, and detected crop stress types. Stress classes are listed as text below each.

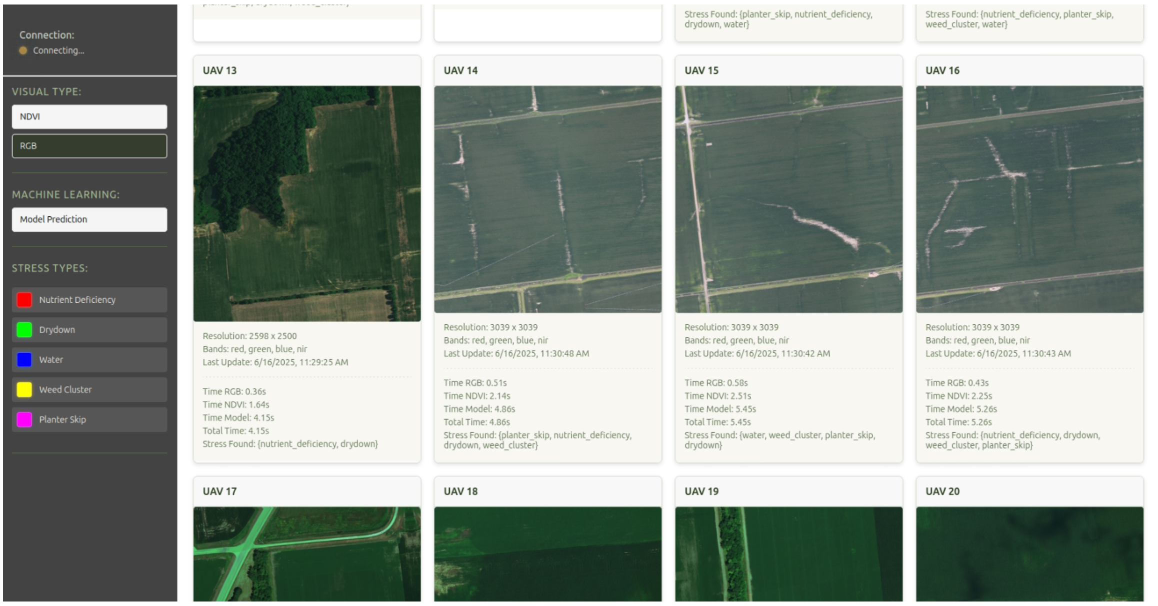

Figure 5.2: Dashboard interface for real-time UAV monitoring.

3. Segmentation and Classification Performance

Our U-Net models were trained on multiple input modalities and evaluated using Dice coefficient and classification accuracy. The RGB+NIR model achieved the best overall performance on both training and test sets.

Table 5.3: Model performance comparison across input types, resolution, and backbones.

Input Type	Resolution	Dice (Train/Test)	Class Acc. (Train/Test)
NDVI	256×256	0.57 / 0.52	0.94 / 0.91
RGB	256×256	0.62 / 0.54	0.95 / 0.92
RGB + NIR	256×256	0.66 / 0.58	0.96 / 0.92
NDVI	512×512	0.51 / 0.50	0.90 / 0.92
RGB	512×512	0.62 / 0.60	0.95 / 0.92
RGB + NIR	512×512	0.72 / 0.70	0.97 / 0.95

To gain deeper insight into model behavior, we examined class-wise segmentation performance for the RGB+NIR model at both 256×256 and 512×512 resolutions. Table 5.4 presents the Dice coefficient per semantic class: nutrient deficiency, drydown, water, weed cluster, and planter skip.

To gain deeper insight into classification effectiveness, we also computed precision and recall for the best-performing model configuration—RGB+NIR input at 512×512 resolution. This model achieved:

- Precision: 0.9477
- Recall: 0.9482

Table 5.4: Dice coefficient per class for RGB+NIR models at different input resolutions.

Class	RGBN-512 Dice	RGBN-256 Dice
Nutrient Deficiency	0.5797	0.5691
Drydown	0.6027	0.6062
Water	0.8293	0.7107
Weed Cluster	0.5861	0.5750
Planter Skip	0.7779	0.5030

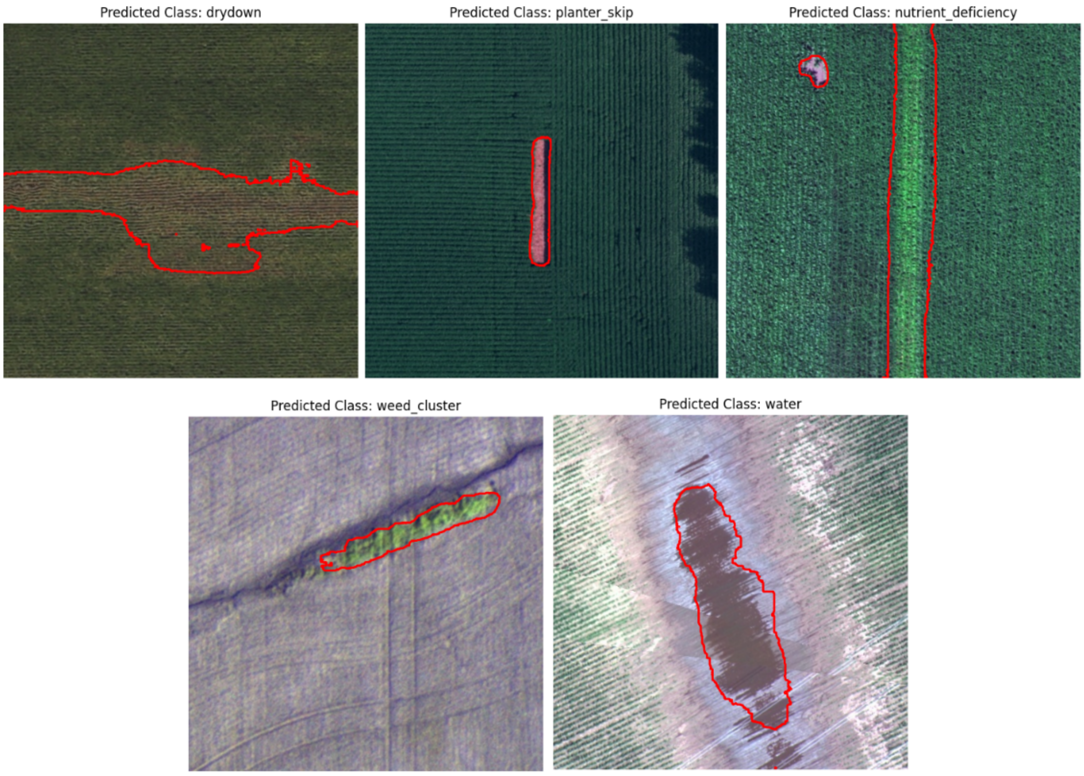

Figure 5.3: Example predictions for each target class: nutrient deficiency, drydown, water, weed cluster, and planter skip. The overlay masks (in color) illustrate the model's segmentation output on RGB input.

These metrics reflect the model's strong ability to identify true positive crop stress conditions while maintaining low false positive rates. Figure 5.4 shows the confusion matrix for this configuration, illustrating high classification fidelity across most stress classes.

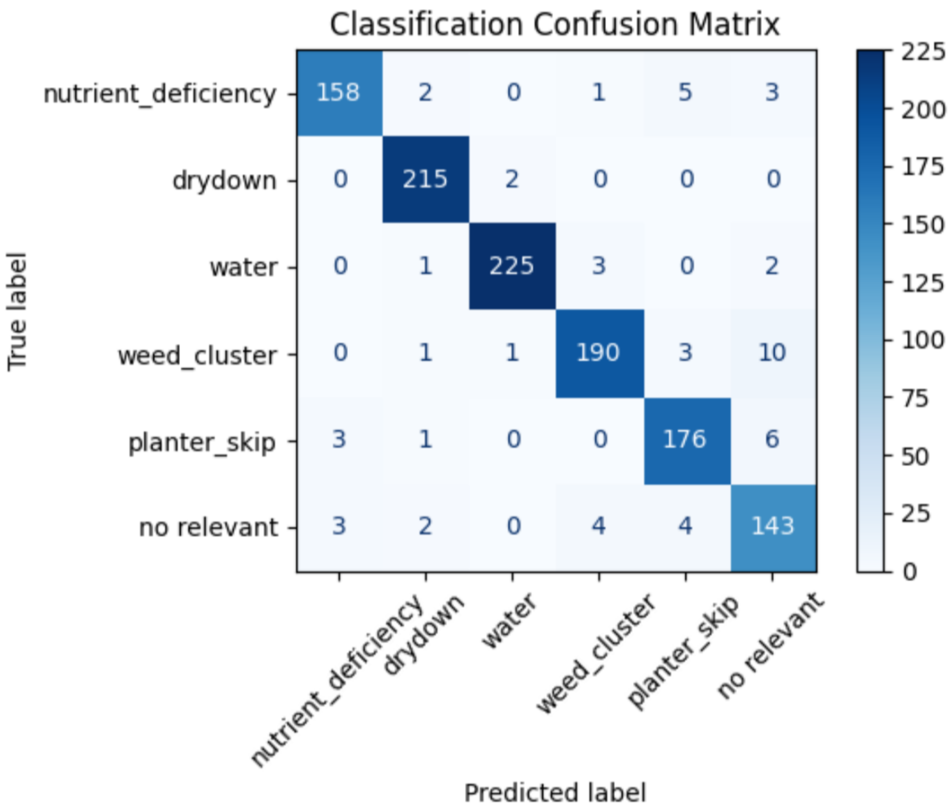

Figure 5.4: Confusion matrix for the RGB+NIR model at 512×512 resolution. The diagonal dominance indicates strong agreement between true and predicted labels.

5.3 Discussion

The experimental results demonstrate a trade-off between system speed and segmentation accuracy depending on the choice of input type, model complexity, and resolution. Lightweight models, such as the Light U-Net using NDVI input, offered low inference latency and good classification accuracy. However, segmentation quality was lower, particularly for fine-grained or spectrally complex classes.

In contrast, models trained on combined RGB and NIR inputs at higher resolutions (512×512) performed better in capturing detailed crop stress patterns. The RGB+NIR U-Net achieved a Dice score of 0.70 on the test set, indicating stronger generalization and improved detection of spatially localized classes such as planter skips and water. However, this improvement came with a noticeable increase in inference time, especially when the number of Kafka producers increased beyond 15. These delays reflect the computational demands of processing higher-resolution, multi-channel images in real time.

Kafka was used as the backbone for real-time image streaming and system coordination. It proved reliable under varying load conditions, maintaining transmission times below 2.5 seconds per image in most cases. This confirms its effectiveness for UAV-based data pipelines requiring low-latency delivery. However, we observed that when inference was pushed toward the producer side (i.e., simulated UAVs acting as producers), performance became more sensitive to the number of producers and available GPU resources.

Finally, segmentation performance varied across classes. High Dice scores for water and planter skips indicate that distinct spatial or spectral features are captured well, especially at higher resolutions. In contrast, lower scores for nutrient deficiency and weed clusters highlight the challenges of segmenting classes with subtle or overlapping patterns.

5.4 Key Findings

- Model performance scales with input complexity: Combining RGB and NIR channels improves segmentation performance, especially at 512×512 resolution.
- Real-time streaming with Kafka is effective: Kafka handled high-frequency image ingestion with low latency and supported multi-producer scenarios efficiently.
- Inference placement affects scalability: Shifting inference to the producer side can reduce response time per stream but may overload GPU resources under high producer counts.
- Higher resolution boosts detection of narrow or spatially structured classes: Planter skips and water regions were segmented more accurately at higher input resolutions.

5.5 Implications for Precision Agriculture

These findings highlight the potential of deep learning–based pipelines to support precision agriculture by providing real-time, high-resolution insights into crop health. In particular:

- Learned features from RGB+NIR models outperformed traditional vegetation indices like NDVI or NDWI for detecting certain stress conditions, such as water-related stress.
- The system's ability to segment narrow and spatially distinct patterns (e.g., planter skips) with high accuracy suggests that these models can assist in early detection and targeted intervention strategies.
- Kafka-based streaming provides a practical foundation for integrating UAVs in large-scale crop monitoring setups, ensuring timely data processing across distributed farms or fields.

5.6 Strengths, Limitations, and Challenges

Strengths

- The modular, containerized system design ensures easy deployment and scalability.
- Lightweight models like Light U-Net (NDVI) showed competitive classification accuracy, enabling deployment on resource-constrained UAVs.

- Kafka enabled efficient real-time streaming across multiple UAVs, demonstrating robustness under high-load simulations.

Limitations

- Segmentation performance declined for classes with overlapping features or irregular boundaries, such as weed clusters and nutrient deficiency.
- MongoDB write delays, though not critical, introduced a minor lag in historical data storage under high-throughput conditions.
- The system does not currently support multilabel outputs, which limits its ability to detect co-occurring stress conditions within the same region.

Challenges

- Improving segmentation for visually ambiguous or spectrally similar classes remains difficult without introducing additional contextual cues or temporal data.
- Maintaining inference performance while scaling to more UAVs requires either GPU scheduling or model compression strategies.
- Incorporating domain-specific constraints (e.g., growth stages, weather patterns) could help improve class discrimination and reduce false positives in operational settings.

5.7 Conclusion

This chapter presented a comprehensive evaluation of a containerized real-time crop monitoring system deployed under simulated UAV field conditions. Through rigorous experimentation, we established critical trade-offs between segmentation accuracy, computational efficiency, and system scalability. The RGB+NIR U-Net model at 512×512 resolution demonstrated superior segmentation performance (Dice=0.70), particularly for spatially distinct classes like planter skips and water stress. However, this came at significant computational cost, with inference latency increasing up to 25 s under high-load scenarios ($>$ 15 producers). Conversely, lightweight models like NDVI-based U-Net provided viable low-latency alternatives (3.1–4.8 s) for classification-focused deployments while sacrificing segmentation granularity.

Kafka proved highly effective for real-time image streaming, maintaining a transmission latency of under 2.5 seconds, especially on the consumer side, where it demonstrates the ability to handle multiple Kafka topics almost simultaneously.

CHAPTER 6

CONCLUSION AND FUTURE WORK

6.1 Summary of Contributions

This thesis presented the design, implementation, and evaluation of a real-time crop monitoring system that combines UAV multispectral imagery, deep learning, and scalable streaming infrastructure. Addressing key limitations in existing precision agriculture systems, this work contributed the following:

- A lightweight dual-output U-Net model tailored for NDVI inputs, capable of simultaneously performing binary segmentation and multiclass classification of crop stress conditions.

- A curated and balanced dataset derived from the Agriculture-Vision 2021 dataset, targeting five critical crop stress types (nutrient deficiency, drydown, weed cluster, water, and planter skip), enabling effective model training and evaluation.

- A Kafka-based streaming pipeline simulating real-time UAV data ingestion and inference. The system handles concurrent producer streams and delivers processed output—including NDVI maps, stress masks, and metadata—for real-time visualization and storage.

- Empirical validation of the proposed model and infrastructure. Results showed that the lightweight U-Net achieved competitive segmentation accuracy (as measured by Dice coefficient) while maintaining real-time throughput, even as the number of producers scaled from 5 to 30.

Together, these contributions demonstrate the feasibility of real-time, scalable crop stress detection using modern geospatial data and efficient deep learning models.

6.2 Future Research Directions

While the system developed in this thesis shows promise for real-time monitoring, several avenues remain open for enhancement and broader adoption:

- Integration of temporal features: Currently, predictions are made on single NDVI frames. Incorporating temporal sequences (e.g., through ConvLSTM or 3D CNNs) could improve stress detection by capturing crop development trends over time.
- Multiple Consumers for Parallel Inference: Future work could explore deploying multiple Kafka consumers, each running different models or processing different image types (e.g., NDVI, RGB, multispectral). This approach would distribute the computational load, enable parallel inference, and allow ensemble predictions or task-specific pipelines, improving both scalability and prediction robustness.
- Alternative Streaming Platforms (e.g., Apache Spark): While Apache Kafka provides robust messaging and decoupling, future systems could integrate stream processing frameworks such as Apache Spark Structured Streaming to support complex, stateful operations and advanced analytics. Spark's distributed architecture may also enhance scalability and allow tighter integration with downstream machine learning and database components.
- Transfer learning for underrepresented crops: Expanding the system to support various crops and regions through fine-tuning and transfer learning could increase its generalizability and real-world impact.

In summary, this thesis lays the foundation for an operational real-time monitoring system in smart agriculture. With further development, it can evolve into a robust decision-support tool for precision farming at scale.

BIBLIOGRAPHY

[1] Kamran Razzaq and Mahmood Shah. Machine learning and deep learning paradigms: From techniques to practical applications and research frontiers. Computers, 2025.

[2] A. Monteiro, S. Santos, and P. Gonçalves. Precision agriculture for crop and livestock farming-brief review. Animals (Basel), 11(8):2345, Aug 2021.

[3] N. Zhang, M. Wang, and N. Wang. Precision agriculture—a worldwide overview. Computers and Electronics in Agriculture, 36(2-3):113–132, 2020.

[4] Sjaak Wolfert, Lan Ge, Cor Verdouw, and Maarten-Jan Bogaardt. Big data in smart farming – a review. Agricultural Systems, 153:69–80, 2017.

[5] Anne-Katrin Mahlein. Plant disease detection by imaging sensors – parallels and specific demands for precision agriculture and plant phenotyping. Plant Disease, 100(2):241–251, 2016.

[6] E. Raymond Hunt, W. Don Hively, Scott J. Fujikawa, David S. Linden, Craig S.T. Daughtry, and Gregory W. McCarty. Acquisition of nir-green-blue digital photographs from unmanned aircraft for crop monitoring. Remote Sensing, 4(11):2909–2924, 2019.

[7] R. Ehsani and C. Divaker Durairaj. Spatial food and agricultural data. Encyclopedia of Life Support Systems (EOLSS), 2000. Available online at https://www.eolss.net/sample-chapters/c10/E5-17-03-02.pdf.

[8] Sumaira Ghazal, Arslan Munir, and Waqar S. Qureshi. Computer vision in smart agriculture and precision farming: Techniques and applications. Artificial Intelligence in Agriculture, 13:64–83, September 2024.

[9] Farmonaut Team. Farmonaut: Geospatial analytics for sustainable agriculture and forestry. https://www.farmonaut.com, 2025. Accessed on 2025-05-09.

[10] Mrutyunjay Padhiary, Pranjal Saikia, Pankaj Roy, Nabiul Hussain, and Kundan Kumar. A review on advancing agricultural efficiency through geographic information systems,

remote sensing, and automated systems. Cureus Journal of Engineering, 2:es44388–024–00559–7, 2025.

[11] R.P. Sishodia, R.L. Ray, and S.K. Singh. Applications of remote sensing in precision agriculture: A review. Remote Sensing, 12(19):3136, 2020.

[12] David J. Mulla. Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps. Biosystems Engineering, 114(4):358–371, 2013.

[13] GeeksforGeeks. Supervised vs unsupervised vs reinforcement learning. https://www.geeksforgeeks.org/supervised-vs-reinforcement-vs-unsupervised/. Accessed: 2025-05-15.

[14] Lumenalta. The 5 different types of machine learning paradigms, explained. https://lumenalta.com/insights/5-types-of-machine-learning. Accessed: 2025-05-15.

[15] T. Zhang, Z. Xu, J. Su, Z. Yang, C. Liu, W.-H. Chen, and J. Li. Ir-unet: Irregular segmentation u-shape network for wheat yellow rust detection by uav multispectral imagery. Remote Sensing, 13(19):3892, 2021.

[16] E. Chakhvashvili, M. Machwitz, M. Antala, and et al. Crop stress detection from uavs: best practices and lessons learned for exploiting sensor synergies. Precision Agric, 25:2614–2642, 2024.

[17] Jun Wang, Yan Wang, Guang Li, and Zheng Qi. Integration of remote sensing and machine learning for precision agriculture: A comprehensive perspective on applications. Agronomy, 14(9):1975, 2024.

[18] Fatemeh Alaieri. Precision agriculture based on machine learning and remote sensing techniques. Engineering, Technology & Applied Science Research, 14(3):14206–14211, June 2024.

[19] M. T. Linaza, J. Posada, J. Bund, P. Eisert, M. Quartulli, J. Döllner, A. Pagani, I. G. Olaizola, A. Barriguinha, T. Moysiadis, and L. Lucat. Data-driven artificial intelligence applications for sustainable precision agriculture. Agronomy, 11(6):1227, 2021.

[20] I. Neri, S. Caponi, F. Bonacci, G. Clementi, F. Cottone, L. Gammaitoni, S. Figorilli, L. Ortenzi, S. Aisa, F. Pallottino, and M. Mattarelli. Real-time ai-assisted push-broom hyperspectral system for precision agriculture. Sensors, 24(2):344, 2024.

[21] A. Theofilou, S. A. Nastis, M. Tsagris, S. Rodriguez-Perez, and K. Mattas. Design and implementation of a scalable data warehouse for agricultural big data. Sustainability, 17(8):3727, 2025.

[22] Tej Bahadur Shahi, Cheng-Yuan Xu, Arjun Neupane, and William Guo. Machine learning methods for precision agriculture with uav imagery: A review. Electronic Research Archive, 30(12):4277–4317, 2022.

[23] J.A.O.S. Silva, V.S.d. Siqueira, M. Mesquita, L.S.R. Vale, T.d.N.B. Marques, J.L.B.d. Silva, M.V.d. Silva, L.N. Lacerda, J.F.d. Oliveira-Júnior, J.L.M.P.d. Lima, and et al. Deep learning for weed detection and segmentation in agricultural crops using images captured by an unmanned aerial vehicle. Remote Sens., 16:4394, 2024.

[24] Lala El Hoummaidi, Abdelkader Larabi, and Khan Alam. Using unmanned aerial systems and deep learning for agriculture mapping in dubai. Heliyon, 7(10):e08154, 2021.

[25] Do Vu, Lương Tuấn, Minh Hoang, and Tuấn Tăng. Real-time flood segmentation from uav imagery using deep learning with apache kafka and spark structured streaming. 07 2024.

[26] M. Li, R. R. Shamshiri, C. Weltzien, and M. Schirrmann. Crop monitoring using sentinel-2 and uav multispectral imagery: A comparison case study in northeastern germany. Remote Sensing, 14(17):4426, 2022.

[27] Tabitha KM. Integration of drones and satellite imaging in crop monitoring and precision farming. Adv Crop Sci Tech, 12:754, 2024.

[28] Narayanaswamy Jeevan, Anil K, Chethan Kumar K B, Nandini H M, and Yerradoddi Sindhu Sree. CROP HEALTH MONITORING USING DRONES: A MODERN APPROACH TO PRECISION AGRICULTURE. 2024.