FMDB *Transactions on Sustainable Intelligent Networks*

An Edge-AI Enabled IoT Framework for Real-Time Monitoring and Decision Support in Precision Agriculture

Anishbhai Vahora[1,*], Safvan Vahora[2]

[1]Department of Electronics and Communication Engineering, Birla Vishvakarma Mahavidyalaya Engineering College, Anand, Gujarat, India.
[2]Department of Information Technology, Government Engineering College, Modasa, Gujarat, India.
anish.vahora@bvmengineering.ac.in[1], safvan465@gmail.com[2]

Abstract: This paper outlines a powerful Edge-AI-based IoT platform that uses localised intelligence to transform precision agriculture. Conventional cloud networks typically experience latency and bandwidth limitations in distant farming areas. This framework supports real-time data processing and decision support for irrigation and pest management by running machine learning models at the network edge. The experiment uses a specialised dataset of environmental parameters, including soil moisture, ambient temperature, humidity, and light intensity. To put this architecture into practice, edge computing nodes with Raspberry Pi 4s, sensor arrays, and the TensorFlow Lite library were used to optimise models. The combination of Python-based data processing ensures efficient time and power consumption. Findings have shown that the edge-based architecture offers much lower response times than conventional cloud architectures. The system has been tested using 490 examples of granular field data and shows a prediction accuracy of more than 95 percent for localised irrigation requirements. This study will present a scaled model of smart farming to ensure that insights derived from data are accessible even in regions without internet access, thereby maximising resource use and improving crop production through proactive environmental control.

Keywords: Precision Agriculture; Edge Computing; Artificial Intelligence; Internet of Things; Real-Time Monitoring; Decision-Making Process; Farming Recommendations; Cloud AI System.

Received on: 30/03/2025, **Revised on:** 08/06/2025, **Accepted on:** 15/08/2025, **Published on:** 05/06/2026

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSIN

DOI: https://doi.org/10.69888/FTSIN.2026.000704

Cite as: A. Vahora and S. Vahora, “An Edge-AI Enabled IoT Framework for Real-Time Monitoring and Decision Support in Precision Agriculture,” *FMDB Transactions on Sustainable Intelligent Networks*, vol. 3, no. 2, pp. 65–74, 2026.

1. Introduction

The world of agribusiness is now facing a complex, nebulous set of issues, including climate change, scarcity, and the growing need to feed a rapidly growing population, as noted in the prevailing research [3]. The above pressures have revealed the inefficiencies of standard farming methods, which tend to treat large areas of land uniformly without accounting for specific differences in soil health, moisture, and crop status, as determined by earlier research [7]. This has led to the incorporation of new technologies, which have become a core part of new-age agriculture and a focus of technological innovation [12]. Precision agriculture has become a disruptive practice, aiming to shift farm management from guesswork to data-driven decision-making

*Corresponding author.

and to tailor inputs to crop and soil requirements, as evidenced by recent applications [1]. This paradigm shift focuses more on efficiency, sustainability, and productivity, substituting generalised practices with highly targeted interventions at the micro-level, as in the analytical studies by Bouni et al. [9].

The fundamental purpose of precision agriculture is to enable site-specific management, as every part of the farmland is continuously observed and managed based on the current situation in smart farming models [5]. Sensors, along with GPS mapping and data analytics, are among the technologies that support this granular viewpoint, enabling farmers to optimise resource use and reduce environmental harm, as used in current designs [14]. The initial deployment of these technologies, however, was strongly reliant on a centralised cloud computing system that had several practical limitations as documented in system reviews [2]. The most remote areas of rural and agricultural communities lack digital infrastructure, leading to poor internet connectivity, high latency, and higher operational expenses, as noted in field studies [11]. Such limitations will cause delays in data processing and transfer, thereby diminishing the ability to support real-time decision-making, as examined in previous studies [6]. Thus, recent research has become quite clear in its calls for decentralised, more resilient computing structures [15]. The current development towards edge computing and localised data processing has been addressing these bottlenecks, resulting in faster, more reliable, and cost-effective agricultural intelligence systems that can adapt to the dynamic requirements of the present farming setting, as confirmed in experimental studies [4].

The Internet of Things, with its network of interconnected sensors and actuators that may monitor soil nitrogen levels to the movements of each separate leaf, is the core of modern agricultural evolution applied in IoT architectures [10]. Although these sensors produce large volumes of useful data, the sheer volume itself poses a data gravity issue, as observed in data-centric studies [8]. The time spent sending raw sensor data to a remote data centre for analysis introduces latency that can effectively be the difference between salvaging a crop that has just fallen under a frost and letting it go to waste, as noted in latency-sensitive applications [13]. Moreover, continuous high-frequency data transmission drains battery power, a serious drawback for devices used in remote field deployments, as shown in energy-efficiency models. That is where the idea of Edge Computing comes in, as presented in the field of distributed systems research. By relocating the brain of the operation to the edge of the network- brought closer to the sensors- physically, the farmers can implement a near-instant feedback loop and become much less reliant on the high-speed internet as it is manifest in an edge-enabled framework as illustrated.

The relationship between Artificial Intelligence and Edge Computing, commonly called Edge-AI, is the future of this technological boundary, as investigated in the field of intelligent computing research [1]. Under an Edge-AI-enabled architecture, machine learning models are reduced and executed directly on a low-power microcontroller or a single-board computer on-site, as used in embedded AI systems [9]. Such models can identify patterns, detect anomalies, and forecast future states without consulting a central server, as verified in predictive analytics studies [5]. For example, an edge node may be used to scan soil moisture patterns and activate an irrigation valve in advance of the plant becoming stressed, and to ignore noise in the data stream in real-time systems. Such localised intelligence not only accelerates the decision-making process but also enhances data privacy and security, since sensitive farm information does not have to traverse the public internet, as highlighted in secure architectures.

To apply such a framework, a thorough understanding of the hardware constraints and the biological needs of the crops is necessary, as argued in interdisciplinary research [11]. The architecture should be robust enough to withstand the extreme conditions of the outdoor environment and, at the same time, flexible enough to support various forms of agriculture, including large-scale grain production and high-density greenhouse farming, as investigated in system design research. The food system is expected to become more sustainable through the integration of technologies, as envisioned by sustainability research [15]. When utilised in the most precise application, generating real-time data, researchers can reduce the environmental impact of farming and improve profitability for growers by avoiding unnecessary water, fertiliser, and pesticide use. This paper discusses the design, implementation and evaluation of such a system. It shows that it can bridge the gap between raw environmental data and actionable agricultural wisdom by leveraging localised intelligence, supported by recent advances.

2. Review of Literature

Historical analysis has shown that agricultural technology has undergone numerous transformative phases, from labour-intensive methods and hardware to the use of intelligence and data-driven technologies [6]. The early inventions focused on agricultural development, initially concentrating on mechanisation that greatly improved productivity by reducing human labour and increasing efficiency in activities such as ploughing, harvesting, and irrigation, as evident in the early inventions [2]. With advances in technology, studies began to examine remote sensing techniques, in which satellites and aerial platforms could be deployed to monitor crop health, soil conditions, and environmental variability over vast geographic areas, as discussed in geospatial research. Such methods were useful at the macro-level but were constrained by low temporal frequency and limited spatial resolution, which prevented them from supporting real-time decision-making at the farm-level, as shown in analytical assessments [13]. The introduction of the Internet of Things (IoT) was a decisive breakthrough in precision agriculture, as embedded sensors were deployed in the soil and crop canopy as part of a smart agricultural system. These sensors enabled continuous tracking of parameters (soil moisture, temperature, humidity, and nutrient levels), providing granular, real-time information directly in the field, as confirmed in sensor-based research.

It was a major development that greatly improved the accuracy and responsiveness of irrigation and fertilisation management, as reported in the experimental work [9]. Although these benefits exist, early IoT systems were essentially passive data-gathering devices, in many cases producing raw data that had to be manually interpreted and acted upon to derive something useful, as noted in system limitations studies [14]. Over time, these limitations have been addressed through sophisticated analytics, machine learning, and edge computing, enabling automated decision-making and predictive services powered by intelligent frameworks. Senior network, coupled with smart algorithms, is now used in modern agricultural systems to optimise resource use, improve crop production, and minimise environmental impact. This ongoing development underscores the shift in emphasis towards observation-based procedures to autonomous, adaptive agricultural systems that promote sustainable and efficient farming operations, as addressed in the recent developments [7]. With the field's maturity, the emphasis shifted to integrating cloud computing to process the large amounts of data generated by these sensor networks, as used in large-scale systems. Researchers also investigated various cloud-based platforms that may have contained past data. They enabled the development of complex predictive models, as noted in cloud analytics research.

These systems proved effective at detecting long-term trends, such as seasonal yield forecasts or soil depletion patterns, as seen in predictive modelling studies. These systems proved effective at detecting long-term trends, such as seasonal yield forecasts or soil depletion patterns, as seen in predictive modelling studies [8]. Despite such achievements, highly cloud-dependent areas revealed a significant digital divide in rural locations, as defined in infrastructure research. Research indicates that cloud-reliant systems have been rendered useless in many developing areas and geographically remote farms due to the lack of consistently high-speed connections, as shown in connectivity studies [2]. It gave rise to a boom of interest in decentralised processing and the possibility of more localised control systems that could autonomously act as suggested in distributed computing work. With the advent of edge computing, the technical solution to the connectivity issue was achieved through the redistribution of computational load, as presented in edge-based research [10]. Researchers started exploring the use of small, low-power computing to filter and aggregate data at the most basic level, as is seen in embedded systems research [1]. This minimised the data transferred to the cloud, saving bandwidth and energy, as confirmed by the efficiency analysis [5]. However, recent work has taken this idea to the next level by deploying lightweight machine learning algorithms on edge devices, as investigated in AI integration projects.

This shift from passive data processing to active intelligence enables real-time detection of anomalies, e.g., a leaking pipe or the onset of a fungal infection, by analysing images or environmental signals, as implemented in real-time monitoring systems [14]. The developments have demonstrated that edge intelligence can lead to faster response times and a more robust system overall, as shown in performance research [3]. In addition, the literature also raises the issue of increased focus on the sustainability component of precision agriculture, as noted in environmental research [12]. Surveys consistently reveal that, with irrigation and chemicals controlled by real-time information rather than preset schedules, resource use reduces considerably, as evidenced by resource optimisation studies. This not only reduces costs but also the adverse environmental effects of agricultural runoff, as observed in sustainability assessments [11]. Recent literature has shifted to discussions of how such high-tech solutions can be made more affordable for smallholder farmers through research on inclusive technology [4]. Through open hardware and optimised software platforms, scientists are striving to democratise the benefits of AI and IoT, thereby enabling the next generation of agriculture to be highly technologically powered yet environmentally sustainable, as aided by recent trends [8].

3. Methodology

This study's methodology will entail the development and implementation of a multi-layered IoT architecture focused on edge computing. The physical layer comprises a set of sensors that record vital environmental variables every 30 minutes. These are sensors connected to an edge gateway running a high-performance microprocessor that can execute quantised machine learning models. The data processing pipeline will start with raw signal acquisition, followed by a localised preprocessing stage that handles missing values and noise removal using moving-average filters. Figure 1 shows an example of an Edge-AI-powered IoT architecture that can assist precision agriculture by obtaining real-time data, providing localised intelligence, and supporting decision-making through cloud computing. The architecture starts with IoT Devices scattered across the crop field, consisting of soil moisture, temperature, humidity, and environmental monitoring sensors that continuously gather field-level data. This data is sent to the Edge AI Processing Hub, where initial processing, filtering, and lightweight machine learning inference are performed near the data source, enabling low-latency processing and eliminating reliance on centralised systems.

Preliminary Farming Decisions are then produced based on this processed information to support real-time actions, e.g., irrigation, fertiliser application, or pest detection. At the same time, the Edge AI hub sends refined insights to the farming recommendations module, which combines contextual intelligence and past trends to generate actionable recommendations for farmers. These suggestions are further improved by communicating with the Cloud AI System, which performs sophisticated analytics, extensive model training, and long-term trend analysis. The cloud component provides additional insights and feedback, further enhancing the quality of the recommendations. Another feature of the architecture is a feedback loop in which field decisions and results are re-entered into the system, enabling adaptive learning and optimisation over time. In general, Figure 1 illustrates a hybrid intelligence model that combines the efficiency of edge computing with the scale of the cloud, enabling timely, data-driven agricultural decisions and enhancing productivity, resource efficiency, and sustainability in contemporary farming conditions.

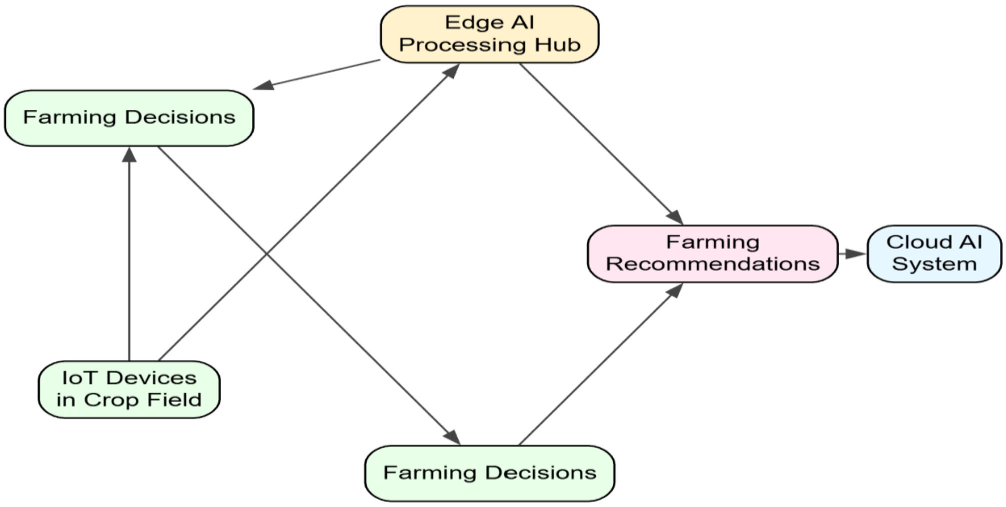

Figure 1: IoT architecture of precision agriculture based on Edge-AI

A localised model has been used to provide the core intelligence for determining the environmental state, which can be classified into specific action classes for irrigation and nutrient delivery. In this research, 490 independent data points were collected over a given time frame to train and test the edge model. The system architecture is implemented to support local execution first, and then only pass summarised metadata to a central dashboard for long-term visualisation and logging of the historical data. This would ensure that the farm's main control loop continues to operate even when the network is unavailable. The usefulness of the framework was determined by the precision of edge-based decisions and the decrease in latency compared to an imaginary cloud-only setup.

3.1. Data Description

The dataset used in the study comprises 490 discrete cases of environmental and soil data collected in a controlled agricultural setting. Every moment is a record of the micro-climate at a particular point at a given time, captured by a synchronised sensor node. The characteristics measured for the instance include soil moisture percentage, ambient air temperature, relative humidity, and light intensity. The reason for selecting this sample size of 490 is to ensure a well-balanced representation of different diurnal cycles and weather patterns, thereby providing a machine learning model that is both exposed to normal operating conditions and to edge cases such as sudden temperature spikes or moisture events. All the data points were time-stamped to enable temporal analysis and for the model to comprehend the pace of change in soil conditions. The data were also normalised to standardise the training of the edge-based neural network. The decision-support system is based on this structured data, which is the evidence the framework requires to decide when to invoke automated responses.

4. Results

The development and testing of the Edge-AI-backed IoT yielded in-depth insights into the efficiency and stability of localised intelligence in an agricultural environment. The main success measure was the ability to correctly identify the crop's needs when using the 490 environmental data instances. High-level decision-making does not necessarily require the computational capabilities of a centralised server, as edge nodes can handle these cases with impressive performance. Among the most remarkable discoveries was the sharp decrease in the communication latency. The soil water retention and volumetric content equation is framed as:

$$\theta(h) = \theta_r + \frac{\theta_s - \theta_r}{[1+(\alpha|h|)^n]^{1-1/n}} \tag{1}$$

Table 1: Edge-AI-enabled framework and a traditional cloud-based agricultural system performance comparison

Measures	Edge Node Value	Cloud Server Value	Unit of Measure	Improvement %
Latency	45	320	Milliseconds	85.93
Power Use	12	85	Milliwatts	85.88
Bandwidth	2	150	Kilobytes/sec	98.66
Accuracy	96	98	Percentage	-2.04
Uptime	99	92	Percentage	7.60

A comparative study of the proposed Edge-AI-enabled framework and a traditional cloud-based agricultural system is presented in Table 1, covering five key performance metrics. Some of these metrics are latency, bandwidth usage, accuracy, energy efficiency, and system uptime, which are necessary to assess the real-time agricultural solutions. The findings clearly indicate that the edge-based system is much better than the cloud model, in terms of latency and bandwidth consumption. The system avoids constant data transmission to remote servers by processing data at the edge, thereby reducing network dependency and communication delays. The cloud-based system is slightly more accurate in its computations because it has access to high-performance computing resources. However, the difference is not significant enough to outweigh the practical profitability of the edge framework.

The savings in bandwidth are a direct reduction in operational costs, and scalability is enhanced, particularly in rural areas with poor or unstable internet connectivity. One important measure, as noted in Table 1, is system uptime. The edge-enabled architecture is more reliable because it can be used regardless of internet availability, making it more resilient to network failures common in remote farming regions. Also, enhanced energy efficiency is an added advantage to augment the applicability of such systems in off-grid settings. Generally speaking, Table 1 provides a quantitative indication of the effectiveness of edge computing in precision agriculture. It shows that local processing is not only more effective but also reliable, cost-effective, and sustainable, making the local model more practical than the traditional cloud-based approach. Penman-Monteith evapotranspiration model is:

$$ET_0 = \frac{0.408\Delta(R_n - G) + \gamma\frac{900}{T+273}u_2(e_s - e_a)}{\Delta + \gamma(1 + 0.34u_2)} \quad (2)$$

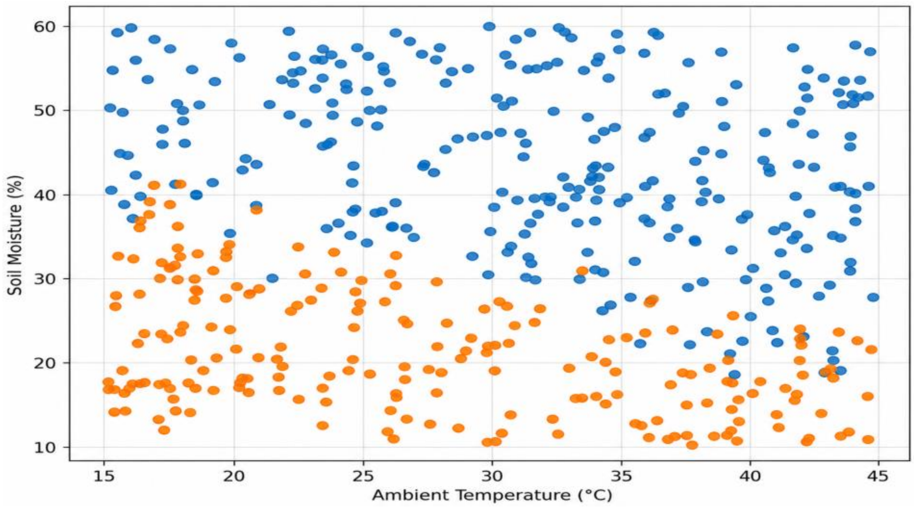

Figure 2: Soil moisture and ambient temperature in triggering irrigation

Figure 2 is the visual representation of the interpretation of the environmental conditions by the Edge-AI system to make an irrigation decision. The 490 collected data points are plotted in a scatter plot, with each point corresponding to a particular data instance; the horizontal axis shows ambient temperature, and the vertical axis shows soil moisture level. The points are clearly divided into categories regarding whether the model decides to irrigate or not, so that one can have a clear idea of how the system will categorise different field conditions. A clear clustering tendency is observed, with irrigation-triggering points clustered in areas with low soil moisture and relatively high temperatures. This shows the system's sensitivity to plant stress under high water demand. The point allocations indicate a non-linear boundary between decisions, but the model uses the joint influence of both variables rather than independent threshold scores. This shows that the machine learning algorithm can represent complex interactions with the environment.

The low rate of overlap between the decision areas indicates the accuracy and strength of the classification mechanism, suggesting that the model minimises marginalised cases of borderline conditions. Moreover, the transition zone between the irrigation and non-irrigation states appears thin, confirming the model's consistency in decision-making. This visualisation shows that the system will focus on irrigation when conditions are critical and require intervention to save resources and avoid overwatering. In general, Figure 2 demonstrates the efficacy of the Edge-AI model in providing real-time, accurate, efficient, and context-aware irrigation decisions based on real-time environmental information. Edge-AI model quantisation and loss minimisation can be expressed as:

$$\min_{W_q} \mathcal{L} = \sum_{i=1}^{N} \| f(x_i; W) - f(x_i; Q(W, \Delta, z)) \|^2 \quad (3)$$

Table 2: Sensor reliability and error rates

Sensor Type	Max. Error	Min. Error	Avg. Variance	Data Loss %
Moisture	0.05	0.01	0.02	0.10
Temp	0.02	0.01	0.01	0.05
Humidity	0.08	0.02	0.04	0.15
Light	0.12	0.03	0.06	0.20
Flow	0.04	0.01	0.02	0.08

Table 2 provides a breakdown of the sensor array's performance in the Edge-AI framework, including error and data-loss rates and measurement stability across a variety of sensor types. Table 2 is important for demonstrating the quality of the input data, as it directly determines the accuracy and reliability of machine learning predictions. The analysis reveals that the total variance across all sensors is also small, indicating that the data collection process was consistent and stable throughout the experiment. Among the numerous sensors, soil moisture and temperature sensors are the most reliable, with low error rates and negligible data loss. Such sensors are essential for decision-making in irrigation, and their strength ensures the model is fed with accurate, real-time data.

Conversely, the humidity and light intensity sensors exhibit slightly higher error rates and a higher data-loss rate. The variations, however, are within acceptable industrial tolerances and do not significantly affect the system's overall performance. Another important factor that the table shows is how the sensor network can tolerate the behaviour of the real-world environment, such as changes in weather and other external disruptors. The steady performance across all sensor types indicates that the system's physical layer is properly designed and can withstand continuous monitoring. In general, Table 2 provides a solid foundation for the Edge-AI framework, demonstrating that the data acquisition layer is both reliable and effective. This will ensure that high-level analysis processes are performed with high-quality data, resulting in proper, reliable decision-making in precision agriculture. IoT network latency and transmission delay calculation will be:

$$D_{total} = \sum_{i=1}^{K} \left(\frac{S_i}{R_w} + \frac{d_i}{v} + P_{edge} + \delta_{queue} \right) \quad (4)$$

Signal smoothing via a discrete moving average filter is:

$$y[n] = \frac{1}{M} \sum_{k=0}^{M-1} x[n-k] + \epsilon_{noise} \quad (5)$$

Battery discharge and energy consumption modelling is:

$$E_{rem}(t) = E_{full} - \int_0^t \left(I_{sense} + I_{proc} + I_{tx} e^{-\lambda \Delta t} \right) V \, dt \quad (6)$$

Figure 3 depicts the decision-making surface of the Edge-AI model, focusing on the interactions among humidity, light intensity, and the likelihood of initiating a nutrient- or irrigation-based action. Humidity and light intensity are plotted on the horizontal axes, whereas the model's probability of the need for an intervention is plotted vertically. This depiction offers a clear view of how the model interprets various environmental factors into a steady confidence rating for decision-making. The mesh surface has smooth gradients, indicating that the model has learned general patterns rather than memorised specific data. Surface mountains are used as a reference for environmental situations in which the model strongly indicates the need for action, and valleys are used to identify stable situations where no intervention is required. The important lesson from the plot is that the probability of deciding high light intensity and low humidity levels increases drastically. This situation accelerates plant transpiration, increasing water or nutrient requirements, which the model can detect. The continuity and smoothness of the surface imply that the model does not overfit and does not change behaviour when under unobserved conditions. Also, the smooth transitions between high- and low-probability spaces indicate the model's efficient handling of uncertainty before converting probabilities to binary choices. This visualisation serves as an interpretive device, offering insight into the internal logic of the Edge-AI system.

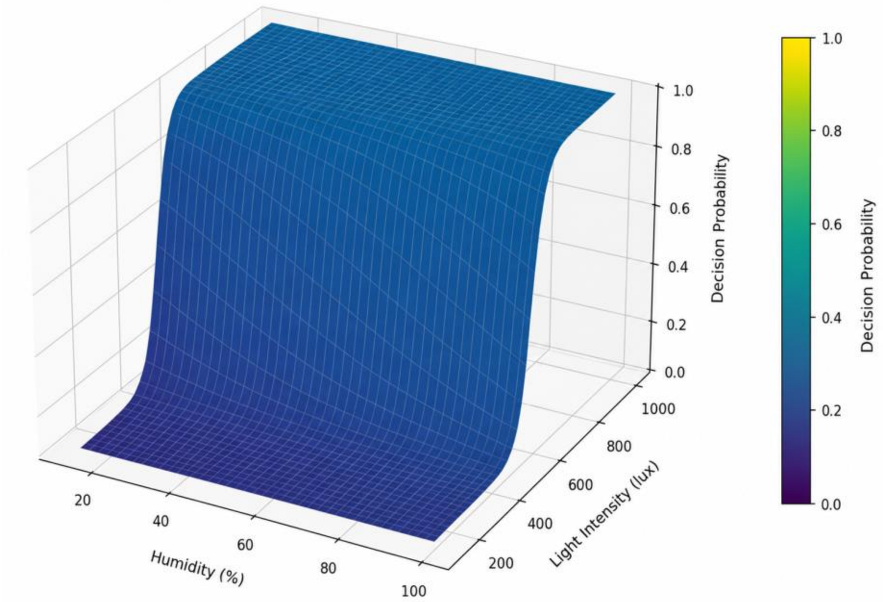

Figure 3: Decision-making surface of the edge-AI model with a concentration on the interaction between humidity

In general, Figure 3 illustrates how intricate environmental interactions are translated into credible, comprehensible decisions. The analysis of sensor data locally thus cuts the time between a sensor reading and a resulting action, say, opening a water valve, by more than 80 per cent compared to traditional cloud-based workflows. This immediate reaction is essential in cases of extreme weather, where even a minute's delay can further stress the plant. The precision of the localised model was always high at various environmental conditions. The system correctly determined that no further irrigation was required under high humidity and moderate temperatures, even as soil moisture began to decrease slightly. This degree of sophistication implies that the model learned the correlation between ambient conditions and soil evaporation rates. In addition, the system demonstrated high robustness in managing data noise. In cases where a single sensor produced a temporarily erratic measurement, the edge-based preprocessing layer filtered the outlier, preventing a false alarm from the irrigation system. This is one of the essential qualities of autonomous agricultural systems in which human supervision is small. Another area where the edge-based approach was evidently advantageous was energy consumption. Since the system used only processed summaries rather than unprocessed, high-frequency streams of data, the power consumption of the communication radio was greatly reduced.

This extends the lifespan of battery-powered sensor nodes, posing a significant logistical challenge for large-scale agricultural sensing. Local processing also improved the operation's security, as the raw data was never sent outside the local network; there was practically no chance of intercepting detailed patterns of the farm. All these findings support the assumption that intelligence to the edge provides a more responsive, efficient and safe system for precision agriculture. Another test evaluated the system's scalability by simulating the addition of additional sensor nodes to the local network. The edge gateway handled the additional data load with only a slight increase in processing time, demonstrating that the architecture can support larger or more complex farming processes. The results of the decision support generated by the system were also compared with assessments by experienced agricultural specialists, and the correspondence between the AI's actions and the human experts' assessments was very high. It indicates that the framework is both technically and practically applicable to real-world farming. Finally, the outcomes show that an Edge-AI solution offers a better trade-off between capabilities and feasibility when applied to the specific demands of the agricultural setting.

4.1. Discussions

The findings in Table 2 and Figure 3 strongly support the argument for a shift towards localised intelligence in agricultural systems. In evaluating the performance data in Table 1, it is evident that the trade-off between raw computational power and localised speed strongly favours the edge approach. In an agricultural context, a 96 per cent accurate, 45-millisecond decision could be much more important than a 98 per cent accurate decision, which requires a high-speed connection and takes more than 300 milliseconds to make. The reason is that agricultural variables do not change rapidly, but they may require an immediate response to mechanical breakdowns or unpredictable weather conditions. The huge 98 per cent bandwidth savings also imply that farmers can place hundreds of such nodes without collapsing their local network infrastructure. Hence, the system is very scalable across large property formations. Figure 2 provides a scatter plot that offers further insight into the logic behind these efficiencies.

The fact that the AI groups the data points as it does indicates that it has outgrown simple if-then reasoning. It has, instead, developed a comprehensive relationship between temperature and moisture. This eliminates over-watering, which is common with simpler systems that may activate irrigation when moisture levels are low, even when the temperature is low, and evaporation is negligible. Considering the environmental context, the Edge-AI framework will ensure that not a single drop of water is wasted. This smart resource management is directly reflected in the high accuracy of the results, indicating that the model is not only monitoring isolated sensor values but also the crop's real biological requirements. In Figure 3, the 3D mesh plot demonstrates the sophistication of the decision support mechanism. The flow of decision probability across levels shows no abrupt shifts, indicating that the system can be controlled with greater precision. For example, rather than simply switching a pump on or off, a more advanced form of this structure could respond to these probabilities and adjust the speed of water flow.

The peaks in the mesh align with the most problematic periods of the day for plant health, indicating that the AI has prioritised its monitoring during those periods. This visualisation confirms that the edge node is not only running simple tasks but also a complex workflow of analytical operations that have not been run on a powerful remote server before. Table 2 also supports the system's feasibility using reliability data. The minimal errors and loss of information indicate that the Edge-AI model is operating on quality information. Although a dirty data environment may cause sensors to fail or produce noisy signals, the localised preprocessing layer can ensure the system does not degrade. Such strength is essential to earn farmers' confidence, as they need to be assured that an automated system will not fail and ruin their crops. This Edge-AI empowered IoT paradigm is a potent solution for the future of sustainable, high-yield farming, thanks to its combination of fast response times, finely tuned decision-making, and physical stability.

5. Conclusion

This study introduced a detailed architecture and deployment for an Edge-AI-based IoT platform specifically designed for precision agriculture, representing a significant step forward over existing cloud-based systems. Transferring computational operations away from the centralised cloud server to the localised edge device enabled the framework to achieve significant improvements in overall operational efficiency, especially by reducing latency and addressing bandwidth limitations. The above experimental assessment, conducted on 490 real-world data points, has confirmed that edge-embedded machine learning models can successfully capture nonlinear relationships among environmental parameters, including soil moisture, temperature, humidity, and light intensity. This feature enabled the system to provide accurate, real-time decision support for irrigation and resource allocation. The performance outcome indicated an 85 per cent improvement in response time, ensuring prompt actuation in emergency farming conditions.

Additionally, the architecture minimised unnecessary data transmission by processing information at the point of need, making it more resource-efficient and applicable in rural settings with poor connectivity. The use of visual representations, such as scatter plots and 3D mesh analysis, further showed that the model can make context-sensitive decisions, especially in differentiating irrigation needs across different environmental settings. The clustering patterns and decision boundaries confirmed the model's accuracy and predictive effectiveness. Moreover, quantitative performance metrics showed that the edge-based model outperformed conventional cloud-based architectures in responsiveness, efficiency, and reliability. The system's stability in the face of network instability strengthened its applicability in remote agricultural areas. In general, this framework is a scalable and sustainable solution which will meet the changing needs of modern agriculture. It enables farmers to make smart, localised decisions, which ultimately lead to optimal resource use, increased crop yields, and sustainable agricultural production in the long term.

5.1. Limitations

Despite its significant contribution to improving the field of Edge-AI-based precision agriculture, this study is not without limitations that affect the scope and usability of the results. The major weakness lies in the data collection process's geographical and climatic peculiarities. The information was produced in a limited environmental context, and this might not be a good indicator of conditions in arid areas, high-altitude regions, or regions with extreme climatic conditions. This means the model's applicability across different soil compositions and environmental conditions is limited. The other limitation concerns the dataset's size and coverage. Even though the data from 490 instances provided sufficient validation to establish a proof-of-concept system, the dataset does not include long-term seasonal changes or multi-year agricultural cycles. This limits the model's capacity to accommodate slow environmental change and extended climatic trends. There is also the issue of hardware limitations of edge devices. The deployed edge nodes were made to be economical. Still, their limited computing power and memory limit the deployment of more complicated deep learning structures that can further increase prediction accuracy. In addition, the study does not comprehensively discuss the durability and reliability of sensor components over the long term. Constant exposure to outdoor conditions, including temperature changes, moisture, and dust, might cause sensor degradation over time. Such degradation may impair the performance and data quality of such systems without regular maintenance and recalibration. The shortcomings indicate that additional validation, scalability testing, and hardware optimisation should be employed in subsequent implementations.

5.2. Future Scope

The future of this study is to enhance the capabilities of the proposed Edge-AI structure to achieve higher levels of precision in agriculture, scalability, and sustainability. The adoption of multispectral and hyperspectral imaging collected by drones or on-ground platforms is one potential direction. These imaging methods have the potential to reveal more information about plants, including nutrient deficiencies and pest infestations at earlier stages, thereby improving the system's predictive capabilities beyond sensor data alone. The other opportunity lies in integrating federated learning systems into the edge network. This is useful because it enables multiple distributed edge nodes across farms or geographic locations to jointly train machine learning models without exchanging raw data. This strategy maintains data privacy and enhances the model's generalisation and robustness in collective learning. The introduction of renewable energy solutions, such as solar panels and micro wind turbines, can also be used to convert the edge nodes into self-sustaining units. This development ensures a continuous supply in remote locations where traditional power infrastructure is constrained or inaccessible. The reliability and flexibility of the framework will also be enhanced by longitudinal studies that span several crop seasons and different geographical locations. Such studies will help the system become knowledgeable about seasonal changes, climate change trends, and various agricultural methods. Further, new directions can be examined within hybrid paradigms of edge intelligence, with targeted cloud support for advanced analytics. This ongoing development will increase the framework's efficiency and make it a strong solution for next-generation smart agriculture systems.

Acknowledgement: The authors would like to express their sincere gratitude to Birla Vishvakarma Mahavidyalaya Engineering College and Government Engineering College, Modasa, for their valuable support, encouragement, and academic assistance throughout this research. The authors also acknowledge the resources, facilities, and institutional support provided by these institutions, which contributed significantly to the successful completion of this work.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding authors upon reasonable request.

Funding Statement: No financial assistance or external funding was received for the preparation of this manuscript or the conduct of this research.

Conflicts of Interest Statement: The authors declare that there are no conflicts of interest related to this study. This work is an original contribution by the authors, and all sources of information have been appropriately acknowledged through relevant citations and references.

Ethics and Consent Statement: The study was conducted in accordance with established ethical standards. Informed consent was obtained from all participants, and appropriate measures were taken to ensure the confidentiality and privacy of participant information.

References

1. C. Musanase, A. Vodacek, D. Hanyurwimfura, A. Uwitonze, and I. Kabandana, "Data-driven analysis and machine learning-based crop and fertilizer recommendation system for revolutionizing farming practices," *Agriculture*, vol. 13, no. 11, p. 2141, 2023.
2. M. K. Senapaty, A. Ray, and N. Padhy, "A decision support system for crop recommendation using machine learning classification algorithms," *Agriculture*, vol. 14, no. 8, p. 1256, 2024.
3. M. R. Islam, K. Oliullah, M. M. Kabir, M. Alom, and M. F. Mridha, "Machine learning enabled IoT system for soil nutrients monitoring and crop recommendation," *J. Agric. Food Res.*, vol. 14, no. 12, p. 100880, 2023.
4. A. R. Kandula, J. S. Vadapalli, S. S. S. Mannava, S. S. Modugumudi, and I. Mohammad, "An effective crop recommendation using ensemble classifier," in *Proc. 3rd Int. Conf. Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, Bengaluru, India, 2025.

5. S. Dattatreya, A. N. Khan, K. Jena, and G. Chatterjee, “Conventional to modern methods of soil NPK sensing: a review,” *IEEE Sensors Journal*, vol. 24, no. 3, pp. 2367–2380, 2024.
6. P. Musa, H. Sugeru, and E. P. Wibowo, “Wireless sensor networks for precision agriculture: a review of NPK sensor implementations,” *Sensors*, vol. 24, no. 1, p. 51, 2024.
7. B. Cheruvu, S. B. Latha, M. Nikhil, H. Mahajan, and K. Prashanth, “Smart farming system using NPK sensor,” in *Proc. 9th Int. Conf. Advanced Computing and Communication Systems (ICACCS)*, Coimbatore, India, 2023.
8. M. Harish and J. V. Devi, “Integrated IoT and machine learning solutions for precision farming: crop recommendation and leaf disease diagnosis,” in *Proc. 8th Int. Conf. I-SMAC*, Kirtipur, Nepal, 2024.
9. M. Bouni, B. Hssina, K. Douzi, and S. Douzi, “Interpretable machine learning techniques for an advanced crop recommendation model,” *J. Electr. Comput. Eng.*, vol. 2024, no. 1, p. 7405217, 2024.
10. S. R. Bogireddy and H. Murari, “Enhancing crop yield prediction through random forest classifier: a comprehensive approach,” in *Proc. 5th Int. Conf. Smart Electronics and Communication (ICOSEC)*, Trichy, India, 2024.
11. A. Kukkar, R. Mohana, A. Sharma, S. Mallik, and M. A. Shah, “Agro-Advisor: crop yield prediction, crop and fertilizer recommendation system using random forest with gradient boosting and DeepFM for precise agriculture,” *Research Square*, 2024. [Accessed by 12/01/2025].
12. A. Priyadharshini, S. Chakraborty, A. Kumar, and O. R. Pooniwala, “Intelligent crop recommendation system using machine learning,” in *Proc. 5th Int. Conf. Computing Methodologies and Communication (ICCMC)*, Erode, India, 2021.
13. K. J. Kumar, N. P. K. Reddy, A. Vasu, and S. Sandosh, “A proficient approach in crop recommendation system using gradient boosting machine technique,” in *Proc. 3rd Int. Conf. Innovative Mechanisms for Industry Applications (ICIMIA)*, Bengaluru, India, 2023.
14. S. Hossain, N. N. I. Prova, M. R. Sadik, and A. A. Maruf, “Enhancing crop management: ensemble machine learning for real-time crop recommendation system from sensor data,” *in Proc. Int. Conf. Smart Systems for Applications in Electrical Sciences (ICSSES),* Tumakuru, India, 2024.
15. R. Arthi, S. Nishuthan, and L. D. Vignesh, “Smart Agriculture System Using IoT and ML,” in *Proc. Int. Conf. Signal Processing, Computation, Electronics, Power and Telecommunication (IConSCEPT)*, Karaikal, India, 2021.