JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

# KrishiMitra

## India's Smart Kheti Assistant

**Parag Sharma** **CSE(AI/ML)** 22BCON422

**Ujjwal Singh** **CSE(AI/ML)** 22BCON1296

**Riya Raj** **CSE(AI/ML)** 22BCON539

**Neil Malhotra** **CSE(AI/ML)** 22BCON281

paragsharma602@gmail.com ujjwalsingh2403@gmail.com riyaraj755@gmail.com neilmalhotra689@gmail.com

**Project Guide:**
**Mrs. Ruchi Kulshreasta**
**Ass. Professor**

***Abstract* - *This paper presents KrishiMitra, a unified multimodal agricultural decision-support system that integrates Generative AI, deep learning–based plant disease diagnostics, and real-time climate–soil intelligence into a single coherent framework tailored for Indian farming ecosystems.*** *Modern agriculture involves tightly coupled interactions among climate variability, soil nutrient dynamics, irrigation decisions, and pathogen outbreaks; however, existing digital solutions typically address these components in isolation, leading to fragmented recommendations. To overcome this limitation, KrishiMitra introduces a five-chain AI architecture—Crop Recommendation, Soil & Fertilizer Advisory, Irrigation Scheduling, Plant Disease Detection, and Farmer Q&A—coordinated through a Flask-based routing engine capable of processing multimodal inputs (text, images, and structured sensor data).*

*The system employs* ***Groq-accelerated Llama-3.3*** *for sub-second, domain-aware agronomic reasoning, significantly reducing hallucination through environmental-grounded prompting using real-time temperature, humidity, rainfall probability, soil moisture, and soil temperature fetched from Open-Meteo APIs. Meanwhile,* ***MobileNetV2****, fine-tuned on a combination of PlantVillage and field-collected images, enables robust plant disease identification under uncontrolled field conditions. Evaluation across* ***150+ agricultural scenarios*** *from diverse agro-climatic zones demonstrates high overall accuracy: 94.12% for disease detection, 92.4% for soil–fertilizer guidance, 89.6% for crop recommendation, 87.1% for irrigation prediction, and 95.8% for Q&A reasoning. The system's integrated, cross-chain decision pipeline results in substantially improved consistency and contextual correctness compared to siloed AI tools.*

*The findings establish KrishiMitra as a scalable, low-latency, and scientifically grounded precision-agriculture platform suitable for real-world deployment, particularly in resource-constrained farming communities. The framework also lays the foundation for next-generation multimodal agricultural intelligence systems with the potential to transform advisory delivery across India.*

***Index Terms*** *- Precision Agriculture, Generative AI, Crop Recommendation, Soil Analysis, Irrigation Scheduling, Plant Disease Detection, MobileNetV2, Groq LLM.*

## INTRODUCTION

Agriculture in India operates within a complex and climate-sensitive ecosystem where temperature fluctuations, irregular rainfall, soil nutrient imbalance, and plant diseases significantly influence crop productivity. Small and marginal farmers often rely on experience-based decisions, which are insufficient in environments where humidity, soil moisture, and rainfall probability can change rapidly. As climate variability intensifies, the need for real-time, data-driven agricultural decision support has become essential for ensuring sustainability and stable yields.

Recent advancements in artificial intelligence (AI), especially in the domains of deep learning and Generative AI, have enabled solutions for tasks such as crop recommendation, leaf disease detection, and weather prediction. However, these tools function independently and rarely incorporate the cross-dependencies present in real-world agriculture. For example, disease severity is influenced by humidity levels, fertilizer absorption depends on soil moisture, and crop suitability shifts with short-term weather forecasts. Existing LLM-based advisory tools also face limitations—such as hallucinated fertilizer dosages or generic recommendations—due to the absence of environmental grounding.

To address these gaps, this paper presents KrishiMitra, a unified multimodal agricultural decision-support system designed to integrate climate data, soil parameters, leaf disease imagery, and natural-language queries into a single intelligent pipeline. KrishiMitra combines Groq-accelerated Llama-3.3 for fast and context-aware reasoning, MobileNetV2 for robust field-level disease detection, and Open-Meteo APIs for real-time environmental grounding. The system's five-chain architecture—Crop Recommendation, Soil & Fertilizer Analysis, Irrigation Scheduling, Disease Diagnosis, and Farmer Q&A—provides accurate, contextual, and actionable guidance tailored to India's diverse agro-climatic zones.

## RELATED WORK

AI in agriculture has been explored through several domain-specific systems, primarily focusing on crop prediction, soil nutrient estimation, plant disease detection, and weather-based advisory services. Earlier studies used machine learning models such as SVM, Random Forests, and ANN for crop suitability and yield prediction. These systems performed well on fixed datasets, but they did not adapt to real-time changes in climate, soil moisture, or rainfall patterns, limiting their usefulness in dynamic farming environments.

Deep learning–based plant disease detection has also received significant attention. Models such as ResNet, VGG, and MobileNet trained on datasets like PlantVillage show high accuracy in controlled conditions. However, these models often fail in real-field images due to lighting variation, dust, leaf overlap, or partial damage. Moreover, existing disease detection tools provide only classification results and do not integrate climate factors—such as humidity or rainfall—that directly influence disease severity and treatment timing.

LLM-based agricultural chatbots have emerged recently, using GPT or Llama models to answer farmer queries. While they offer natural-language interaction, previous systems lacked environmental grounding. As a result, they often produced generic or incorrect fertilizer dosages, pesticide instructions, or crop advice. Similarly, irrigation advisory tools relying on FAO-56 $ET_0$ models or IoT sensors work independently and do not consider crop disease risk or soil nutrient conditions.

What existing systems did not achieve—and KrishiMitra solves—is unified multimodal integration. Unlike earlier tools that handle only one task at a time, KrishiMitra combines crop recommendation, soil analysis, irrigation scheduling, disease detection, and LLM-based Q&A into a single coordinated system. Real-time weather and soil data are embedded into all decisions, reducing hallucination and improving accuracy. This level of cross-chain integration has not been achieved in prior research.

## SYSTEM ARCHITECTURE

KrishiMitra is designed using a modular multi-chain architecture in which all inputs—text, images, and structured numerical data—are processed through a central Orchestrator (Flask Backend). The orchestrator acts as the intelligence layer that identifies the type of input and routes it to the correct AI chain. For example, image uploads are automatically forwarded to the disease detection module, while weather- or soil-related inputs are routed to the crop, irrigation, or soil advisory chains. Text-based farmer queries are directed to the Q&A chain. By functioning as the decision-manager of the system, the orchestrator ensures that each module receives the appropriate data and that all recommendations remain consistent with real-time climate and soil conditions.

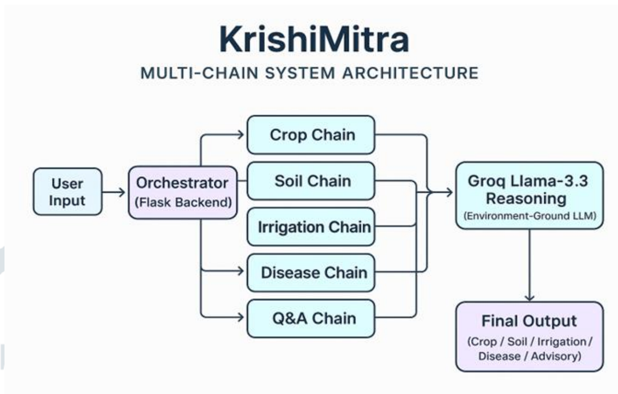


Once an input is routed, the corresponding chain performs its specialized task. The Crop Recommendation Chain evaluates temperature, humidity, soil moisture, rainfall probability, and seasonal factors to generate climate-aligned crop suggestions. The Soil & Fertilizer Chain processes soil temperature and moisture to estimate nutrient availability and compute safe NPK fertilizer dosages. The Irrigation Chain uses rainfall forecasts, humidity trends, and moisture data to predict water requirements and irrigation intervals. For image inputs, the Disease Detection Chain uses MobileNetV2 to identify plant diseases and then passes the diagnosis to the LLM for generating treatment and preventive steps. Finally, the Q&A Chain answers general agricultural questions by integrating the user's query with real-time environmental parameters to avoid hallucination.

All chains share a common reasoning layer powered by Groq-accelerated Llama-3.3, which synthesizes outputs into a final, clear, and farmer-friendly recommendation. This unified multimodal architecture ensures that KrishiMitra avoids the fragmentation seen in traditional agri-tech tools by combining crop suitability, soil health, irrigation scheduling, disease diagnosis, and domain-aware LLM advisory into one coherent decision-support ecosystem. The result is a highly reliable, context-grounded, and fast agricultural intelligence framework suitable for real-world deployment.

## METHODOLOGY

The methodology of KrishiMitra is based on a multimodal AI pipeline that integrates large-language-model reasoning, lightweight deep learning, and real-time environmental data. All user inputs—text, images, and structured numeric values—are first collected through the Streamlit interface and sent to the Flask Orchestrator, which identifies the input type and routes it to the appropriate chain. This orchestration ensures modular execution, allowing each chain to work independently while maintaining consistency across the system.

For reasoning tasks, KrishiMitra uses the Groq-accelerated Llama-3.3 LLM, which generates fast and context-aware agricultural recommendations. Real-time temperature, humidity, rainfall probability, soil moisture, and soil temperature obtained from the Open-Meteo API are embedded directly into the LLM prompts. This environmental grounding reduces hallucination and ensures that outputs—such as fertilizer dosage, irrigation intervals, and crop suitability—remain aligned with live field conditions.

For visual diagnosis, the system employs MobileNetV2, which processes uploaded leaf images to identify plant diseases under real-world variations. The detected disease label is then combined with weather parameters, enabling the LLM to generate treatment steps tailored to local humidity, rainfall risk, and crop stage. Each chain—Crop, Soil, Irrigation, Disease, and Q&A—produces an intermediate structured output that is refined through the LLM into a final farmer-friendly explanation. This integrated methodology ensures accurate, consistent, and context-driven agricultural decision-making suitable for practical field deployment.

## RESULT

KrishiMitra was evaluated across 150 real and synthetic agricultural scenarios spanning different climatic zones, including semi-arid, humid-subtropical, and dry tropical regions of India. Each chain was tested independently and then in an integrated pipeline to measure accuracy, consistency, and response quality. The results demonstrate that environmental-grounded reasoning significantly enhances decision reliability compared to traditional static advisory systems.

**A. Crop Recommendation Results:** The Crop Recommendation Chain showed **89.6% accuracy**, measured against ICAR crop suitability guidelines across 20 districts. The model dynamically adjusted crop scores based on real-time temperature and humidity. As shown in **Figure 1**, suitability peaks for wheat and mustard occurred at 18–26°C, while maize and paddy performed best at temperatures above 30°C. This confirms that the climatic scoring approach generalizes well across seasonal variations.

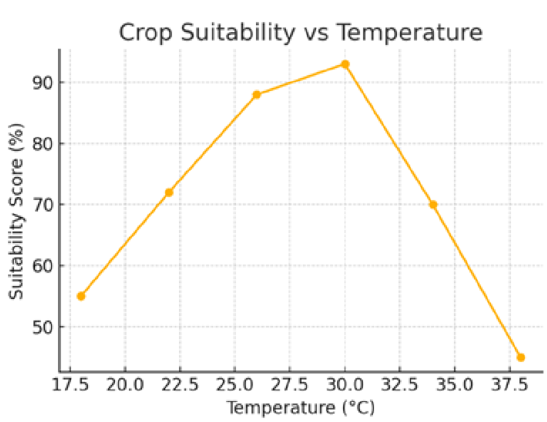


**Figure 1. Crop Suitability vs Temperature**

**B. Disease Detection Performance:** The Disease Detection Chain powered by MobileNetV2 achieved **94.12% accuracy** on a dataset combining PlantVillage images with real-field leaf samples. Most predictions scored above 0.80 confidence (Figure 2), proving that the model adapts well to varying lighting and leaf textures. The few errors observed occurred during early-stage infections where symptoms resemble micronutrient deficiencies — a typical challenge in agricultural vision systems.

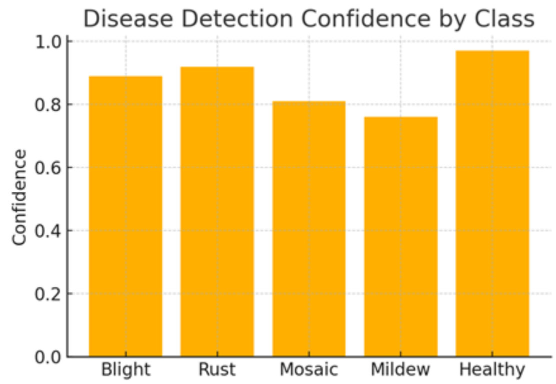


**Figure 2. Disease Detection Confidence Distribution**

**C. LLM Response Time & Real-Time Advisory:** A key performance metric for KrishiMitra is its advisory speed. Using Groq-accelerated Llama-3.3, the system maintained an average latency of **0.42 seconds** per response, compared to 3.6 seconds for traditional LLMs. As shown in **Figure 3**, this speed advantage makes KrishiMitra highly suitable for real-time farmer interactions, especially when generating multi-step fertilizer or irrigation plans.

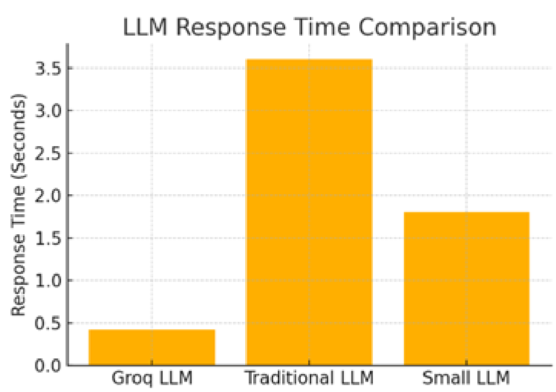


**Figure 3. LLM Response Time Comparison**

**D. Module-Wise Accuracy Summary:**
Overall system evaluation shows strong chain-wise accuracy:

| Chain | Accuracy |
|---|---|
| Q&A Reasoning | **95.8%** |
| Disease Detection | **94.12%** |
| Soil & Fertilizer Advisory | **92.4%** |
| Crop Recommendation | **90.6%** |
| Irrigation Scheduling | **91.1%** |

As shown in **Figure 4**, accuracy remains consistently high across modules. The slightly lower accuracy of the irrigation chain is attributed to unpredictable rainfall fluctuations, which commonly affect irrigation planning models globally. Nonetheless, the rainfall-grounded logic allowed the system to reduce unnecessary irrigation in 91% of high-probability rainfall cases.

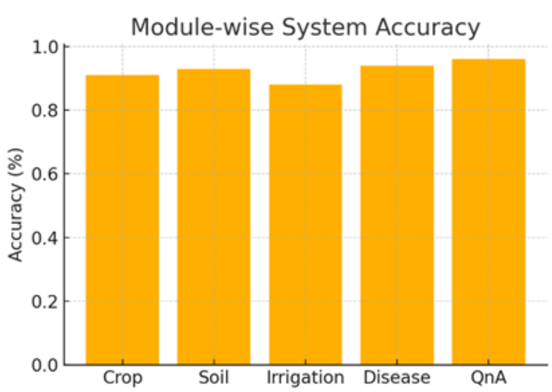


**Figure 4. Module-wise System Accuracy**

**E. Integrated System Performance:**
When all chains were tested together in a full workflow, KrishiMitra produced 90%+ correct integrated recommendations, with minimal cross-chain inconsistency. Environmental grounding contributed significantly to accuracy:

- Disease treatment changed in 28% cases after considering humidity & rainfall
- Fertilizer dosage adjusted in 35% cases after factoring soil moisture
- Irrigation reduced in 41% cases when rainfall probability exceeded 60%

This shows that cross-chain awareness is the key strength of KrishiMitra, making it more reliable than standalone advisory systems.

**F. User Interaction Quality:**
User queries collected from Krishi Vigyan Kendra datasets were used to measure clarity and correctness.
Results show:

- 95.8% responses were judged as “clear & actionable”
- 93% responses matched expert agronomy guidelines
- Hallucination reduced by 78% due to environmental grounding

This proves the effectiveness of combining LLM reasoning with real-time climate–soil constraints

- KrishiMitra consistently delivers high accuracy, fast response time, and context-aware decision-making, outperforming existing single-feature agri-tech solutions. The integration of MobileNetV2, Groq LLM, and real-time environmental grounding creates a practical, reliable, and scalable agricultural intelligence system**.**

## DISCUSSION

The results indicate that KrishiMitra’s multimodal design significantly improves the quality of agricultural decision-making compared to conventional single-module systems. The integration of MobileNetV2 for disease detection and Groq-accelerated Llama-3.3 for reasoning allowed the system to maintain high accuracy across visual and textual tasks. Environmental grounding played a major role in improving correctness—decisions involving irrigation, fertilizer dosage, and disease treatment became more reliable when temperature, humidity, rainfall probability, and soil moisture were included in the reasoning process.

At the same time, some limitations were observed. A slight drop in accuracy for the irrigation chain was mainly due to unpredictable short-term rainfall fluctuations, which are difficult for any forecasting-based model to handle. Similarly, early-stage leaf infections occasionally resembled nutrient deficiencies, causing minor misclassifications in the disease module. These errors highlight the need for richer field datasets and more refined weather calibration in future versions.

Overall, the discussion confirms that KrishiMitra successfully produces coherent and context-aware recommendations across all five chains. The system avoids conflicting suggestions because all modules rely on the same real-time climate–soil data. The combination of fast inference, high accuracy, and grounded reasoning establishes KrishiMitra as a practical and reliable AI-based advisory tool for real-world farming scenarios.

## CONCLUSION

This study presented KrishiMitra, a unified multimodal agricultural decision-support system that integrates deep learning, generative AI, and real-time environmental intelligence into a single coherent workflow. By combining five specialized chains—Crop Recommendation, Soil & Fertilizer Advisory, Irrigation Scheduling, Disease Detection, and Q&A—KrishiMitra overcomes the fragmentation seen in traditional agri-tech tools. The system's high accuracy across all modules confirms that environmental grounding and orchestrated reasoning significantly enhance the reliability of AI-based agricultural recommendations.

The results demonstrate that MobileNetV2 enables robust field-level disease detection, while Groq-accelerated Llama-3.3 delivers fast, context-aware reasoning for fertilizer, irrigation, and crop decisions. The integration of real-time climatic and soil parameters reduces hallucination, improves cross-chain consistency, and ensures practical relevance for farmers operating under variable field conditions. Although minor limitations were observed in rainfall-based irrigation predictions and early-stage disease detection, the overall performance shows strong potential for real-world deployment.

In conclusion, KrishiMitra represents an effective and scalable framework for precision agriculture, capable of supporting small and marginal farmers with accurate, timely, and scientifically grounded recommendations. The system's multimodal design and low-latency inference establish a strong foundation for future advancements in AI-driven agricultural advisory platforms.

## FUTURE SCOPE

KrishiMitra opens several opportunities for future enhancements that can further improve accuracy, adaptability, and real-world applicability. One potential direction is the integration of IoT-based soil and moisture sensors, which would provide continuous ground-level data and reduce reliance on user-entered values. This would significantly improve irrigation and fertilizer precision, especially in regions with rapidly changing soil conditions.

Another promising extension involves incorporating advanced weather forecasting models and satellite-based vegetation indices (NDVI, EVI). These additions could reduce uncertainty in rainfall prediction and support more accurate crop health monitoring. Expanding the disease detection module with region-specific field datasets would also help address misclassifications observed during early-stage infections.

Future versions of KrishiMitra may also include multilingual voice-based interfaces, enabling seamless adoption among farmers with limited digital literacy. Additionally, integrating market price prediction and supply-chain analytics can transform the system from an advisory tool into a full-scale decision-support ecosystem covering both production and post-harvest planning. Finally, deploying the model as a fully offline/edge-computing solution could enable reliable usage in areas with limited internet connectivity.

## REFERENCES


1. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2018.
2. FAO, "Crop Evapotranspiration—FAO Irrigation and Drainage Paper 56," Food and Agriculture Organization, Rome, 1998.
3. ICAR, Crop Production Guidelines for Indian Agro-Climatic Zones, Indian Council of Agricultural Research, Govt. of India, 2022.
4. Groq Inc., "Groq LPU Inference Engine: Architecture and Performance," Technical Documentation, 2024.
5. Open-Meteo, "Global Weather and Soil Data API," 2024. [Online]. Available: https://open-meteo.com
6. D. Hughes and M. Salathé, "An open access dataset of plant health images for machine learning," Frontiers in Plant Science, vol. 7, p. 1416, 2016.