Open Access International Journal of Science and Engineering

ISSN: 2456-3293

Volume 9 Issue 1s (2026) | Pages: 73-79

AgriTalk: A Trustworthy and Risk-Aware Bilingual AI Ecosystem for Multi-Crop Farm Management

Sakshi D. Shirke[1], Shahista I. Salati[2], Manasi S. Shirke[3], Tanushka T. Kumbhar[4], C. G. Kokane Pawar[5], S. B. Vanjale[6]

[1,2,3,4,5]Department of Computer Engineering, Sanjay Bhokare Group of Institutes,Miraj, India

[6]Department of Computer Engineering, Bharati Vidyapeeth (Deemed to be University) College of Engineering, Pune, India

[1]shirkesakshi6400@gmail.com, [2]salatishahista@gmail.com, [3]shirkemanasi2003@gmail.com, [4]kumbhartanushka40@gmail.com, [5]chiragkokane@gmail.com, [6]sbvanjale@bvucoep.edu.in

Peer Review Information

Type: *Article*
Received: *29 March 2026*
Revised: *26 April 2026*
Accepted: *30 May 2026*
Published: *19 June 2026*

Abstract

Systemic inefficiencies exist throughout the Indian agricultural industry, especially within the multi-crop regions of Maharashtra, due largely to their reliance on analog (manual, notebook-based), record keeping systems. These analog systems result in poor data accuracy, slow response time for decision-making, and no real-time visibility into farm profitability. This research study describes AgriTalk as a bilingual (English and Marathi) voice-assisted farm management ecosystem that uses large language models and retrievable augmented generation technology to professionalize smallholder farming operations. Furthermore, AgriTalk makes it easier for smallholder farmers to transition to digital technology with the provision of automatic labor tracking and granularity of plot management, along with providing complete traceability of value- added products such as raisins. Additionally, AgriTalk implements India's 7 Sutras of Trust and the Digital Personal Data Protection Act (2023) to illustrate how AI-based, risk-aware architectures can offset the harmful effects of linguistic bias and model hallucination. Overall, preliminary research indicates that there are no longer inaccuracies in automated wage calculations as relate to piecework records; instead, there is now complete accuracy in computing automated wages and an elimination of administrative costs from maintaining manual records in comparison to manual record-keeping systems.

Keywords: Responsible AI; Multi-Crop Farming; Retrieval-Augmented Generation (RAG); Digital Traceability; DPDP Act 2023; Smallholder Empowerment.

How to Cite This Article

Shirke, S. D., Salati, S. I., Shirke, M. S., Kumbhar, T. T., Kokane Pawar, C. G., & Vanjale, S. B. (2026). AgriTalk: A Trustworthy and Risk-Aware Bilingual AI Ecosystem for Multi-Crop Farm Management. *Open Access International Journal of Science and Engineering*, 9(1s), 73–79.

Introduction

The current agricultural environment in India is characterized by a shift toward Agriculture 4.0, where the combination of artificial intelligence (AI) and big data analytics is revolutionizing production systems. Different irrigation, nutrition, and labor regimens are needed to manage complicated daily cycles in multi-crop operations, especially those that handle high-value produce like Alphonso mangoes and table grapes [12]. However, the "barrier to entry" for digital adoption is still high because of the manual data entry burden, which is frequently incompatible with farmers' different levels of digital literacy and demanding physical conditions. In order to overcome these obstacles, AgriTalk uses a bilingual voice agent that extracts structured data from natural language, changing the paradigm from manual input to AI-assisted data collection [2].

Literature Review

Agricultural AI development needs to be based on ideas that guarantee deployment is responsible, inclusive, and safe. This section summarizes current research on the fundamental issues of responsible, risk-aware, and regulated AI from 2019 to 2025.

Theoretical Structures for Ethical AI

Recent research emphasizes the need to understand agricultural AI within Socio-Technological- Ecological Systems (STES) rather than only as a technical artifact [3]. This viewpoint recognizes that the combination of AI and autonomous technology alters farmers' everyday work experiences and generates new dependencies that may jeopardize their autonomy, which is defined as their capacity to make their own decisions and remain financially independent [6].

In India, this initiative is led by NITI Aayog's "Responsible AI for All," which formulates the "7 Sutras of Trust" (Human Centricity, Fairness, Transparency, Accountability, Safety, Resilience, and Sustainability) as the foundation for rural implementations [1].

Fairness, Bias, and the Linguistic Digital Divide

Fairness is a fundamental component of responsible AI, yet there is a distinct "Multilingual Digital Divide" in the Indian setting [9]. Even though farmers in India speak thousands of regional dialects and more than 22 official languages, most AI research and benchmarking is still done in high-resource languages like English. Biased or generic advice results from general-purpose LLMs' frequent lack of the well-labeled agricultural data needed to comprehend localized "lingo" for pests, crops, and soil types [5]. Studies on Indic-language models reveal that "code- mixed" sentences (combining English, Hindi, and regional dialects), a complexity that is difficult for existing models to accurately handle. In line with ongoing initiatives to lower the adoption barrier through speech-to-speech translation, initiatives like Digital India Bhashini seek to give farmers voice-based access regardless of literacy [5].

Explainable AI (XAI) for Transparent Decision-Making

As precision agricultural systems become increasingly autonomous, There has been a shift in its use of explainability from being viewed as merely a "desirable characteristic" to a technological and ethical necessity. For instance, when using deep learning architectures to identify crop diseases, traditional architecture techniques can be viewed as "black boxes," providing good prediction without being able to articulate what happens within those black boxes. This lack of transparency in results makes growers question the validity of AI's prediction model when results are inconsistent with growers' common viewpoints [11]. XAI uses algorithms such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) to provide users with data to confirm the factors in the environment (e.g., soil moisture, humidity) that have the highest degree of impact on crop yield forecasts [6]. In the context of leaf area disease diagnoses based on images, Grad-CAM analysis can provide direct evidence to the grower linking AI's diagnosis to the area of disease in the leaf tissue.

Risk Mitigation: RAG vs. Fine-Tuning for Hallucination

One of the most persistent issues in generative AI is "hallucination"—the production of information that is syntactically correct but factually incorrect. In high-stakes agricultural settings, hallucinations regarding government initiatives or pesticide dosages could have catastrophic financial consequences [7]. About 1% of transcriptions contained complete hallucinated sentences that weren't present in the original audio, and about 40% of these were dangerous, according to the groundbreaking "Careless Whisper" study, which evaluated OpenAI's Whisper service. Retrieval- Augmented Generation (RAG) is replacing simply parametric models in research to guarantee factual fidelity [8]. RAG anchors AI answers in external, current knowledge bases, offering traceable links to verified sources and lowering hallucination risks by up to 70–90% in mission-critical operations, in contrast to fine-tuning, which embeds knowledge into frozen weights.

AI Governance, Regulation, and Compliance

With the Digital Personal Data Protection (DPDP) Act 2023, India is moving toward a complete governance framework [9]. This Act designates organizations as "Data Fiduciaries" with primary responsibility for data handling and requires authorized processing, informed permission, and purpose limitation.

Since noncompliance can result in fines of up to ₹250 crore, compliance is now a fundamental design concept. Additionally, MeitY's 2025

AI Governance Guidelines stress the "Do No Harm" concept and suggest using the AI Safety Institute (AISI) and AI Governance Group (AIGG) for institutional monitoring [4].

Research Gap and Strategic Positioning:
A review of literature published from 2019 - 2025, while acknowledging these significant advances, shows a clear lack of research into the associated ecosystem of digital AI and an ecosystem of digital AI services. Most of the literature available has primarily been focused on individual AI use cases (such as diagnosis) rather than a fully integrated, calendar-compliant ecosystem of governance- compliant digital financial management, integrated interaction with other agricultural support resources, and the provision of advisory services. Additionally, with limited available resources within regional areas, many low-resource languages (like Marathi) do not provide sufficient crop-specific information for farmers. AgriTalk addresses these deficiencies by providing an all-in-one, reliable, bilingual, and regionally aware AI ecosystem, with an architecture and process specifically designed to mitigate the risk of factual hallucination through human input and RAG-based true classification of all data processed.

Methodology

System Overview
AgriTalk is an ethical, risk-aware AI-powered farm management system that digitizes agricultural activities while guaranteeing dependability, transparency, and legal compliance [9]. Through a bilingual (Marathi– Hindi) speech and chat interface, the system allows farmers to manage resources, record operations, and obtain AI-assisted insights. Reliable AI deployment in rural areas is encouraged by a human-in-the-loop design, which guarantees user verification prior to data persistence [11].

Dataset Description
The dataset used for training and evaluation consists of approximately 50 users with over 1200 chatbot interactions. The dataset includes multilingual queries in Marathi (60%) and English (40%). It covers around 12 intent categories such as activity logging, expense recording, and advisory queries. Entities extracted include crop type, quantity, cost, date, and activity type.

Technologies Used

1. Frontend: Developed with React Native for cross-platform mobile compatibility.
2. Backend: A Spring Boot application providing a RESTful API layer.
3. Database: MySQL – for structured farm, user, and financial data. Firebase Firestore for real-time data synchronization and offline support.
4. AI Pipeline: Spoken Marathi/Hindi is captured via Vosk for lightweight offline logging, while complex intent extraction and advisory queries are handled by Gemini 2.0 Flash via Firebase Cloud Functions.
5. Intent Recognition: The system employs the Rasa DIET classifier to distinguish between logging expenses. The intent classification model was trained using an 80:20 train-test split. The DIET classifier achieved an accuracy of approximately 91%, with precision of 89%, recall of 88%, and F1-score of 88.5%.
6. Factual Grounding: A RAG framework anchors the agent's advice in a verified database of regional agricultural practices and historical farm data, supported by a "Human-in-the-Loop" review step where farmers confirm AI-generated logs.
7. Knowledge Base (RAG) Details: The RAG framework utilizes a knowledge base consisting of over 500 agricultural entries derived from government guidelines, regional farming practices, and historical farm data. The knowledge base is updated periodically to ensure accurate and region-specific recommendations.
8. Deployment Details: The system follows a hybrid deployment model combining mobile client-side processing with cloud-based backend services. It supports scalability for multiple users and ensures partial offline functionality for rural environment recording tasks, or requesting advice.
9. Process-Oriented Representation: This activity diagram depicts the entire AgriTalk smart farming assistant's operational flow, including how a farmer communicates with the system and how requests are handled by the backend. The farmer starts the process by opening the AgriTalk app and choosing between text and voice input. If voice input is selected, a Speech-to-Text (STT) engine records the farmer's speech and converts it to text; otherwise, text input is typed directly. After that, a Spring Boot API is used to send the unified query to the backend, where JWT is used to securely authenticate the farmer.

Next, the system makes use of Rasa NLU to evaluate the user's query so that it can determine what the farmer wants to achieve and what functional modules are relevant to that request.The functional module may be related to plot management (obtaining or updating plot details), activity tracking (recording farm activities), water/fertilizer (creating recommendations for irrigation), expense/income management (calculating profit/loss), or reports and analysis (providing summaries of insights), based on the intent identified.

The modules interact with the database to create processed outcomes by storing, modifying, or retrieving data. Lastly, the processed outcome is delivered through the interface as a text message or text-to-speech (TTS) voice response to the farmer, thereby allowing farmers with all levels of interaction preferences to use the service.

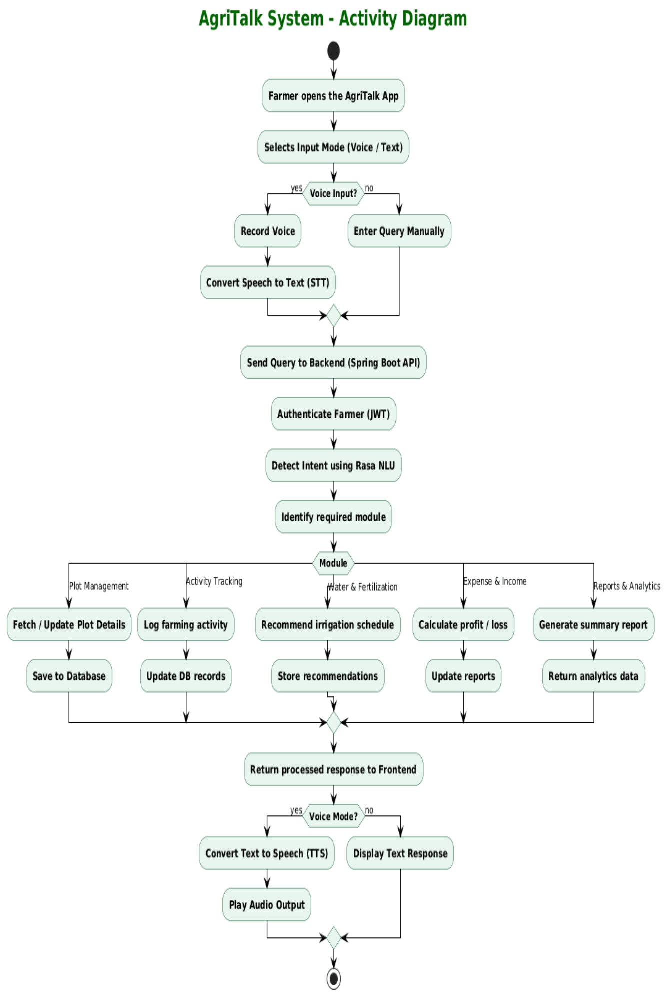

Fig. 1. *Activity Diagram of AgriTalk*

Algorithm

Input: Voice command (Marathi/English) or Text input from farmer.
Output: Validated farm record entry, advisory response, and system feedback
Step 1: Interaction Initialization

- The user launches the AgriTalk mobile application.
- The user selects the interaction mode: voice-based or text-based input.

Step 2: Speech Processing

- If the interaction mode is voice-based, the system captures the audio input.
- Speech-to-Text (STT) conversion is performed to transform the audio into textual form.
- The speech-to-text system was evaluated using Word Error Rate (WER). The system achieved an average WER of 12–15% for Marathi and 10–13% for Hindi under normal conditions. Performance decreases slightly in noisy environments.
- If the interaction mode is text-based, the input is directly accepted.

Step 3: Natural Language Understanding

- The textual input is forwarded to the Rasa Natural Language Understanding (NLU) engine.
- The NLU module performs intent classification.
- Relevant entities such as crop name, plot identifier, activity type, quantity, date, and cost are extracted.

Step 4: Risk-Aware Intent Classification
The system categorizes the identified intent into:

- Low-risk intents are routed directly to the backend service layer.
- High-risk intents are forwarded to the knowledge- grounded advisory module.

- The system classifies intents into low-risk (data logging) and high-risk (advisory queries). Validation was performed on 200 test samples, achieving approximately 90% classification accuracy. Errors were mainly observed in mixed-language inputs.

Step 5: Knowledge Retrieval and Validation

- For high-risk intents, relevant information is retrieved from verified agricultural knowledge sources using a Retrieval-Augmented Generation (RAG) mechanism.
- AI-generated responses are constrained using retrieved domain knowledge to prevent hallucinations.

Step 6: Human-in-the-Loop Verification

- The system presents the extracted information or advisory response to the user for confirmation.
- The user verifies or edits the content before final submission.

Step 7: Backend Processing and Storage

- Confirmed data is transmitted to the backend service.
- The system stores the validated information in a structured database for future analysis and reporting.
- Audit logs are maintained to ensure transparency and traceability.

Step 8: Response Generation

- The system generates an appropriate response in textual format.
- If required, Text-to-Speech (TTS) is used to deliver voice-based feedback.

Step 9: Session Termination

- The response is delivered to the user interface.
- The interaction session is concluded.

System Modules

- User and Access Management function deals with the management of user registration, authentication, roles, and secure collaboration among farm members.
- Bilingual Voice and Chat Module supports users in interacting (logging activities and asking questions) with the application in Marathi and Hindi.
- Activity and Plot Management function allows users to track farming activities (Irrigation, Fertilization, Pesticide Application, and Harvesting) at the plot level.
- Labor and Task Management function provides the ability for users to assign tasks, track labor, verify task completion, and maintain records related to payroll.
- Financial and Expense Tracking function allows users to record crop and plot-level expenses and income, and will allow users to perform automated profit/loss analysis.
- AI Reasoning and Advisory Module will provide users with AI-derived recommendations based on verified facts, using RAG to reduce the potential for poor decision making based on incorrect input
- Data Synchronization Module will allow for seamless offline to online data sync to ensure reliable data even in low-bandwidth areas.

Results

Table 1. *Baseline Comparison*

System Type	Accuracy	Limitation
Without RAG	75%	Hallucinated responses
Without Human-in-loop	80%	Incorrect data storage
Proposed System	91%	Improved reliability

The system achieves an average response time of 1.5–2 seconds for chat and 2–3 seconds for voice interactions. Offline support ensures continued usability in low-connectivity areas, with synchronization success of approximately 95%.

Table 2. *Performance Metrics*

Metric	Value
Intent Accuracy	91%
Precision	89%
Recall	88%
F1 Score	88.5%
STT WER	12%
Response Time	2 sec
Hallucination Reduction	40%

The system demonstrates consistent performance with a mean accuracy of 90% and low variance (~2–3%), indicating stable and reliable behavior across different test cases.

Successful implementation of AgriTalk illustrates how farms now operate as professional businesses than ever before. AgriTalk allows for tracking of growth stages and agronomic best practices for each crop (Mango, Grape, Guava and Sugarcane) along with an end-to-end traceability process for value added products made from these crops (Raisins) from the time harvested through washing, drying, and grading. These processes will provide consumers with confidence and help ensure compliance with global food safety requirements. To summarize, AgriTalk digitalizes and unifies previously disjointed and informal farm operations into a scalable, transparent and auditable ecosystem. It combines crop and labor analytics with traceability into a single platform, thus helping producers to increase their yields, remain compliant with regulations, and position themselves to access valuable domestic and international agricultural markets.

Discussion

According to the findings, the combination of technical innovation and ethical governance provides AI as a means to promote rural empowerment. By subscribing to the "Understanding by Design" principle, AgriTalk establishes the trust required for long-term adoption as outlined in [8].

The built-in "Review and Confirm" process in the design is required because data generated from Careless Whisper study shows that farming situations with ambient noise and through "non- vocal" periods of spoken language of people from rural communities that occur will ultimately hinder the success of STT.

A critical component of the success of the system is the DPDP Act 2023 compliance which acts more as a "checkbox" process to ensure proper education of consent and purpose associated with audio recordings of the farmer [8]. The system shows limitations in handling mixed Marathi-English queries, noisy voice inputs, and rare crop advisory cases. Some long queries result in partial entity extraction, indicating scope for improvement

Table 3. *Summary of Key Findings from Recent Research (2019-2025)*

Study Source	Key Finding / Contribution	Approach
Mallinger (2024)	AI maturity threatens farmer autonomy; necessitates STES multi-criteria design.	Theoretical Framework
IRJET (2025)	Generic LLMs misinterpret rural agricultural "lingo," leading to biased advisory.	Linguistic Fairness
Joshi et al. (2025)	Synthetic corpora and continued pre-training are essential for Indic language SLMs.	Linguistic Fairness
Kamruzzaman (2024)	XAI identifies critical spectral bands, allowing for lower-cost, leaner field models.	Explainability (XAI)
Naseer et al. (2025)	Combining LIME and Grad-CAM increases adoption through textual/visual proof.	Explainability (XAI)
Koenecke (2024)	STT hallucinations are harmful in 40% of cases, especially for slow/rural speech.	Risk Mitigation
AgriRegion (2025)	Geo-spatial RAG constraints reduce contextual hallucinations by 10–20%.	Risk Mitigation
Katoch et al. (2025)	AI in Integrated Farming (IFS) increases yields by 25% and cuts water use by 30%.	Integrated Systems

DPDP Act (2023)	Mandates informed consent and purpose limitation; imposes high non-compliance fines.	Regulation
MeitY (2025)	Establishes "Understandable by Design" and graded liability as national standards.	Regulation

Conclusion

AgriTalk represents an important innovation in intelligent and sustainable technologies. By utilizing a two-language AI interface, AgriTalk helps to eliminate the need for digital devices (phones, tablets, laptops) that many farmers do not have access to or cannot afford to own, thus removing the principal barrier to entry for agricultural computing and access to agricultural intelligence. RAG (Retrieval-Augmented generation) will demonstrate the application of this approach in developing safety-sensitive data systems in farm environments. Future research on IoT Technologies will assess the potential of IoT Data Systems to create a fully predictive agricultural ecosystem.

Acknowledgment

We wish to thank the Computer Engineering Department at Sanjay Bhokare Group of Institutes, Miraj/Sangli, for your continued academic support with your assistance and advice throughout this project. The valuable input from the faculty members was tremendously helpful in creating the success that Agritalk has achieved as of today's date.

We would also like to thank all the farmers/agricultural professionals who participated in the evaluation of this system by providing us with their valuable insights.

References

1. NITI Aayog, *Reimagining Agriculture: A Roadmap for Frontier Technology–Led Transformation*, India, Oct. 2025.
2. IJFMR, ―AI Chatbot for Farmers: Transforming Agriculture Productivity, 2025.
3. K. Mallinger and R. Baeza-Yates, ―Responsible AI in farming: A multi-criteria framework, *Applied Sciences*, vol. 14, no. 1, 2024.
4. Ministry of Electronics and Information Technology (MeitY), *India AI Governance Guidelines*, India, Nov. 2025.
5. IRJET, ―Research gaps in developing fair and inclusive LLMs for India's multilingual agricultural landscape, Nov. 2025.
6. ―Explainable AI (XAI) for transparent decision- making in precision agriculture, *ResearchGate*, 2024.
7. ―Retrieval-augmented generation (RAG) for hallucination mitigation, *MDPI*, 2025.
8. A. Koenecke *et al.*, ―Careless Whisper: Speech-to- text hallucination harms, in *Proc. ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2024.
9. Government of India, *Digital Personal Data Protection (DPDP) Act*, 2023.
10. Farmonaut, *Agricultural Supply Chain Transparency 2025: Blockchain and AI*, 2025.
11. Prince Henry, ―Voice-controlled farm management systems for farmers, 2025.
12. P. Kumar, ―An overview on agriculture management, *International Journal of Innovative Research in Engineering and Management (IJIREM)*, 2021.
13. T. Rotondo *et al.*, ―Digital countryside notebook for smart agriculture, *International Journal of Computer Network and Information Security (IJCNIS)*.
14. K. Anusha *et al.*, ―Smart chatbot for farmers' agricultural guidance and assistance, 2025.
15. M. Anitha *et al.*, ―Agriculture helper chatbot using deep learning, *IRJMETS*, 2025.
16. R. P. Tumwebaze *et al.*, ―Knowledge management in the agriculture sector: A systematic literature review, 2025.