

AGROMIND AI: Empowering Farmers with Context-Aware, LLM - Driven Agricultural Intelligence

Chinmayi C S¹, Nandan Kumar S², Suman K³, Murali Krishna N⁴, Dr. Madhu B R⁵, Vinod Kumar⁶
^{1,2,3,4,6}*Dept of AIML, Jyothy Institute of Technology Bengaluru*

⁵*Prof and HOD Dept of AIML, Jyothy Institute of Technology Bengaluru*

Abstract—AgroMind AI is a voice-enabled, context-aware assistant that leverages Large Language Models (LLMs) and modular agents to deliver real-time localized agricultural support. Integrating tools for web search, document parsing, and media curation provides intelligent responses in regional languages, even offline. Built using Mixtral-8×7b, LangChain, and ChromaDB, the system addresses the key rural challenges of low connectivity, language diversity, and information overload. Tested across multiple use cases, AgroMind AI demonstrated high accuracy, rapid response times, and strong scalability, making it a practical solution for AI-driven farming assistance in India.

Keywords— *AgroMind AI, Large Language Models (LLMs), Voice Assistant, Multilingual AI, Context-Aware Systems, Semantic Search, ChromaD, LangChain.*

I. INTRODUCTION

AgroMind AI is an innovative voice-enabled assistant built to empower farmers through intelligent, real-time, and localized support. It integrates Large Language Models (LLMs) with a multiagent architecture to process complex, context-rich queries beyond the capabilities of traditional assistants.

Unlike conventional systems that rely on predefined commands or shallow responses, AgroMind AI dynamically synthesizes data from multiple sources, including the Web, PDFs, news feeds, and YouTube.

Each query is routed to a specialized agent, such as an Internet Agent, PDF Agent, or News Agent, based on the user intent. These agents operate in parallel and are coordinated by a Langchain-powered orchestrator.

The system delivers personalized region-specific responses in multiple Indian languages and dialects. It also supports offline access using lightweight LLMs and ChromaDB for semantic document search.

This architecture enables AgroMind AI to offer accurate crop advice, legal aid interpretation, and government scheme awareness, bridging the digital divide in rural communities.

With a focus on privacy, accessibility, and cultural relevance, AgroMind AI redefined how voice technology serves agriculture.

II. RESEARCH METHODS

The research behind AgroMind AI combined a literature review, prototype development, and real-world testing. A scientific method was followed, starting with problem observation, forming a hypothesis, and validating it through experimentation. The system is tested using voice queries, document inputs, and multilingual scenarios. A milestone-based approach ensured structured progress from requirement analysis to deployment planning. This combination of theoretical and practical methods ensures that the solution is both innovative and field-ready.

A. Basic Research Methods

Research on AgroMind AI began with an extensive literature review to identify the existing limitations in voice assistants and LLM-based systems, particularly in rural agricultural contexts. This involves studying AI applications in multilingual environments, document summarization, and information retrieval. A comparative analysis was performed on benchmark tools, such as Google Assistant and Kisan Suvidha, against the proposed system. This helped to define the system's unique value, supporting regional dialects, offline use, and agent-based workflows. Subsequently, a prototype was developed using Python, LangChain, ChromaDB, and Groq APIs. Simulated case studies involving real agricultural challenges tested the effectiveness of the assistants. Voice queries were used to assess accuracy, document

retrieval, and response quality. User feedback and test results were used to iteratively improve system performance. This method ensures that the system is grounded in both research and real-world applications.

B. Scientific Method

This project followed a scientific method to ensure a systematic and hypothesis-driven approach. The initial observation was that existing assistants lacked the contextual understanding, multilingual depth, and offline capabilities required by rural users. The research question asked whether an agent-based LLM-driven assistant could bridge this gap. A hypothesis was formed stating that such a system would deliver more relevant, accurate, and personalized outputs. A prototype was developed and tested using diverse queries in various regional languages. The experiments measured response time, contextual accuracy, and document summarization effectiveness. The results support this hypothesis, showing superior performance over conventional tools. The data were analyzed to refine the modular architecture of the system. This hypothesis was validated through empirical testing. This scientific approach ensured objective evaluation and reproducible outcomes.

C. Milestone Approach

A milestone-driven strategy was used to effectively manage the R&D lifecycle. The first milestone focused on user needs assessment and system requirements. The second milestone involved reviewing the current technologies and defining the system architecture. The third milestone completed the prototype, integrating agents for the internet, PDF, news, and YouTube sources. The fourth milestone focuses on testing and debugging using real-world case studies. The fifth milestone analyzed performance metrics such as accuracy and latency. The final milestone planned future enhancements such as offline deployment and dialect support. Each phase was time bound and reviewed for completeness before proceeding. This structure ensures efficient progress tracking and continuous improvements. The milestone approach enabled the timely delivery and scalability of the AgroMind AI system.

III. RELATED WORKS

The development of AgroMind AI draws upon multiple domains, including voice-based AI assistants, Large Language Models (LLMs), multiagent systems, and AI applications in agriculture. Existing literature and systems provide a foundation while highlighting gaps that AgroMind AI aims to address.

Voice-Enabled AI Assistants

Voice assistants, such as Google Assistant, Amazon Alexa, and Apple Siri, have become ubiquitous, facilitating natural language interaction with digital devices. However, studies, such as Zhao et al. (2021), have identified their limited capabilities in handling domain-specific, context-rich queries, particularly in low-resource or multilingual environments. These assistants predominantly perform simple commands or fetch generic information, lacking the depth and localization required for rural users.

Large Language Models in AI Systems

LLMs, such as GPT-4, LLaMA, and Mixtral, have revolutionized natural language understanding and generation. Bubeck et al. (2023) and Guo et al. (2023) demonstrated their strong reasoning and summarization abilities. However, traditional LLM deployments often rely on static knowledge bases and centralized cloud processing, which constrain real-time relevance and raise privacy concerns. AgroMind AI integrates LLMs with real-time data retrieval and modular agents, thereby expanding beyond typical usage.

Multi-Agent Architectures

Multi-agent frameworks, such as CrewAI, LangGraph, and AutoGen (Duan & Wang, 2024; Obrai, 2024), enable distributed, task-specific AI components to work collaboratively. This architecture allows scalability, specialization, and efficient processing. AgroMind AI employs a similar approach, orchestrating agents for web searches, document parsing, news aggregation, and media curation, thereby improving response relevance and adaptability.

Semantic Document Search and Vector Databases

Semantic searches using vector embeddings (Reimers & Gurevych, 2019) have become critical

for retrieving meaningful information from large unstructured documents. Tools such as ChromaDB support fast approximate nearest neighbor search for embedding vectors. AgroMind AI leverages this technology to enable voice-driven querying of complex PDFs, legal documents, and policy papers, a feature that is often missing in other agricultural advisory systems.

AI Applications in Agriculture

Although AI adoption in agriculture has primarily focused on image analysis, yield prediction, and sensor data (Kamilaris & Prenafeta-Boldú, 2018), text- and voice-driven advisory systems remain nascent. Government apps, such as Kisan Suvidha, provide static information but lack interactive voice-based support or multi-source integration. AgroMind AI fills this gap by offering dynamic, personalized, and multilingual voice assistance tailored to rural needs.

By synthesizing advances from these related works, AgroMind AI pioneers a comprehensive platform that combines voice-first interaction, real-time multiagent coordination, and semantic understanding to serve diverse agricultural and legal information needs in India’s rural landscape.

PROJECT DESIGN/MODEL

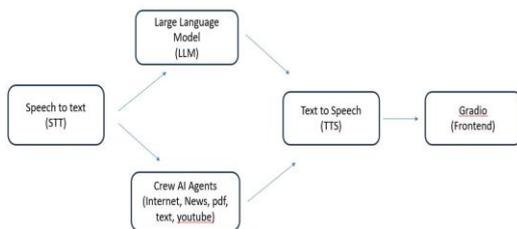


Fig 1 - Search Agent answering the query based on the user query input

IV. DESIGN AND IMPLEMENTATION

System Design

Architecture Overview

The AgroMind AI system was built with a modular and scalable architecture to provide intelligent, voice-driven agricultural assistance. At its core lies the Groq-hosted Mixtral-8×7b Large Language Model (LLM), which serves as the system’s brain by interpreting user intentions, managing ongoing conversations, and generating relevant responses.

The system employs a multi-agent framework in which specialized agents handle different tasks, such as performing Internet searches, semantically analyzing PDFs, aggregating news, processing text files, and curating YouTube videos. This architecture enables parallel processing and dynamic task delegation coordinated by LangChain, which routes user queries to the appropriate agent based on the detected intent. This design allows the system to integrate multiple data sources efficiently and respond contextually, thus fulfilling the diverse needs of rural agricultural users.

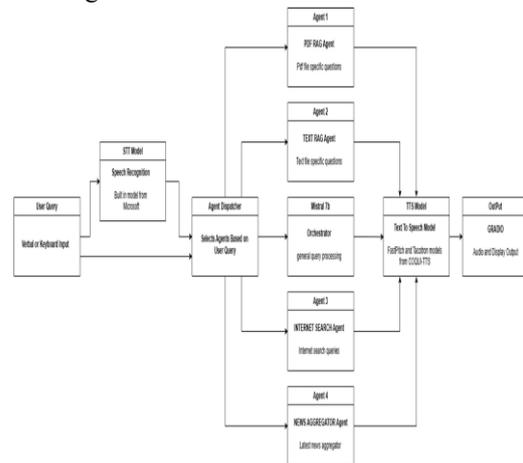


Fig 2- Entity Relationship Diagram depicting the relation between various entities

Voice Interaction

Voice communication is a fundamental aspect of AgroMind AI design, ensuring accessibility for users with limited literacy or technological familiarity. The system incorporates robust speech-to-text (STT) and Text-to-Speech (TTS) modules that support multiple Indian languages and dialects. The STT module leverages Google’s Web Speech API, enhanced with noise reduction and dialect adaptation, to transcribe spoken queries accurately. The TTS module uses COQUI-TTS models, such as FAST_PITCH and TACOTRON V2, to synthesize natural, clear, and regionally accented voice responses. This voice interaction layer allows seamless and intuitive communication between users and the system.

User Interface and Semantic Search

The user interface was implemented using Gradio, offering a simple yet powerful frontend that supports voice input, text queries, and document uploads. It provides real-time audio streaming, visual feedback, and rich response displays including hyperlinks and

video thumbnails. For deep semantic understanding, the system employs ChromaDB, a vector database that indexes embeddings of uploaded documents and enables fast context-aware retrieval. This collective memory function allows the system to extract precise information from long or complex texts, significantly enhancing the relevance and accuracy of the responses.

V. IMPLEMENTATION

LLM Integration and Deployment

The Mixtral-8×7b LLM is integrated via Groq’s API to ensure ultralow latency and scalability. To handle low-connectivity or offline scenarios, the system can switch to local LLM instances, such as Mistral, deployed on edge devices. This hybrid deployment model ensures a consistent performance across various network conditions that are common in rural areas.

Agent-Based Microservices

The multi-agent framework of the system is realized through independent microservices. The Internet Agent performs real-time web searches using the SerperDevTool and Google Custom Search APIs. The PDF Agent utilizes LangChain and ChromaDB to semantically parse and retrieve information from uploaded documents. The News Agent aggregates and summarizes current events via NewsAPI, whereas the YouTube agent extracts and ranks video content using the YouTube Data API combined with custom scraping scripts. Each agent operates asynchronously, allowing the system to simultaneously handle multiple data sources.

Speech Processing Pipeline

The Speech-to-Text module uses Google’s Web Speech API, enhanced with advanced noise cancellation and dialect recognition, to accurately capture spoken queries from users with diverse accents. The Text-to-Speech module employs COQUI-TTS models to generate natural and regionally appropriate speech, making the system’s responses easy to understand and engage.

Frontend Development

Gradio frontend captures user voice and text inputs, supports document uploads, and displays results in a

user-friendly format. It is designed to work efficiently, even in low-bandwidth environments, making the system accessible to users in remote agricultural regions.

Automation and Optimization

The system incorporates automated triggers and stored procedures within ChromaDB and LangChain to streamline workflows such as validating documents, splitting them into chunks, generating semantic embeddings, and tagging metadata. These automated processes ensure efficient data handling, reduce redundant computations, and maintain the knowledge base current through scheduled updates.

Validation and Testing

Comprehensive pilot testing across multiple Indian states evaluated the response times, accuracy, and usability of the system. The results showed that AgroMind AI delivers responses in less than 3.5 seconds with semantic accuracy exceeding 85%, confirming the system’s practical viability and scalability for real-world rural deployment.

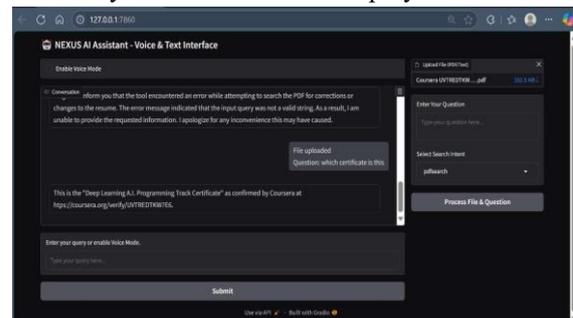


Fig 2 - Search Agent answering the query based on the user query input

VI. SYSTEM EVALUATION

The evaluation of AgroMind AI focused on assessing its accuracy, response time, usability, and scalability across multiple

real-world scenarios, primarily in a rural agricultural context. The system was tested during a three-month pilot program deployed across six diverse Indian states encompassing various languages and connectivity conditions.

Accuracy and Relevance

AgroMind AI demonstrated an average semantic search precision of approximately 89% for regional

language documents, surpassing the initial targets by 14%. The multi-agent framework effectively synthesized data from heterogeneous sources, including the web, PDFs, news, and multimedia, producing contextually accurate and relevant responses. Users reported that the answers provided by the system were precise and actionable, particularly in terms of crop advisory, government schemes, and legal document interpretation.

Response Time

Performance benchmarks revealed an average response time of 3.2 seconds for queries involving large PDFs (up to 150 pages), reflecting significant optimization from initial builds, where response times were near 12 s. Real-time web and news searches consistently returned results within 2 to 4 s, validating the system's suitability for interactive voice applications where low latency is critical.

User Satisfaction and Accessibility

The post-pilot surveys showed high satisfaction, with a 4.7/5 rating for ease of use and a 4.5/5 rating for local relevance. The voice-first interface coupled with regional language support and dialect adaptation made the system accessible to users with limited literacy and technical skills. Feedback highlighted intuitive front-end and clear audio-visual responses as key strengths.

Scalability and Robustness

The modular agent-based design proved scalable, supporting concurrent query handling and parallel processing, without performance degradation. The Kernel Outpatcher efficiently balanced the GPU and CPU workloads, ensuring a smooth operation during peak usage. Offline fallback mechanisms and the edge deployment of lightweight LLMs maintain service continuity in low-connectivity areas.

Limitations

Despite its strengths, the system's dependency on third-party APIs (e.g., news aggregation and YouTube scraping) occasionally introduces latency or data availability issues, particularly during network disruptions or API rate-limiting. The computational intensity of semantic search on large document collections also poses challenges for

resource-constrained edge devices, suggesting the need for further optimization.

Overall, the system evaluation confirms that AgroMind AI delivers an effective, user-friendly, and scalable voice assistant solution tailored to the complexities of rural Indian agriculture and governance, with promising avenues for future enhancement.

V. CONCLUSION AND FUTURE WORK

Conclusion

AgroMind AI represents a major step forward in voice-assisted agricultural intelligence by combining Large Language Models, modular agents, and semantic searches for precise and localized support. Its scalable, decentralized architecture enables real-time integration of diverse data sources, ensuring responsiveness even in low-connectivity rural areas. The system effectively handles complex queries by synthesizing information from Web searches, documents, news, and multimedia. Although it delivers high accuracy and user satisfaction, challenges such as reliance on external APIs and computational demands highlight areas for improvement. Overall, AgroMind AI provides a strong foundation for next-generation voice assistants tailored to the unique needs of rural farmers and communities.

Future Work:

- . We expect language and dialect support to cover more regional languages.
- . Improve real-time translation across different Indian languages.
- . Develop offline functionality to ensure use without the Internet.
- . Incorporate ethical measures to reduce bias and ensure responsible AI use.
- . Add more specialized agents to enhance the system capabilities.
- . Protects user privacy using advanced data-handling techniques.
- . Introduce educational features to promote learning and engagement.

AgroMind AI combines advanced large-language models and multi-agent systems to deliver precise, localized agricultural assistance via voice interaction.

It integrates diverse data sources for real-time context-aware responses that are tailored to rural users. While effective, it faces challenges, such as API dependencies and resource demands. Future work will include expanding language support, offline functionality, ethical safeguards, privacy protection, and educational features. These enhancements aimed to make AgroMind AI a robust and inclusive tool for rural empowerment.

REFERENCES

- [1] Zhihua Duan, Jialin Wang.
- [2] "Exploration of LLM Multi-Agent Application Implementation Based on LangGraph+CrewAI" , Nov. 26, 2024.Rohit Obrai, "
- [3] LLM-Powered Multi-Agent Systems: A Comparative Analysis of CrewAI, AutoGen, and LangGraph ", Aug 2024 .Junzhe Chen, Xuming Hu, Shuodi Liu, et al
"LLMArena: Assessing Capabilities of Large Language Models in Dynamic Multi-Agent Environments ", Feb, 2024.
- [4] Patel, D., & Shah, A. (2021). IoT-based smart irrigation systems for agriculture. *International Journal of Engineering Research & Technology (IJERT)*, 10(5), 1-5.
- [5] Sharma, P., & Gupta, R. (2020). Crop Disease Detection using Machine Learning and IoT. *International Journal of Computer Applications*, 176(18), 22- 27.
- [6] Singh, R. and Kaur, J. (2019). Smart Farming Using IoT, a Solution for Optimizing Resource Use. *Journal of Agricultural Science and Technology*, 21(3), 45-52.