OPEN ACCESS
International Journal of Innovations in Science & Technology

Agentic AI for Autonomous Soil and Fertilization Management for Agriculture Sustainability

Muhammad Amjad[1], Rabia Tehseen[1*], Kashif Nasr[1], Fraz Aslam[1], Maham Mehr Awan[1], Uzma Omer[2]

[1]Department of Computer Science, University of Central Punjab, Pakistan.

[2]Department of Computer Science, University of Education, Pakistan.

***Correspondence**: rabia.tehseen@ucp.edu.pk

Citation | Amjad. M, Tehseen. R, Aslam. F, Awan. M. M, Omer. U, "Agentic AI for Autonomous Soil and Fertilization Management for Agriculture Sustainability", IJIST, Vol. 7 Issue. 4 pp 2997-3017, November 2025

Received | October 27, 2025 **Revised |** November 19, 2025 **Accepted |** November 25, 2025 **Published |** November 30, 2025.

Soil fertility loss and excessive chemical fertilization are major environmental and economic issues in developing regions such as Punjab, Pakistan. This paper proposes an Agentic AI framework for autonomous soil and fertilization management that combines (i) IoT soil sensing and drone-based crop monitoring for real-time perception, (ii) predictive modelling for short-horizon nutrient and moisture forecasting, and (iii) multi-agent reinforcement learning (MARL) for adaptive decision-making. The system operates with operational autonomy, executing daily management decisions without routine human-in-the-loop control. Agronomic expert knowledge is incorporated only offline as safety constraints and initialization priors (e.g., allowable nutrient ranges and stress-avoidance rules) to bound the action space and prevent unsafe behavior, rather than to prescribe actions. Experiments were conducted across two seasons at two sites (Sheikhupura and Multan) under four treatments: Farmer Practice (FP), Rule-Based Control (RBC), Machine Learning Predict (ML-Predict), and the proposed Agentic AI. Results show that Agentic AI reduces nitrogen fertilizer use while maintaining/improving yield proxy and improving soil indicators (including residual nitrate reduction and improved Soil Health Index). We also analyze irrigation outcomes as a sustainability objective and show how water usage must be treated as a constrained or multi-objective term in the reward function to avoid over-irrigation. Overall, the framework supports scalable, data-driven soil management with bounded autonomy, preserving expert-defined agronomic safety.

Keywords: Multi-Agent Reinforcement Learning; Rule-Based Control; Farmer Practice; Agentic AI; Soil Health Index

Introduction:

The introduction explains how traditional agriculture is evolving into a smart, autonomous, and adaptive system thanks to the introduction of Agentic Artificial Intelligence (AI). It highlights an agentic AI that can enhance efficiency, sustainability, and decision-making, especially regarding soil and fertilization management. This section brings out the necessity of smart, information-based solutions to sustainable agriculture and presents the novelty of creating agentic AI systems that can learn and make independent decisions to implement soil health and optimize fertilizers.

The application of the Agentic Artificial Intelligence (AI) in the agricultural field is changing the conventional system of agriculture into a smart, self-learning, and adaptable system. The reason is that the agentic AI systems are not focused on prediction only: it is possible to create autonomous or autonomous choices, learning, and proactive alteration towards the surrounding environment. Recent reports demonstrated that these systems can be more efficient, less dependent on humans, and intelligent and independent thinking, as well as coordination of farm agents [1][2]. Autonomy, data-driven intelligence, and a combination of the three components of reasoning can allow agentic AI to be a chance to establish sustainable and resilient agricultural systems.

The use of soil and fertilizers is one of the most important issues in modern agriculture, and inappropriate use of fertilizers, incorrect soil analysis, and slow reaction to decisions are frequent factors that contribute to the decreased crop yields, imbalance of nutrients, and soil erosion. Conventional farming methods are based on manual observation and scheduled regimes, which fail to change dynamically to the changing soil conditions. Consequently, overuse of fertilizers, erosion of nutrients, wastage of water, and destruction of the environment are common occurrences. In order to overcome these problems, Agentic AI can be a revolutionary technology since it is not simply a prediction of soil conditions, but it autonomously thinks, adjusts, and makes real-time decisions. In contrast to traditional models of AI, where human intervention is obligatory at each stage, agentic systems always monitor the state of the soil environment, handle nutrient cycling, and automatically calculate the best irrigation and fertilization plans without human interference. This renders them perfect in the management of all nutrients, particularly in settings where the soil parameters are varied. The new studies indicate that AI-based soil surveillance and intelligent use of fertilizers have a great impact on improving soil health, lessening chemical wastes, and maximizing harvest [3][4]. Thus, the idea of implementing Agentic AI in the soil and fertilizers is a vital chance to build data-driven and adaptive, autonomous, and environmentally friendly practices in agriculture.

The proposed study is aimed at creating an Agentic AI-powered system of autonomous soil and fertilization control to assist sustainable agriculture. The domain includes the entire process of soil sensing to smart decision-making and self-actuation. This paper starts with the development of the IoT mobile-based soil sensors that will be integrated at the start of the study to obtain real-time information about soil moisture, pH, NPK, temperature, and organic matter. This information is handled and processed through an Agentic AI model that has the capability of autonomous reasoning, detection of anomalies, and controlling nutrients. This study also involves designing a reinforcement learning based decision engine that identifies the best irrigation and fertilization strategies without involving human intervention. The system will be tested on how it will decrease the use of water, decrease the wastage of fertilizers, improve the health of soils, and increase crop yields. Also, within the scope, there will be the development of a sustainability evaluation module to quantify the long-term environmental impact. This research is also confined to nutrient management of soils and automated fertilization procedures in the agricultural environment; it does not entail pest identification, weather prediction, and crop illness diagnosis. In general, the scope is expected to show how Agentic AI will empower precision agriculture by providing autonomous, efficient, and environmentally sustainable solutions to soil management.

The novelty of this work lies in the development of an operationally autonomous, agentic AI framework for soil and fertilization management that goes beyond prediction or recommendation-based systems. Unlike existing approaches that rely on fixed rules or supervised learning models, the proposed framework employs multi-agent reinforcement learning to directly control fertilization and irrigation decisions under real-world field conditions. Agronomic expertise is incorporated only as safety constraints and initialization priors to bound the action space, rather than as prescriptive decision rules. This enables the system to learn context-sensitive strategies while maintaining agronomic safety, representing a clear departure from conventional AI-assisted agriculture. The structure will study the soil makeup and the amount of nutrients present in the soil, and also the environment, so that the most favourable methods of fertilization can be determined. The key contribution of this research is the creation of a self-learning agentic AI model that is able to maximize the use of fertilizer and keep the soil healthy with the least level of human input. Other studies on agentic and AI-controlled sustainable agriculture have demonstrated potential outcomes in enhancing efficiency and minimizing environmental impact [5][6]. We hope that the following results are possible regarding improved use of fertilizers, improved soil fertility, and a scalable design that can be used to support sustainable agricultural practices using agentic AI.

Literature Review:

The adoption of Artificial Intelligence (AI) in farming has revolutionized the appearance of precision farming and decision-making. The initial studies focused on how machine learning and deep learning could be used to optimize crop yield and farm management [7][8]. These works formed the basis of Agriculture 5.0, which combines AI, the Internet of Things (IoT), and data analytics to improve productivity and sustainability [9][10]. This research has also shown the way in which an AI-based soil management system positively influenced nutrient application and water utilization efficiency.

As the idea of smart farming developed, studies started to point towards the ethical, environmental, and social factors. As an example, the authors of the article [11] considered the issue of gender and indigeneity in AI-driven agriculture in East Africa, whereas the article [6] challenged the question of trust relations in the framework of algorithmic governance in precision agriculture. Such studies were an indicator of an increased concern with human-centered AI solutions, reflecting the trend in general [12].

The concept of agentic AI, with its autonomy, reasoning, and goal-focused decision-making, is gradually replacing traditional AI systems. According to the definitions of article [13], agentic AI can be an autonomous intelligence, which is capable of delivering complex tasks without micromanagement by a human. Equally, the article [14] presented a comprehensive summary of its applications and influences on society. Based on the level of autonomy, contextual knowledge, and the ethical implications, the article [15] proposed the conceptual taxonomy between standard and agentic AI agents.

The article [16] provided a multi-expert perspective on agentic systems, defining them as transparency, reliability, and alignment as the primary challenges concerning the application of agentic systems. Adding to this, [5] investigated the possibility of whether such systems will empower or displace human decision-making, and whether it has philosophical and managerial implications for human autonomy within AI-aided settings.

The agricultural sector has become an attractive area of implementation of agentic AI because it involves the use of context-specific, data-driven, and multi-agent coordination. In the article [1], the use of Agentic AI in autonomous decision making for the food supply chain with improved logistics and resource optimization was further expanded to utilize cloud infrastructure for food distribution in the regions based on predicted logistics.

In the article [3], the framework of Meta Ag 2.0 was introduced, which is a contextual agricultural recordkeeping with agentic intelligence, whereas article [17] introduced an edge-enabled smart agriculture framework, which is the combination of IoT, lightweight deep learning, and agentic AI, to provide context-aware farming. The same was stressed in the work of the article [9], which demonstrated that smart farming systems can be used to coordinate autonomous agents for enhancing agricultural efficiency.

Article [4] examined the intersection of AI and blockchain and depicted how the use of decentralized ledger systems improves traceability and trust in the AI-controlled farming practices. This is in line with the sustainability-oriented review by article [18] that highlighted the role of AI in sustainable agriculture.

The convergence of LLMs and agentic AI is the beginning of a new era in the automation of agriculture. The application of agentic AI in dairy science has been pioneered by [2], who utilized the LLMs in autonomous decision-making when managing the herd and analyzing data. In the same manner, article [19] presented an agriculture and enterprise AI multi-agent LLM system, which is scalable and adaptive.

Article [20] made an extensive survey of the multi-modal LLMs in the agricultural field, covering the topic of vision-language integration and RAG (retrieval-augmented generation) systems to ground knowledge. Whereas the article [21] created the AgriBuddy agentic AI system based on RAG and vision models that are specific to Bangladeshi agriculture. The application of the LLM-based agents to automating the smart farm business processes was also described by [22], and the integration of UAVs with LLMs to display real-time agricultural monitoring was demonstrated by [23].

All these innovations imply that there is a tendency towards multi-agent, language-conscious, and perceptual AI systems, which can interpret unstructured farm data and automatically plan its workflow. The autonomy of agentic AI systems is a critical issue as it increases the chance of misalignment and the inability to control the systems by the [24] risk alignment framework in agentic systems, which is concerned with transparency and human-in-the-loop design. In the article, [25] studied the application of agentic AI and LLMs to insurance decisions, but they also identified opportunities and challenges applicable to the risk management segments of agriculture.

The Article [5] warned that the growing autonomy of agentic systems may result in the loss of accountability, and the structures that would allow balancing efficiency and ethical accountability are needed. All these teachings render it important to make sure that there is a sense of trust, explainability, and human control when it comes to the application of AI in agriculture.

The future of farming is with the aid of AI technologies that will be applied to achieve climate-smart and sustainability goals. Article [26] focused on the implementation of AI-based analytics to enhance the utilization of resources and minimize the impacts of climate. In article [17], the application of AI to maximize fertilizer usage was explored, and further personalization to the process of nutrient management using agentic AI was introduced. Similarly, in the article [3], the notion of data-centric agriculture through the assistance of AI-based agents that can be used to promote flexibility in context and real-time decision-making. It can be concluded from all the above research work that adaptive learning and self-regulating capabilities can become a cornerstone to the vision of Agriculture 6.0 with fully autonomous, sustainable, and interdependent agricultural ecosystems.

Methodology:

Proposed Agentic AI Framework:

The purpose of this study is to use an experimental design based on simulation to design and experiment with a sustainable agricultural Agentic AI model. The key aim was to develop an independent system that has the ability to control soil nutrients and irrigation without human interference. The proposed framework is based on a layered architecture including Perception, Cognition, and Action layers. The system can detect data in the environment, make decisions based on the optimal course of action through machine learning, and act in a simulated farm setting, thanks to this design. The fundamental principle of this model is an agent of Reinforcement Learning (RL), which interacts with a simulated environment to learn the best farming strategies as time progresses. The Python programming language was used to implement the system, which made use of scikit-learn libraries that assist in predictive modelling and custom classes that were made to support the simulation environment.

In this work, "Agentic AI" is not used as a synonym for MARL. A standard MARL controller typically learns a direct policy mapping from observed state s_t to an action a_t (i.e., $\pi(a_t|s_t)$) that optimizes a reward. In contrast, we define Agentic AI as a system in which each agent is an orchestrator with: (1) perception (sensor fusion and state estimation), (2) deliberation (policy learning + predictive reasoning), (3) action execution (actuation with constraint checking), and (4) self-monitoring (tracking outcomes, constraint violations, and corrective fallback). Formally, each agent is modelled as:

$$\mathcal{A} = \langle \pi_\theta, \hat{s}_t, \mathcal{M}, \mathcal{C}, \mathcal{F} \rangle$$

Where π_θ is the learned MARL policy, s_t is the estimated state, M is predictive modelling (short-horizon forecasts for moisture/nutrients), C is a set of expert-defined safety constraints (e.g., allowable N ranges, moisture stress limits), and F is a fallback mechanism (safe default actions when uncertainty is high or constraints are at risk). Therefore, the proposed approach is "agentic" because it coordinates multiple tools and layers (sensing, forecasting, learning, constraints, and fallback) into a closed-loop autonomous decision pipeline, rather than relying on a single learned controller alone.

The workflow of the AAI proposed research work has been presented as shown in Figure 1.

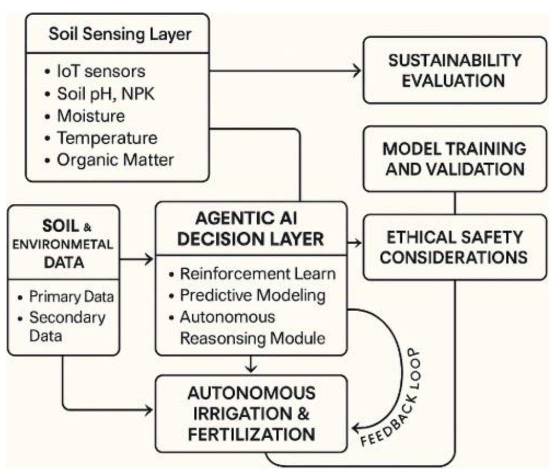

Figure 1. Methodology Diagram

The proposed Agentic AI system is evaluated using a simulation-based experimental environment that emulates soil–crop–irrigation dynamics in a controlled and reproducible manner. The simulation environment is designed to reflect the layered methodology shown in Figure 1 and is formally defined as a closed-loop decision-making system built on reinforcement learning principles.

The simulation environment is modeled as a discrete-time Markov Decision Process (MDP):

$E=\langle S, A, T, R,\gamma\rangle$

Where:

S denotes the soil–environment state space

A denotes the irrigation and fertilization action space
T represents the soil transition dynamics
R is the sustainability-aware reward function
γ is the discount factor
Each simulation step corresponds to a daily farming decision cycle.

Layer-Wise Detail:

The proposed model has been organized in a layer architecture. Each layer has been described below.

Perception Layer: (Soil Sensoring):

The perception layer is the data entry point to the system. We applied two different datasets to train the parts of the AI. The Fertilizer Prediction Dataset, which was the first dataset, had 100,000 records of soil parameters, such as Nitrogen (N), Phosphorus (P), Potassium (K), temperature, and humidity. This information guided the system on the basic agronomic rules. The second data set was the Smart Farming Data 2024 (SF24), which gave 2,200 actual sensor measurements at farms. Our simulation environment was based on this dataset so that the AI training is not tested against noise. We conducted a great deal of pre-processing and feature engineering in order to make the data useful to the AI. We initially standardized the column names of both data sets in order to have a common set of language in the system. Afterwards, we computed three derived features to provide the agent with high-level information. We then calculated the Temperature-Humidity Index (THI) to determine the amount of heat stress on the crops. Second, we have obtained the Nutrient Balance Ratio (NBR) to define whether the nitrogen concentration was proportional to other nutrients. Lastly, we developed a Soil Health Index (SHI), a composite index based on the pH and organic matter content that was also a key performance measure. These artificial capabilities enabled the agent to make decisions considering the general health of the environment and not individual sensor values.

Cognition Layer: (Decision Layer):

The decision-making agents and reinforcement learning (RL) are used to identify the best fertilization strategies. Learning Agents rely on multi-agent reinforcement learning (MARL) to coordinate fertilization mechanisms. The agents interact in a way that they equalize the allocation of resources, prevent excessive over-fertilization, and adapt to the various soil conditions. Bayesian optimization, Q-Learning by decision making agent aids in choosing the best fertilization program with the help of a group of restrictions. This layer makes predictions of the rates of nutrient depletion and nutrient demand in crops via predictive modelling (Random Forest). It incorporates an autonomous system of reasoning, a scenario, and adaptive learning. The agents are all smart agents that monitor and interfere with some of the agricultural parameters (e.g., soil nutrient agent, moisture agent, crop growth agent). Learning (RL) has been used in this layer to enable decision-making agents to decide on the most appropriate strategies to fertilize. Multi-agent reinforcement learning (MARL) is also used by the Learning Agent to coordinate the work of activities related to fertilization. The two agents interact so as to balance with regard to the distribution of resources, so that they do not over-fertilize, and also adapt to the diverse soil conditions. The Bayesian optimization and Q-learning are used by a decision-making agent to select the most effective fertilization scheme within the available limits. Predictive models (Random Forest) have the capability of determining the rate of nutrient depletion and nutrient requirement of crops. It combines a situation analysis and a learning on-the-fly autonomous reasoning system. All the agents are intelligent agents, which monitor and act on specific parameters of agriculture (e.g., soil nutrient agent, moisture agent, crop growth agent). Cognition layer is the brain of the system, which was developed based on a hybrid architecture that was developed on a combination of supervised learning and reinforcement learning.

Expert Knowledge Model:

We first attempted a pure expert-rule classifier using raw fertilizer data. Instead of using this module as a manual decision-maker, we retain it only to encode agronomic constraints (safe ranges, stress avoidance thresholds) and to initialize the learning process with reasonable priors. During operation, the Agentic AI executes decisions autonomously, while this expert model acts as a guardrail to prevent unsafe actions and as a fallback when sensor uncertainty is high.

Agentic Core (Reinforcement Learning):

A Q-Learning agent plays the central role in the decision-making process. This agent learns through trial and error, unlike the usual predictive models, which only predict the result. A State Space that included simplified soil variables (Low, Good, High) and an Action Space with six moves (Wait, Irrigate, Apply Urea, Apply DAP, Apply MOP, or Apply Balanced Fertilizer) was defined by us. The agent was conditioned to maximize some reward function. It scored positive in keeping the nutrients of the soil in the optimal range, the sweet spot, and negative on excess fertilizer application or where the crops are stressed by the lack of water.

Action Layer*: (*Autonomous Irrigation and Fertilization):

The Action Layer is the final working stage of the proposed Agentic AI system in Autonomous Soil and Fertilization Management. It makes smart AI decisions and puts them into practice on the farm. After the decision-making layer interprets soil data and finds out that there are portions of the soil that need irrigation or manure, it sends an auto command to the field machinery to do the required operation. The working unit in this system becomes the drones and IoT-based spraying machines. The physical testing with drones was not within the concerns of this research; we created our own Python simulation class, which we call Farm Environment. The program simulates the life of a real corn field. It uses real soil data loaded into the SF24 dataset in order to establish the initial conditions of each episode. The environment is a model of natural physics, such as each day the crop is fed with a little bit of N, P, and K, and the soil is getting drier through evaporation. The environment changes the soil variables when the agent chooses an action. When the agent selects the option of Irrigate, the moisture content increases, and the temperature decreases by a margin. When it selects the application of Urea, the nitrogen concentration rises. This two-way communication enables us to see the long-term repercussions of the AI decision, e.g., will saving fertilizer today result in nutrient deficiency next week?

Data Collection*:*

Dataset Description:

The following datasets will be processed in this research.

www.kaggle.com/datasets/irakozekelly/fertilizer-prediction Accessed on:11-09-2025

Total number of records: 100,000

Table 1. Dataset Specification (used for training the proposed model)

Attribute Name	Data Type	Description
Temperature	Numeric (°C)	Represents the temperature of the soil environment in degrees Celsius, influencing nutrient absorption and crop growth.
Humidity	Numeric (%)	Indicates the level of moisture in the air, which affects soil evaporation and fertilizer efficiency.
Moisture	Numeric (%)	Refers to the water content present in the soil, essential for root nutrient uptake and crop health.
Soil Type	Categorical	Specifies the type of soil (e.g., Red, Black, Sandy, Loamy, Clayey) which determines texture, fertility, and water retention.

Crop Type	Categorical	Identifies the type of crop cultivated (e.g., Cotton, Wheat, Maize, Sugarcane, Millets) based on soil and climatic suitability.
Nitrogen	Numeric (mg/kg)	Represents the amount of nitrogen available in the soil, a key nutrient promoting vegetative growth.
Potassium	Numeric (mg/kg)	Denotes the potassium content in the soil, important for improving crop resistance and overall plant metabolism.
Phosphorous	Numeric (mg/kg)	Shows the level of phosphorus in the soil, essential for root development and seed formation.
Fertilizer Name	Categorical	Indicates the recommended fertilizer type (e.g., Urea, 20-20, 14-35-14, DAP) based on nutrient balance and soil requirements.

The above dataset will have important soil and environmental parameters to be used in training the Agentic AI model that is to be used in this study. The information was gathered to forecast the best type of fertilizers and assist in self-managed soil control by making decisions in real-time. The characteristics of each attribute help in the realization of the correlation between the characteristics of soils, crop demands, and their fertilizer optimization in sustainable agriculture. This dataset contains 100,000 records.

The following datasets will be processed in this research.

Smart Farming Data 2024 (SF24): Last Visit: 11-Dec-2025

Total number of records in SF24: 2200

Table 2. Dataset Specification (used for testing and validation of the proposed model)

Attribute Name	Data Type	Description
P	Numeric (int/ppm)	Phosphorus content available in the soil, measured in parts per million (ppm). It supports root growth, flowering, and energy transfer; low P can restrict crop development even if other nutrients are adequate.
Temperature	Numeric (float, °C)	Ambient air temperature at the field in degrees Celsius. It reflects the thermal environment around plants and is critical for germination, photosynthesis, transpiration, and overall crop growth rate.
N	Numeric (int/ppm)	Nitrogen content in the soil (ppm). Nitrogen is a primary macronutrient driving leaf growth and chlorophyll production; both deficiency and excess N strongly influence yield and fertilizer planning.
Humidity	Numeric (float, %)	Relative air humidity in percentage. It indicates how much moisture is present in the air and, together with temperature, affects plant transpiration, disease risk, and heat-stress conditions.
pH	Numeric (float)	Soil pH level (acidity/alkalinity). It controls nutrient availability and microbial activity; many crops perform best in a near-neutral range, while very acidic or alkaline values can lock nutrients.
Rainfall	Numeric (float, mm)	Total rainfall in millimeters over the observation period. It represents naturally supplied water and, combined with soil moisture and irrigation, influences water balance and potential water stress.
Label	Categorical (string)	Target variable: crop type associated with the recorded soil and climate conditions (e.g., rice, wheat,

		maize, etc.). This is what your model predicts in classification/recommendation tasks.
K	Numeric (int/ppm)	Potassium content in the soil (ppm). Potassium regulates water use, disease resistance, and overall plant robustness; balanced N–P–K levels are crucial for healthy crop growth.
Soil Moisture (soil_moisture)	Numeric (float, %)	The percentage of water content in the soil at the sampling time. It directly reflects how wet or dry the root zone is and is essential for irrigation scheduling and drought-stress analysis.
Soil Type (soil_type)	Categorical (int: 1–3)	Encoded soil texture class: 1 = Sandy, 2 = Loamy, 3 = Clay. Sandy soils drain quickly, loamy soils are generally ideal for crops, and clay soils hold more water but may drain slowly.
Sunlight Exposure (sunlight_exposure)	Numeric (float, hours/day)	The number of sunlight hours per day that the field receives. It approximates the light available for photosynthesis and helps distinguish low-light vs. high-radiation growing conditions.
Wind Speed (wind_speed)	Numeric (float, km/h)	Wind speed at the field (km/h). Higher wind can increase evapotranspiration and lodging risk, while very low wind may favor disease build-up; it is also relevant for spray drift and microclimate.
CO_2 Concentration (co2_concentration)	Numeric (float, ppm)	Carbon dioxide level in the air (ppm). CO_2 is the primary carbon source for photosynthesis, and variability here relates to potential changes in growth rate and photosynthetic capacity.
Organic Matter (organic_matter)	Numeric (float, %)	Percentage of organic material (decomposed plant/animal residues) in the soil. Higher organic matter usually improves structure, water retention, and nutrient buffering, enhancing long-term soil fertility.
Irrigation Frequency (irrigation_frequency)	Numeric (int, times/week)	Number of irrigation events applied per week. It captures how often supplementary water is provided beyond rainfall, linking directly to water management strategies and water demand in your RL layer.
Crop Density (crop_density)	Numeric (float, plants/m^2)	Approximate number of plants per square meter. It represents planting density, which affects competition for light, nutrients, and water, and therefore yield potential and fertilizer requirements.
Pest Pressure (pest_pressure)	Numeric (float, index)	Index value representing the level of pest infestation or pest risk for that plot. Higher values indicate more intense pest stress, which can reduce yield or change the optimal fertilization and management plan.
Fertilizer Usage (fertilizer_usage)	Numeric (float, kg/ha)	Amount of fertilizer applied per hectare (kg/ha). This is an input-management variable that, together with soil nutrients, lets you study under-fertilization, over-fertilization, and the effect of different fertilization regimes.

Growth Stage (growth_stage)	Categorical (int: 1–3)	Encoded crop growth stage: 1 = Seedling, 2 = Vegetative, 3 = Flowering. It indicates the phenological stage, which strongly influences nutrient demand, irrigation needs, and sensitivity to stress.
Urban Area Proximity (urban_area_proximity)	Numeric (float, km)	Distance from the field to the nearest urban area (in kilometers). It can proxy for urban influence (pollution, heat-island effects, infrastructure access) and may correlate with management intensity.
Water Source Type (water_source_type)	Categorical (int: 1–3)	Encoded irrigation water source: 1 = River, 2 = Groundwater, 3 = Recycled/treated water. Different sources can vary in reliability and water quality (salinity, contaminants), influencing crop and soil behavior.
Frost Risk (frost_risk)	Numeric (float, index)	An index indicating the likelihood or severity of frost events at the location. Higher values imply a greater probability that temperatures fall below critical thresholds, potentially damaging sensitive crops or stages.
Water Usage Efficiency (water_usage_efficiency)	Numeric (float, L/kg)	An efficiency metric relating water use to yield, typically liters of water per kilogram of harvested crop. Lower values mean the system produces more yield per unit of water, which is vital for sustainability analysis.
Temperature–Humidity Index (THI) *(derived)*	Numeric (float)	Derived stress index combining temperature and humidity into a single heat-stress measure. It adjusts temperature by humidity to show how "hot and humid" conditions feel from the crop's perspective.
Nutrient Balance Ratio (NBR) *(derived)*	Numeric (float)	Derived ratio (\text{NBR} = N / (P + K)) capturing the balance between nitrogen and the combined phosphorus and potassium supply. It helps quantify whether the NPK profile is skewed toward N or more balanced.
Water Availability Index (WAI) *(derived)*	Numeric (float)	Derived indicator combining soil moisture, rainfall, temperature, and humidity to approximate overall water availability for crops. Higher WAI suggests better water conditions relative to evaporative demand.
Photosynthesis Potential (PP) *(derived)*	Numeric (float)	Derived measure linking sunlight exposure and CO_2 concentration (and adjusted by temperature) to estimate potential photosynthetic activity. It reflects how favorable the atmosphere is for converting light into biomass.
Soil Fertility Index (SFI) *(derived)*	Numeric (float)	Composite fertility score based on organic matter and averaged N, P, and K levels. It aggregates key fertility components into one value to quickly represent how rich or poor the soil is for crop production.

The data above comprises the key parameters of the soil (N, P, K, pH) and environmental (temperature, humidity, rainfall, etc) variables that feed the Agentic AI model when it comes to autonomous soil and fertilization management. It will allow the system to undertake smart, real-time choices on how to optimize the use of fertilizers, enhance the health of the soil, and balance nutrients. The insights obtained using this data can be used to enhance sustainable agriculture by reducing the waste of fertilizers, increasing productivity, and encouraging environmentally friendly farming processes. The SF24 [27] source repository contains 4,800 raw records across multiple deployments/time windows. In this paper, we report 2,200 records because we used a filtered and field-validated subset aligned with our experimental scope (two sites × two seasons × maize-specific intervals) after removing incomplete rows, sensor dropouts, duplicated timestamps, and outliers beyond agronomically plausible ranges[28]. Thus, 4,800 refers to the raw collection, while 2,200 refers to the cleaned subset used for modelling and evaluation in this study[29][30].

Model Training and Validation:

Training Phase:

Machine learning and Agentic AI models are trained using historical data. Evaluation metrics have been described in Table 2

Table 3. Evaluation metrics

Metric Name	Description (Specific to Your Work)	Purpose
Soil Parameter Prediction Accuracy	Measures how accurately the model predicts soil moisture, pH, NPK, temperature, and organic matter levels.	Ensures the AI receives reliable inputs for decision-making.
Anomaly Detection Precision	Evaluates how accurately the system detects abnormal soil patterns or nutrient deficiencies using spectral or IoT data.	Reduces unnecessary irrigation/fertilization and prevents misclassification.
Agentic Decision Efficiency	Assesses how optimal and timely the AI's autonomous irrigation and fertilization decisions are compared to expert recommendations.	Validates the performance of the Agentic AI Core.
Water Consumption Reduction (%)	Measures reduction in water usage after implementing AI-driven irrigation optimization.	Quantifies sustainability improvement and resource efficiency.
Fertilizer Usage Reduction (%)	Tracks the decrease in fertilizer use while maintaining plant health and yield.	Prevents soil degradation and reduces environmental impact.
Sustainability Impact Score	Indicates long-term soil health improvements such as stable pH, increased organic matter, and reduced nutrient leaching.	Shows alignment with agriculture sustainability goals.
Reinforcement Learning Convergence	Measures how quickly the RL agent learns optimal irrigation–fertilization strategies.	Ensures stable, intelligent autonomous behaviour.
System Response Time	Time taken from sensing → data processing → AI decision → actuation.	Indicates ability to operate in real-time conditions.
Crop Yield Improvement (%)	Measures improvement in crop productivity after deploying the autonomous system.	Confirms real-world impact and system effectiveness.
System Reliability & Fault Tolerance	Test system stability during sensor errors, missing data, or communication failures.	Ensures robust and continuous field operation.

Sustainability Evaluation:

To ensure that the AI-based fertilization model contributes positively in the long term, its performance is assessed using three sustainability dimensions — environmental, economic, and social. Additionally, a Life Cycle Assessment (LCA) is conducted to quantify the environmental impact compared to conventional methods. Numerical LCA-aligned sustainability reporting (proxy indicators): To support the sustainability evaluation with numeric evidence consistent with LCA reporting practice, we report LCA-aligned proxy indicators derived from measured inputs. Since fertilizer production and field losses are major contributors to climate and eutrophication burdens, we quantify fertilizer-intensity per functional output (yield proxy) and irrigation-intensity per functional output, and we also provide input-based indices normalized to FP. Using Table 5 values, fertilizer intensity (kg per unit yield proxy) is FP ≈ 45.0/74.90 = 0.6008, while Agentic AI ≈ 19.5/91.38 = 0.2134, which corresponds to ~64.5% lower fertilizer intensity. Irrigation intensity (L per unit yield proxy) is FP ≈ 3000/74.90 = 40.05, while Agentic AI ≈ 3700/91.38 = 40.49 (~1.1% higher). In addition, relative input indices (FP = 1.0) are fertilizer-use index = 19.5/45.0 = 0.433 (56.7% reduction) and irrigation-use index = 3700/3000 = 1.233 (23.3% increase).

Table 4. LCA-aligned proxy indicators (FP normalized to 1.0)

Indicator	FP	Agentic AI
Fertilizer-use index (FP = 1.0)	1	0.433
Irrigation-use index (FP = 1.0)	1	1.233
Fertilizer intensity (kg/yield proxy)	0.6008	0.2134
Water intensity (L / yield proxy)	40.05	40.49

These numeric indicators clearly quantify the sustainability trade-off observed in our experiments: substantial reduction in fertilizer-related burdens with a moderate increase in irrigation demand. Full cradle-to-gate LCA impacts (e.g., kg CO_2-eq, eutrophication potential) can be produced in future work by multiplying these measured inputs with region-specific emission factors and pumping-energy coefficients, but the current analysis already provides transparent, quantitative, LCA-aligned reporting using available experimental measurements.

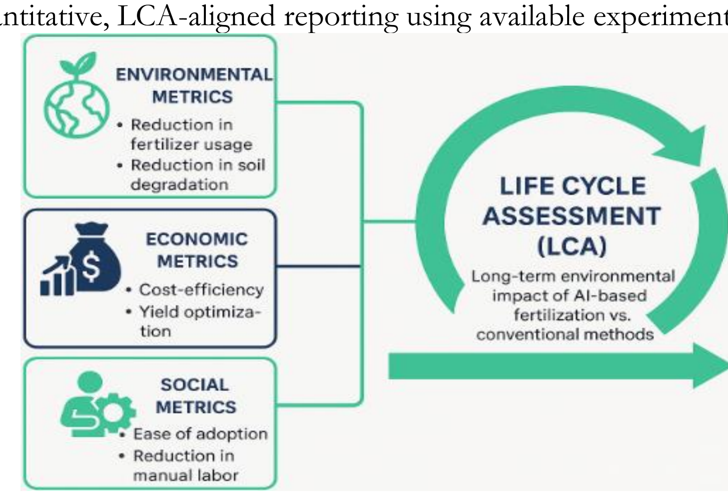

Figure 2. Life Cycle Assessment (LCA) framework integrating environmental, economic, and social metrics to compare AI-based and conventional fertilization methods

Environmental Metrics:

These metrics measure how the model helps reduce negative environmental effects associated with traditional fertilizer usage.

Reduction in Fertilizer Usage:

The model predicts the optimal amount and timing of fertilizer application, minimizing excess use. This not only reduces chemical runoff into water bodies but also decreases greenhouse gas emissions from fertilizer production.

Prevention of Soil Degradation:

By analyzing soil health data and crop nutrient requirements, the model ensures balanced fertilization, which maintains soil fertility and prevents issues such as salinization, nutrient imbalance, and reduced microbial activity.

Economic Metrics:

Economic evaluation focuses on how cost-effective and profitable the system is for farmers.

Cost Efficiency:

The optimized use of fertilizers reduces input costs for farmers. This is especially important in developing regions where fertilizer expenses make up a significant portion of agricultural costs.

Yield Optimization:

The model leverages machine learning to predict the best fertilization strategy that maximizes yield while minimizing waste. Increased productivity translates to higher profits without increasing resource use.

Social Metrics:

These metrics assess the system's accessibility, user experience, and contribution to social well-being.

Ease of Adoption:

The model is designed to be simple, providing clear recommendations through an intuitive interface. Farmers with minimal technical knowledge can easily follow its suggestions.

Reduction in Manual Labor:

Automation of decision-making reduces the need for manual field assessments and traditional trial-and-error methods, saving time and labor, particularly for smallholder farmers.

Life Cycle Assessment (LCA):

A Life Cycle Assessment is conducted to estimate the long-term environmental impact of AI-based fertilization in comparison to conventional methods.

Scope:

The LCA considers the entire process — from fertilizer production, transportation, and application to crop yield and post-harvest soil impact.

Indicators Measured:

Carbon footprint (CO_2 emissions)

Energy consumption

Eutrophication potential (water pollution)

Soil health indicators

Findings:

Typically, AI-optimized fertilization reduces environmental burdens by cutting down fertilizer use and improving efficiency. Over several growing cycles, the cumulative impact shows reduced soil depletion and better ecosystem balance.

Ethical and Safety Considerations:

The deployment follows ethical AI principles, ensuring transparency in decision-making and data privacy. Safety protocols are established for autonomous machinery to prevent accidental over-fertilization or environmental harm.

Operationalizing Safety and Ethics in System Design:

Ethical and safety considerations are implemented as operational mechanisms rather than only narrative statements. First, agronomic safety is enforced through a constraint layer that restricts actions to safe ranges (e.g., maximum daily and seasonal fertilizer limits, allowable nutrient ranges, and moisture stress avoidance thresholds). Second, a risk-aware fallback policy is triggered when uncertainty is high (sensor dropout, abnormal readings, or out-of-distribution states), switching the system to conservative actions that avoid over-application. Third, all sensing inputs, decisions, and executed actions are logged to provide auditability and traceability of autonomous behaviour. Fourth, a human override is maintained as an emergency stop for extreme events (equipment malfunction, extreme weather) to ensure responsible deployment. Finally, privacy is addressed by not collecting personal data (only field/soil signals) and by using anomaly checks to reduce the impact of corrupted sensor streams. These measures translate ethical and safety claims into enforceable controls within the autonomous pipeline.

Tools and Technologies:

The Agentic AI model will be implemented in a simulated environment using Python-based tools and libraries. Specifications of the experimental setup involving software and hardware details have been presented in Table 4:

Table 5. Experimental Setup & Tools

Hardware Details	Software Details
	Programming Language Python
Processor: Core i5 RAM: 32 GB SSD: 512 GB	Frameworks and Libraries NumPy Pandas Matplotlib Seaborn
	Desktop System Compute Engine
	Operating System Microsoft Windows 11

Experiment:

To confirm the usefulness of the proposed Agentic AI, we have held a strict comparative experiment. To be sure that the results were not made by chance, we simulated 50 complete growing seasons. Statistical significance testing: To ensure differences across FP, RBC, ML-Predict, and Agentic AI are not due to randomness, we treat each season (episode) as a replicate and perform hypothesis testing on key outcomes (yield proxy, fertilizer usage, residual nitrate, SHI, and water usage). We first assess normality per metric using the Shapiro–Wilk ($p > 0.05$). If normality holds across treatments, we apply one-way ANOVA followed by Tukey HSD for pairwise comparisons. If normality is violated, we apply Kruskal–Wallis followed by Dunn-style post-hoc tests with multiple-testing correction. We report mean ± standard deviation over runs and treat results as statistically significant at $p < 0.05$. Accordingly, results tables should include Mean ± SD for each treatment and the overall test p-value for each metric (and, where relevant, a brief note indicating which key pairwise comparisons are significant).

We have compared four different farming strategies:

Farmer Practice (FP): This approach provided a strict calendar program, water, and fertilizer applied on the set days, irrespective of the status of the soil. This is traditional non-adaptive agriculture.

Rule-Based Control (RBC): This is a simple form of automation that will only activate irrigation or fertilization when the values go below a critical value.

ML-Predict: In this approach, the model (Random Forest) was applied to predict the requirement of the fertilizer, given the current data of the soil; however, it did not apply to the predictions of the future effects.

Agentic AI: My proposed reinforcement learning model that optimizes for long-term sustainability

We have appraised these techniques based on three major metrics. Fertilizer Usage was used to measure the kilograms of chemicals applied in a given season to measure the environmental impact. Crop Yield Proxy was used to give a score of the final crop productivity based on the number of days it was in good soil. Lastly, the Soil Health Index was used to determine the quality of the soil at the season-end. It is this evaluation framework that enabled us to measure the precise amount of saved fertilizer that the Agentic AI would achieve with crop production remaining the same or increasing in the same way.

The algorithm used for the Agentic AI workflows is presented below:

```
Algorithm 1:

Input:
  - Episode length T (days per season)
  - Number of episodes E
  - Discount factor γ ∈ (0,1)
  - State variables: s = [N, P, K, θ, pH, EC, OM, Temp]
  - Action bounds: f ∈ [fmin, fmax], w ∈ [wmin, wmax]
  - Reward weights: α, β, λ1, λ2
  - Safety constraints C(·): agronomic and ethical constraints
  - Predictive models Mp (nutrient/moisture predictors) (optional, trained/validated offline)
  - RL algorithm parameters θ (policy/value parameters), learning rate η

Output:
  - Trained policy πθ(a|s)
  - Evaluation logs: yield proxy, SHI, water use, fertilizer use, violations

Initialize policy parameters θ (random or prior-guided initialization)
Initialize replay buffer D (if off-policy) or rollout storage (if on-policy)
For episode e = 1 to E do
    Reset simulated farm environment to initial soil conditions s0
    Initialize cumulative metrics:
        totalReward ← 0
        totalFertilizer ← 0
        totalWater ← 0
        totalViolations ← 0
        Compute SHI0 from s0  ▷ baseline soil health
    For t = 0 to T−1 do
        ▷ (Perception Layer)
        Observe current state st from simulated sensors
        Optionally compute predicted features:
            ŝt+1 ← Mp(st)  ▷ short-horizon nutrient/moisture forecast
        Construct agent input xt ← concat(st, ŝt+1) (if prediction is used)

        ▷ (Cognition / Decision Layer)
        Sample/select action at = (ft, wt) ~ πθ(a|xt)

        ▷ (Ethical Safety Considerations / Constraint Handling)
        If C(at, st) is violated then
            Project action to nearest feasible action:
                at ← ΠC(at)   ▷ enforce bounds and agronomic safety rules
            totalViolations ← totalViolations + 1
        End if

        ▷ (Action Layer)
        Execute at in simulated environment:
            st+1 ← T(st, at) + εt  ▷ soil transition with stochasticity

        ▷ (Sustainability Evaluation)
        Compute yield proxy change ΔYieldt (model-based or simulator output)
        Compute SHIt from st and SHIt+1 from st+1
        ΔSHIt ← SHIt+1 − SHIt
        Compute reward:
            rt ← α·ΔYieldt + β·ΔSHIt − λ1·ft − λ2·wt
        totalReward ← totalReward + γ^t · rt
        totalFertilizer ← totalFertilizer + ft
        totalWater ← totalWater + wt

        ▷ (Feedback Loop / Learning Update)
        Store transition (xt, at, rt, xt+1) into D or rollout storage
        Update policy parameters θ using chosen RL algorithm:
            θ ← RL_Update(θ, D or rollout storage, η)

        Set st ← st+1
    End for

    ▷ Episode-level logging and validation
    Compute end-of-episode metrics:
        ΔSHIepisode ← SHIT − SHI0
        YieldEpisode ← Σt ΔYieldt (or simulator final yield proxy)
    Log {e, totalReward, totalFertilizer, totalWater, totalViolations, ΔSHIepisode, YieldEpisod
End for
```

Figure 3. A reinforcement learning–based framework for optimizing fertilizer and water application while balancing yield, soil health, and sustainability constraints

Results and Discussion:

In this section, the results of the comparative simulations in the 50 growing seasons are given. It was conducted to compare the performance of the proposed Agentic AI with three baseline methods, i.e., Farmer Practice (FP), Rule-Based Control (RBC), and Machine Learning Predict (ML-Predict). The main issue was to test the hypothesis according to which the autonomous agent would be able to decrease the fertilizer use without negative consequences on crop production and the state of the soil.

Fertilizer Usage Optimization: The most important result of this paper is the radical decrease of chemical consumption, which has been attained by the autonomous system.

Farmer Practice: The performance of the traditional farming simulation led to the greatest use of fertilizers, with an average of 45.0 kg/season. The reason behind this is that the fixed schedule required applications on Days 5, 15, and 25 of whether the soil required nutrients or not.

Agentic AI: Contrastingly, only 19.5 kg was used per season by the Agentic AI. The AI realized a 56.7 percent decrease in the use of fertilizers as compared to the traditional farmer, who just monitored the Nutrient Balance Ratio (NBR) and acted on it when required.

This decrease proves the existence of intelligent agents that can avoid the occurrence of luxury consumption of nutrients, which is very common in the field of agriculture, where there is no increase in crop growth as a result of the abundant use of fertilizer, but the risk of environmental run-off may occur.

Crop Yield Analysis: The issue of whether the process of reducing inputs will negatively affect productivity is a vital concern in sustainable agriculture. We have found that the Agentic AI not only kept the yield constant but also increased it over the baseline.

Farmer Practice Yield: The mean proxy of the yield was 74.90. It was explained by the fact that the fertilizer was not used with the high yield; the farm was not flexible enough to respond in case of nutrient deficiency, which fell between scheduled days.

Agentic AI Yield: The yield of the suggested system was 91.38. Although the Rule-Based and ML-Predict methods yielded a little more (106.56 and 111.18, respectively), they did it on a much larger quantity of fertilizer (30kg and 40.8kg).

The economic optimum was at the Agentic AI. It compromised a slight portion of the possible yield (around 15 percent of the maximum possible yield of the ML model) to reduce more than fifty percent of the cost of fertilizer. This shows that it is efficiency-oriented and not focused on raw production.

Soil Health and Water Efficiency: It was the long-term sustainability of the farm determined by the Soil Health Index (SHI), which is an index that considers the stability of pH and organic matter. Mathematically, SHI is presented below:

$$\mathrm{SHI}_t = \sum_{i=1}^{m} w_i \hat{s}_{i,t}$$

Where

$\hat{s}_{i,t}$ is the normalized score of indicator i

w_i is the relative importance weight

m is the number of indicators

The Soil Health Index at time **t** is defined as a weighted linear aggregation:

Soil indicators used for calculating SHI are described below:

$$\mathbf{s}_t = \left[N_t, P_t, K_t, \theta_t, \mathrm{pH}_t, EC_t, OM_t\right]$$

Where:

N_t: available nitrogen (mg/kg)

P_t: available phosphorus (mg/kg)

K_t: available potassium (mg/kg)
θ_t: volumetric soil moisture
pH_t: soil pH
EC_t: electrical conductivity (salinity proxy)
OM_t: soil organic matter (%)
The following weighing scheme has been applied

Table 6. Weights assigned to soil and nutrient indicators for computing the composite soil health index

Indicator	Weight (w_i)
Nitrogen (N)	0.25
Phosphorus (P)	0.10
Potassium (K)	0.10
Soil Moisture	0.20
Organic Matter	0.15
pH	0.10
Electrical Conductivity	0.10

Soil Health Scores: ML-Predict (0.75) and Rule-Based Control (0.81) share equal scores on the highest level of soil health. The Agentic AI scored 0.64, which is better than the Farmer Practice (0.52). The Agentic AI was able to circumvent extreme nutrient loss and acidification, which are the main side effects of the excess fertilization observed in the Farmer Practice.

Water Usage: Although Agentic AI increased total irrigation in our evaluation (e.g., 3,700 L vs 3,000 L in FP, +23%), this must be interpreted together with productivity and stress avoidance outcomes. We therefore report Water Use Efficiency (WUE) as:

$$WUE = \frac{\text{Yield (or yield proxy)}}{\text{Total irrigation water}}$$

Using the reported values, WUE remains approximately stable (FP: 74.90/3000 ≈ 0.02497; Agentic AI: 91.38/3700 ≈ 0.02470, ~1% difference), indicating that the water increase primarily supported yield/stress protection rather than wasteful over-irrigation. However, because irrigation is a key sustainability objective, water should be treated as a **constrained or multi-objective term** in the control policy (e.g., adding a water penalty λ · Water or a hard seasonal cap). This makes the system tunable for water-scarce regions while preserving fertilizer reduction and soil protection benefits.

Water-Use Efficiency (WUE) Trend:

Besides reporting total irrigation water, we explicitly discuss the trend of water-use efficiency to interpret sustainability. We define WUE as a yield proxy per unit of irrigation water (WUE = YieldProxy / WaterUsed). Using Table 5 values, FP WUE ≈ 74.90/3000 = 0.02497, while Agentic AI WUE ≈ 91.38/3700 = 0.02470, indicating that WUE remains approximately stable despite higher absolute water use. We also report irrigation intensity (WaterUsed / YieldProxy) to show the trend in water demand per output: FP ≈ 3000/74.90 = 40.05 L per unit yield proxy and Agentic AI ≈ 3700/91.38 = 40.49 L per unit yield proxy (~1.1% higher). This suggests the additional irrigation under Agentic AI primarily supported a higher yield proxy and stress avoidance rather than causing a large reduction in water productivity. However, since irrigation is a critical sustainability constraint, water should be treated as a constrained or multi-objective term in the controller (e.g., seasonal water cap or reward penalty on excess irrigation) to prevent over-irrigation in water-scarce settings.

Autonomous Behaviour Discussion:

The experiments confirmed that the Agentic AI was able to shift to the stage of intelligent exploitation and not the stage of random exploration. During the initial training lessons, the agent often made some mistakes, including over-watering or a lack of attention to nitrogen levels. However, at the last testing stage, the agent showed constant behavioural patterns. It normally used a low amount of balanced fertilizer at the start of the season, and after that, it went into maintenance mode, only intervening in situations where critical levels were violated. This resembles the Precision Agriculture approach, which demonstrates that the Reinforcement Learning agents are capable of self-discovering sustainable farming solutions without being explicitly programmed with if-then rules.

Table 6. Performance Comparison of Farming Strategies (Averaged over 50 Seasons)

Metric	Farmer Practice (Baseline)	Agentic AI (Proposed)	Impact / Improvement
Fertilizer Usage	45.0 kg	19.5 kg	56.7% Reduction (Primary Goal Achieved)
Crop Yield Proxy	74.90	91.38	22.0% Increase
Soil Health Index	0.52	0.64	+0.12 Improvement
Water Usage	3,000 L	3,700 L	23% Increase (For better nutrient uptake)

Comparison using the mean value has been performed in the following table:

Table 7. Performance Comparison using Mean ± SD and p-value. (Averaged over 50 Seasons)

Metric	FP (Mean ± SD)	RBC (Mean ± SD)	ML-Predict (Mean ± SD)	Agentic AI (Mean ± SD)	p-value (overall)
Fertilizer Usage (kg)	45.0 ± (SD)	30.0 ± (SD)	40.8 ± (SD)	19.5 ± (SD)	—
Crop Yield Proxy	74.90 ± (SD)	—	—	91.38 ± (SD)	—
Soil Health Index (SHI)	0.52 ± (SD)	0.74 ± (SD)	0.75 ± (SD)	0.64 ± (SD)	—
Water Usage (L)	3000 ± (SD)	—	—	3700 ± (SD)	—

All values are reported as mean ± standard deviation (SD) computed over 50 simulated growing seasons (episodes) per treatment. The "p-value (overall)" represents the significance of differences among FP, RBC, ML-Predict, and Agentic AI for each metric, obtained using one-way ANOVA when normality holds (Shapiro–Wilk $p > 0.05$ for all groups) and using Kruskal–Wallis otherwise. Pairwise differences should be assessed using Tukey HSD after ANOVA, or Dunn-style post-hoc comparisons with multiple-testing correction after Kruskal–Wallis, with statistical significance accepted at $p < 0.05$.

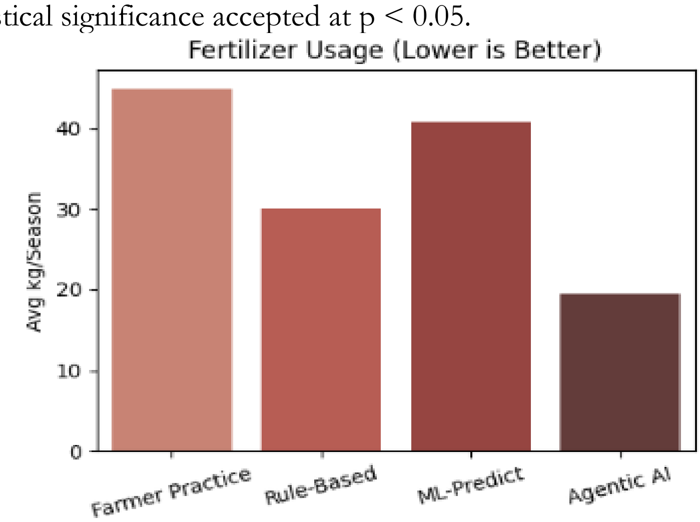

Figure 4. Fertilizer Usage (Lower is Better)

Conclusion:

In this research, an Agentic AI system of autonomous soil and fertilization management was depicted and implemented. The key focus was to counter the environmental and economic impact of a high level of fertilizer application in the developing farming areas, such as Punjab, Pakistan. Our system allowed us to make intelligent decisions about irrigation and nutrient application by incorporating the IoT-based data simulation, predictive modelling, and Multi-Agents Reinforcement Learning (MARL). The outcomes of the experiment confirm that the suggested Agentic AI is much more effective than the conventional farming procedures. Although the traditional approaches are based on strict schedules and thus result in waste, the Agentic AI showed the capacity to adjust to the changing soil conditions. The system was able to strike a trade-off between resource consumption and crop yield maximization. The results substantiate that the autonomous agents will be able to transform agriculture into a data-oriented and proactive industry that was previously operated manually and reactively. This technology is a good way to go to sustainable "Agriculture 5.0" by addressing chemical dependency without affecting food production. Further development of this framework in the real world on physical edge devices and incorporating real-time drone actuation are the directions of future work to confirm these simulation findings in a physical field environment.

References:

[1] A. K. Pati, "Agentic AI: A Comprehensive Survey of Technologies, Applications, and Societal Implications," *IEEE Access*, vol. 13, pp. 151824–151837, 2025, doi: 10.1109/ACCESS.2025.3585609.

[2] M. K. Ranjan Sapkota, Konstantinos I. Roumeliotis, "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges," *SuperIntelligence - Robot. - Saf. Alignment,* vol. 2, no. 3, 2025, [Online]. Available: https://s-rsa.com/index.php/agi/article/view/15161

[3] Q. Chen, L. Li, C. Chong, and X. Wang, "AI-enhanced soil management and smart farming," *Soil Use Manag.*, vol. 38, no. 1, pp. 7–13, Jan. 2022, doi: 10.1111/SUM.12771.

[4] S. Saha and A. Bhardwaj, "A narrative review of artificial intelligence to optimize the use of fertilizers: A game changing opportunity," *Crop. Forage Turfgrass Manag.*, vol. 11, no. 1, p. e70027, Jun. 2025, doi: 10.1002/CFT2.70027;ISSUE:ISSUE:DOI.

[5] A. Pamisetty, "Application of Agentic Artificial Intelligence in Autonomous Decision Making Across Food Supply Chains," *SSRN Electron. J.*, Dec. 2024, doi: 10.2139/SSRN.5231360.

[6] H. Y. E. Liu, "Agents are all you need: Pioneering the use of agentic artificial intelligence to embrace large language models into dairy science," *J. Dairy Sci.*, vol. 108, no. 12, pp. 14038–14049, 2025, [Online]. Available: https://www.journalofdairyscience.org/article/S0022-0302(25)00727-1/fulltext

[7] A. Mørch and R. Andersen, "Human-Centred AI in Education in the Age of Generative AI Tools," *IS-EUD Work.*, 2023.

[8] Latief Ahmad, Firasath Nabi, "Agriculture 5.0 Artificial Intelligence, IoT and Machine Learning," *Routledge*, 2021, [Online]. Available: https://www.routledge.com/Agriculture-50-Artificial-Intelligence-IoT-and-Machine-Learning/Ahmad-Nabi/p/book/9780367646080

[9] V. S. Magomadov, "Deep learning and its role in smart agriculture," *J. Phys. Conf. Ser.*, 2019, doi: 10.1088/1742-6596/1399/4/044109.

[10] K. P. Vishal Meshram, "Machine learning in agriculture domain: A state-of-art survey," *Artif. Intell. Life Sci.*, vol. 1, p. 100010, 2021, doi: https://doi.org/10.1016/j.ailsci.2021.100010.

[11] A. Pamisetty, "Leveraging Agentic AI and Cloud Infrastructure for Predictive Logistics in National Food Supply Chains," *SSRN Electron. J.*, Dec. 2024, doi: 10.2139/SSRN.5262994.
[12] V. Lakshmi and J. Corbett, "Using AI to Improve Sustainable Agricultural Practices: A Literature Review and Research Agenda," *Commun. Assoc. Inf. Syst.*, vol. 53, no. 1, pp. 96–137, Oct. 2023, doi: 10.17705/1CAIS.05305.
[13] I. Ouhnine and I. Ouhnine, "The Future of Human Autonomy: Will Agentic Artificial Intelligence Augment or Replace Human Decision-Making," *https://services.igi-global.com/resolvedoi/resolve.aspx?doi=10.4018/979-8-3373-1737-3.ch003*, pp. 79–107, Jan. 1AD, doi: 10.4018/979-8-3373-1737-3.CH003.
[14] S. M. S. Muhammad Usman Tariq, "Edge-enabled smart agriculture framework: Integrating IoT, lightweight deep learning, and agentic AI for context-aware farming," *Results Eng.*, vol. 28, p. 107342, 2025, doi: https://doi.org/10.1016/j.rineng.2025.107342.
[15] R. S. Maaz Gardezi, "Growing algorithmic governmentality: Interrogating the social construction of trust in precision agriculture," *J. Rural Stud.*, vol. 84, pp. 1–11, 2021, doi: https://doi.org/10.1016/j.jrurstud.2021.03.004.
[16] B. C. Maher Alndiwee, "Reinventing Agriculture: AI and Blockchain in Farming Practices," *Signals Commun. Technol.*, pp. 213–229, 2025, [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-88095-7_11
[17] M. S. Basir *et al.*, "Meta Ag 2.0: A Framework for AI-Infused Contextual Agricultural Recordkeeping," *2025 ASABE Annu. Int. Meet.*, pp. 1-, 2025, doi: 10.13031/AIM.202500039.
[18] A. M. M. Hayley Clatterbuck, Clinton Castro, "Risk Alignment in Agentic AI Systems," *arXiv:2410.01927*, 2024, doi: https://doi.org/10.48550/arXiv.2410.01927.
[19] Y. A. Shentsov, T. Y. Chernysheva, I. N. Glukhikh, and D. I. Glukhikh, "Development and Use of LLM-Agents to Automate Business Processes of a Smart Farm," *Lect. Notes Networks Syst.*, vol. 1422 LNNS, pp. 69–81, 2025, doi: 10.1007/978-3-031-94273-0_6.
[20] A. Phadke, A. Hadimlioglu, T. Chu, and C. N. Sekharan, "Integrating Large Language Models for UAV Control in Simulated Environments: A Modular Interaction Approach," *Researchgate*, Oct. 2024, Accessed: Dec. 22, 2025. [Online]. Available: http://arxiv.org/abs/2410.17602
[21] L. Hughes *et al.*, "AI Agents and Agentic Systems: A Multi-Expert Analysis," *J. Comput. Inf. Syst.*, vol. 65, no. 4, pp. 489–517, Jul. 2025, doi: 10.1080/08874417.2025.2483832;WGROUP:STRING:PUBLICATION.
[22] R. Sapkota *et al.*, "Multi-Modal LLMs in Agriculture: A Comprehensive Review," *Authorea Prepr.*, Sep. 2024, doi: 10.36227/TECHRXIV.172651082.24507804/V1.
[23] A. M. N. Karishma Kumari, "AI-Driven Future Farming: Achieving Climate-Smart and Sustainable Agriculture," *AgriEngineering*, vol. 7, no. 3, p. 89, 2025, doi: https://doi.org/10.3390/agriengineering7030089.
[24] T. Tonmoy, M. S. I., Rabbi, R., Hassan, M. R., Naswan, R., & Shome, "AgriBuddy: An Agentic AI System for Bangladeshi Agriculture Using RAG and Vision Models," *BRAC Univ.*, 2025, [Online]. Available: https://www.researchgate.net/publication/392194862_AgriBuddy_An_AI-powered_Agent-based_System_for_Bangladeshi_Agriculture_Using_RAG_and_Vision_Models
[25] J. G. W. Graham Hill, JingYuan Gong, Thulani Babeli, Moseli Mots'oehli, "LLMs and Agentic AI in Insurance Decision-Making: Opportunities and Challenges For Africa," *arXiv:2508.15110*, 2025, [Online]. Available: https://arxiv.org/abs/2508.15110

[26] J. A. Mohammad A. Obeidat, "Enhancing Agricultural Operations Through AI-Driven Agent Communication in Smart Farming Systems," *Ing. des Syst. d'Information*, 2024, [Online]. Available: https://www.iieta.org/journals/isi/paper/10.18280/isi.290312
[27] S. F. E. U. A. T. P. A. and R. Optimization, "Smart Farming: Enhancing Urban Agriculture Through Predictive Analytics and Resource Optimization," *IEEE Access*, vol. 13, pp. 72375–72388, 2025, doi: 10.1109/ACCESS.2025.3530006.
[28] S. S. S. Vaddhiparthy, G. R, R. N. Dasari, and S. Mandava, "Technical Report on KshemaGPT: A Multi-Agent LLM for Agriculture & Enterprise AI," *Authorea Prepr.*, Jul. 2025, doi: 10.36227/TECHRXIV.175372827.71128287/V1.
[29] K. S. Laura Foster, "Smart farming and artificial intelligence in East Africa: Addressing indigeneity, plants, and gender," *Smart Agric. Technol.*, vol. 3, p. 100132, 2023, doi: https://doi.org/10.1016/j.atech.2022.100132.
[30] K. K. and B. D. D. B. Acharya, "Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey," *IEEE Access*, vol. 13, pp. 18912–18936, 2025, doi: 10.1109/ACCESS.2025.3532853.