Optimizing Agricultural Advisory Services With Multilingual LlaMA And Web Automation

Bharathi MohanG[1], Jayanth Adhitya C M[2], Mithilesh A[3]
[1, 2, 3] Dept of Computer Science and Engineering
[1, 2, 3] Amrita School of Computing
Amrita Vishwa Vidyapeetham Chennai, India

Abstract- *Agriculture is key to global food security, and hence it is imperative to create new solutions that cater to the knowledge deficits of most farmers—especially in developing countries where specialist guidance is frequently not available. In these regions, farmers normally turn to helplines for crucial assistance, but high charges and limited access present major challenges. Automating responses to agricultural questions can alleviate the burden on conventional helpline systems, enabling farmers with timely and accurate data. In addition, integrating real-time weather forecasts and crop disease forecasts into these systems supports farmers in decision-making through proactive, locally-based insights to help mitigate risks and increase productivity. This provides farmers with not just general advice but also context- based recommendations that are specific to their individual environments. Additionally, incorporating artificial intelligence in agriculture offers a bright future avenue, with safety net words—particularly transformers—emerging highly proficient at interpreting complex questions on agriculture and providing appropriate answers. The present paper delves into how large language models (LLMs) are able to ease access to knowledge for farmers by taking advantage of their vast language processing ability. With a rich pool of more than four million queries from Tamil Nadu, India, over a broad topic area in agriculture, this research demonstrates how LLMs are effective at bridging the knowledge gap and giving farmers real-time access to crucial information.*

Keywords- Agriculture, Artificial Intelligence (AI), Large Language Models (LLMs), Transformers, Farmer Inquiries, Natural Language Processing (NLP), Precision Agriculture, Agricultural Information Systems

I. INTRODUCTION

Agriculture is not just an economic pillar but also a key determinant of social and environmental stability in most nations especially in developing nations where farm livelihoods are largely supported by agriculture. Irrespective of its crucial significance, the agricultural sector is subjected to many challenges, such as erratic climatic conditions, unstable market prices, and poor access to timely and pertinent information. Although these challenges impact farmers globally, they are even more acute in nations with underdeveloped infrastructure and limited access to technology. Classic information dissemination tools—agricultural extension services, helplines, and expert consults—are advantageous but constrained by inefficiencies, scalability, and logistics issues. Millions of farmers end up being underserved, depriving them of the critical resources necessary for making well-informed decisions that would strongly increase productivity and sustainability

Large language models (LLMs) and, specifically, transformer- based LLMs are dramatically changing natural language processing and artificial intelligence. Models such as OpenAI's GPT-4, Meta's LLaMA, and Google's Gemini demonstrate strong capabilities in understanding and generating human language across diverse fields, making them well-suited for automating responses to agricultural inquiries. Their application in agriculture has the potential to close existing knowledge gaps by providing farmers with accurate and contextually relevant information. Unlike earlier approaches that relied on keyword matching or basic rule-based systems, LLMs can grasp the subtleties of natural language, enabling more personalized and subtle interactions. However, despite their potential, applications of LLMs in agriculture are still fairly under-explored, with the majority of existing research activities focused in domains such as customer service, healthcare, and education.

This research aims to fill this gap by using a multilingual version of LLaMA, coupled with a Retrieval-Augmented Generation (RAG) system, to mechanize the answering of farmer questions. Focused on a corpus of more than four million agricultural queries from Tamil Nadu, India, the study delves into how LLMs have the potential to transform agricultural advisory services. By integrating a document retrieval process with language model-based generation, the system provides responses that are not only well-coherent but also based on accurate and updated information. This hybrid method successfully addresses one of the key weaknesses of traditional LLMs—hallucination, where models generate coherent- sounding but false or unrelated content answers.

Additionally, adding multilingual support to the LLM is critical in a linguistically diverse nation like India, where farmers speak a vast number of regional languages. This ensures that the system is made available to a wider population, most of whom might not be English or Hindi speaking—the two most widely used languages in Indian government and advisory services. The model's capacity for responding in Tamil, the local language of our farmers, is one key factor that makes it effective, bridging the language barrier that has normally been a communication bottleneck between farmers and agricultural specialists.

For further enhancement of the system's effectiveness, the model supports real-time weather patterns and market prices to allow farmers to make more informed choices. This feature aids them in reacting to abrupt climatic changes and strategizing crop sales in a better way, thus improving their financial stability. By using external APIs and updating the dataset from time to time, the system continues to be adaptive and responsive to evolving farm conditions. The platform also provides an interactive voice-enabled interface for farmers with poor literacy levels. This enables users to communicate with the system using verbal queries and respond with their native language. By lowering the barrier to access, the system enables a larger number of farmers, driving higher adoption and more sustainable agriculture.

II. LITERATURE REVIEW

Anushka Marla et al. studied the creation of a rural chatbot by the name of AgroBot, which utilizes Natural Language Processing (NLP) to handle text-based queries and utilizes Convolutional Neural Networks (CNNs) to detect plant dis- eases based on leaf images. AgroBot provides specific guidance to farmers, hoping to curb losses by 9–21varied linguistic backgrounds. The authors propose future developments, such as increasing the knowledge base to encompass a broader range of crops and pests, and enhancing real-time processing capabilities to encourage wider adoption in rural communities

Samarth Godara et al. developed AgriResponse, a real- time framework designed to generate responses to agricultural queries by building a knowledge base from helpline call logs. The system features three Response Retrieval Models (RRMs), which use input parameters such as state, crop type, and disease or pest names to retrieve relevant answers. AgriResponse achieved a high accuracy rate of 98.5%, significantly improving both the speed and reliability of support for farmers. However, the system is limited by a smaller dataset for less common crops and faces a trade-off between response time and accuracy, with the most accurate model also being the slowest.

Mohammad Zia Ur Rehman et al. introduced KisanQRS, a deep learning-based system designed to automate the query- response process in agriculture. The framework clusters farmer queries based on semantic and lexical similarities, using a trained module to map new queries to these clusters and rank responses by relevance. KisanQRS demonstrates strong performance in query clustering and mapping, enhancing both the accuracy and efficiency of retrieving relevant answers. However, its accuracy can be affected when dealing with queries that have ambiguous semantic meaning. Additionally, the system provides real-time information—such as market prices and weather updates—to support farmers in making informed decisions.

Krish Didwania et al. introduced AgriLLM, a framework that leverages Large Language Models (LLMs) to process and respond to approximately four million real-world farmer queries from Tamil Nadu. AgriLLM demonstrated high ac- curacy in addressing topics such as fertilizer application and nutrient management, highlighting its potential for automating agricultural information dissemination. The model performed exceptionally well in generating precise responses to core agricultural needs but showed comparatively lower effectiveness when handling queries related to government schemes, indicating areas for improvement in the model's comprehension and generalization capabilities.

Abhinand Balachandran enhanced the open-source LLaMA model by integrating 16,000 Tamil tokens and fine-tuning it using Tamil-translated Alpaca and OpenOrca datasets, resulting in the development of Tamil-LLaMA. This model significantly improves Tamil language generation and comprehension, offering more accurate and contextually relevant outputs for Tamil-speaking users. However, Tamil-LLaMA still exhibits limitations in cultural and contextual understanding and lacks detoxification mechanisms, which can result in harmful con- tent or unexpected responses.

Gautam B et al. conducted a comparative study on the performance and scalability of open-source Large Language Models (LLMs) within enterprise-specific Retrieval-Augmented Generation (RAG) systems. The study focused on how these models, when integrated with effective embedding techniques, can enhance accuracy and efficiency for enterprise tasks in- volving specialized datasets. While the findings highlight the potential of LLMs to improve RAG-based systems, the study also points out the limited exploration of long-term performance and adaptability across diverse enterprise scenarios.

The development of interactive systems aimed at improving access to government schemes has emerged as a growing area of research, particularly for monolingual individuals in underserved communities. Previous studies have explored the application of natural language processing (NLP) and chatbot frameworks to simplify complex bureaucratic processes, with platforms like RASA proving effective in managing user queries. Intent classification using machine learning models such as RNNs and SVMs has shown promising results, of- ten achieving accuracy rates above 90%, demonstrating their reliability in understanding user intent. Knowledge graphs have also been widely employed to enhance information retrieval, offering more precise and context-aware recommendations. However, most existing systems have focused on widely spoken languages like Hindi and English, with limited resources available for regional languages such as Kannada. This need highlights the significance of efforts such as Dhvani, which leverage cutting-edge technologies to bridge language gaps. By integrating voice-based interaction, autofill of forms, and high-accuracy scheme suggestions, such systems are a major leap towards empowering disadvantaged communities—keeping pace with the overall trend of AI-based solutions for social good.

III. **BACKGROUND STUDY**

Need for LLMs in Agriculture

Agriculture plays a critical role in ensuring economic stability and food security, especially in developing nations. However, farmers in regions with limited infrastructure often struggle to access timely and reliable information, hindering effective decision-making. Traditional advisory methods are frequently expensive and inefficient, leaving many without actionable support. Large Language Models (LLMs) such as OpenAI's GPT, Meta's LLaMA, and Google's Gemini offer promising solutions by delivering scalable, real-time, and cost-effective assistance. Their ability to comprehend and generate multilingual answers makes them especially suited for linguistically diverse nations such as India.

While their potential is great, Large Language Model (LLM) adoption in agriculture is currently limited. The data that accompanies it indicates the prevalence of different LLM architectures, with the GPT and LLaMA families of models demonstrating consistent growth. This work aims to fill the adoption gap by applying a Retrieval-Augmented Generation (RAG) approach based on a fine-tuned, multilingual LLaMA model. Tailored specifically for the requirements of Indian farmers, the system draws upon a database of more than four million farmer queries to produce context-sensitive, accurate replies in regional languages like Tamil. By aligning with the expanding LLM landscape, this approach enhances accessibility and relevance in regions like Tamil Nadu, ultimately empowering farmers with actionable insights and promoting sustainable agricultural practices.

The upward trajectory of LLM innovation, as reflected in the distribution data, presents a timely opportunity for the agricultural sector to harness AI-driven advancements. Through ongoing fine-tuning and contextual adaptation, LLM-powered systems have the potential to transform how farmers access essential information, promoting greater resilience, efficiency, and productivity within rural communities.

Role of RAG in LLMs

Retrieval-Augmented Generation (RAG) frameworks enhance the capabilities of Large Language Models (LLMs) by addressing common issues such as hallucination and the lack of real-time information. By integrating information retrieval with generative responses, RAG ensures outputs are grounded in factual, up-to-date data, thereby improving accuracy and relevance. This is especially crucial in agriculture, where timely and precise guidance on topics like pest control, crop management, and market trends can significantly impact out- comes for farmers. RAG frameworks retrieve domain-specific content and seamlessly integrate it into the generative process of LLMs, enhancing both accuracy and contextual understanding. In multilingual environments like India, this approach ensures linguistic and cultural relevance, allowing farmers to access timely and localized information in their native languages. By supporting dynamic content—such as weather updates, pest outbreaks, and policy changes—RAG-powered LLMs empower farmers to make informed decisions, promote sustainable agricultural practices, and reduce dependence on traditional advisory systems.

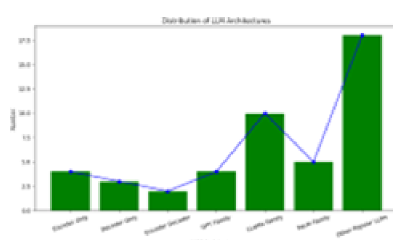

Fig. 1. LLM in Agriculture

Automated Government Form Filling for Aged and Mono-lingual People Using Interactive Tool

Accessing government schemes and completing application forms remains a major challenge for elderly and monolingual individuals, particularly in rural areas. To address this, we propose an interactive, voice-enabled tool integrated with an LLM-powered system that guides users through the process of applying for government services. Utilizing speech-to-text and text-to-speech technologies, the tool communicates in users' native languages, gathers required information, and automatically populates forms in English or the appropriate format. By leveraging the RAG framework, the system provides accurate and contextually relevant recommendations for suitable schemes based on user input. This functionality is particularly useful for those with low literacy or computer skills, so that they can access and use important services on their own and circumvent the bureaucratic obstacles otherwise in their way.

IV. METHODOLOGY

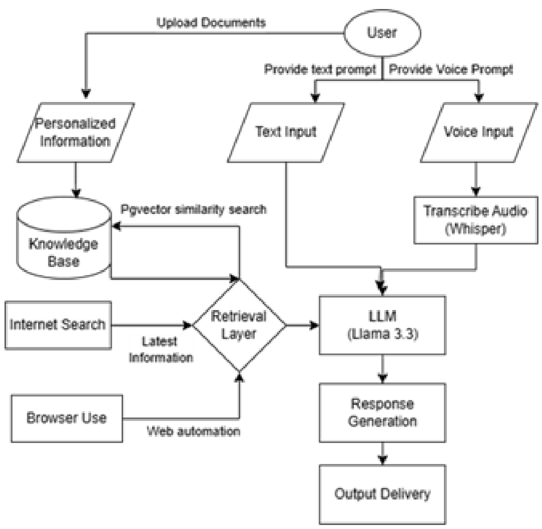

Fig. 2. Model Architecture

A. Dataset and pre-processing

The dataset for this project consists of an extensive set of actual farmer questions from across India, collected through various farm advisory services. Covering the period 2006- 2023, it contains more than 32 million questions on a wide variety of subjects like crop care, pest management, irrigation methods, and weather conditions. These questions, resolved by respective domain experts, are augmented with vital metadata like farmer location, crop, and environmental conditions. The dataset covers several regions and languages, and it is thus typical of India's heterogeneous agricultural geography. This diversity allows the AI model to effectively cater to farmers with varied linguistic, climatic, and agronomic settings. To clean the data to train the multilingual farm advisory system, preprocessing tasks like cleaning, language tagging, and alignment were performed to provide high-quality data as well as easy usability.

The preprocessing stage included several necessary steps to cleanse the dataset so that it would be ready for efficient training. Firstly, data cleaning was done to eliminate incomplete, duplicate, or irrelevant entries. Second, language labeling was used to determine the regional language of every query, allowing the system to process multilingual inputs correctly. Third, noise filtering was done to remove redundant or low-quality data. The dataset was also balanced to provide fair representation of various query types and regions. Lastly, text normalization techniques—such as tokenization and stop-word removal—were used to normalize inputs for effective processing by the AI model. These preprocessing phases were most important in preparation of the dataset for developing an effective and scalable advisory system that could provide accurate, Realtime assistance to Indian farmers

B. *RAG Framework Components*

To make use of a Retrieval-Augmented Generation (RAG) setup effectively, we use a suite of sophisticated tools and elements. These tools allow the system to search for real-time, context-sensitive information, which, when combined with a large language model (LLM), makes it possible to provide precise and elaborate answers to farmers' questions. The following elements are responsible for upgrading the system's ability:

1. *Knowledge Base Search Tool with pgvector:*

The Knowledge Base Search Tool makes use of the pgvector extension for efficient vector-based searches in a pre-indexed database. It fetches relevant entries from user queries, which are then incorporated into the LLM's response generation. This guarantees that answers are correct and based on solid, domain- specific content. The knowledge base covers a broad array of agricultural subjects, such as government schemes, best practices in farming, pest and disease control, and market price trends. This module greatly minimizes the possibility of hallucinations, enhancing the overall reliability of the system

2. *Browser Tool with Selenium and Playwright:*

The Browser Tool makes it possible for the system to dynamically interact with live web pages. Utilizing Selenium and Playwright, the tool can browse a browser, index HTML components, and send commands to them programmatically. It locates page components such as forms, buttons, and data fields utilizing bounding box algorithms, enabling organized navigation. It is especially valuable in fetching actual-time data which is not within the internal knowledge base, i.e., current market prices, weather forecasts or agricultural policy changes. The LLM is capable of executing such tasks as scrolling, clicking, and text typing on its own, thus having the ability to access information or complete forms for the user. This can greatly improve convenience for farmers who have low levels of digital literacy.

By integrating OmniParser, the Browser Tool now also leverages advanced screenshot parsing to further automate navigation. OmniParser applies state-of-the-art computer vision techniques to analyze screenshots of the live browser interface. It extracts structured information—such as bounding boxes and semantic labels for interactive elements—directly from the visual content. This parsed data is then used by the LLM to determine the exact location and function of UI elements, improving the accuracy of navigation commands (e.g., clicking the correct button or filling in forms). In essence, OmniParser bridges the gap between raw visual data and actionable instructions, enabling the system to adapt to dynamic web layouts and automate complex interactions with greater precision.

3. *DuckDuckGo Search Integration Tool:*

The DuckDuckGo Search Integration Tool augments the system's capacity to acquire current, external knowledge in the form of real- time web searches. The tool helps the LLM access appropriate and timely agricultural data, such as recent outbreaks of pests, weather deviations, and market movements. The tool is especially useful for time-critical queries to ensure that farmers access the latest information available. Additionally, DuckDuckGo's privacy-protecting capabilities offer an additional layer of protection when handling possibly sensitive information.

C. Fine-Tuning LLaMA

Fine-tuning the LLaMA model for agricultural use requires tuning the pre-trained model to improve its ability to comprehend and generate domain-specific content. The process starts with compiling a heterogeneous dataset that covers different facets of agriculture, including crop management practices, pest control methods, irrigation methods, and government policies. The dataset is annotated thoroughly to create a domain-specific vocabulary that well represents agricultural jargon and contextual subtleties.

In fine-tuning, the pre-trained LLaMA model is trained further with transfer learning techniques. Critical hyperparameters—learning rate, batch size, and sequence length—are tuned to maximize performance for the agricultural context. This customized training improves the model's comprehension of the linguistic and contextual nuances characteristic of farmer questions.

With the addition of agriculture domain knowledge, the fine- tuned model greatly enhances its value as a decision-support tool for farmers. It facilitates the production of more accurate, relevant, and actionable answers, thus encouraging informed decision making and sustainable farming practices.

D. System Architecture: Retrieval-Augmented Generation (RAG) Framework

The RAG framework architecture is organized to harmoniously provide multiple components that facilitate timely and correct responses to agricultural questions. As the backbone of the architecture, a Large Language Model (LLM) is responsible for creating coherent and context-specific responses.

The process starts with the **User Input Interface**, where farmers can input their questions through an intuitive mobile or web app. When an input is received, the system executes **Query Processing** through Natural Language Processing (NLP) methods to recognize user purpose and identify major entities like crop type, location, or issue

Following this, the **Retrieval Layer** is activated, comprising three essential components:

Knowledge Base Search Tool (pgvector): Enables semantic similarity-based retrieval from a vectorized, domain-specific knowledge base.
Browser Tool (Selenium and Playwright): Facilitates real-time interaction with external web pages to extract live information not present in the knowledge base.
DuckDuckGo Search Integration Tool: Conducts real time, privacy-preserving web searches to access up-to- date agricultural data, including weather forecasts, pest outbreaks, and market trends.

These retrieval tools work collaboratively to collect relevant and updated content corresponding to the user's query. Next, the LLM synthesizes and contextualizes the retrieved information. By combining static knowledge with real-time web content, the model generates responses that are both accurate and grounded in the latest data. Its advanced language understanding capabilities allow it to tailor advice to the farmer's specific context, whether related to pest management strategies, crop rotation practices, or government schemes.

The output thus generated is submitted to the Output Delivery Module, which formats it in an understandable and readable form to be delivered to the user. For users who are low-literacy, the module can consist of a voice-output facility. The system further has a User Feedback Mechanism where farmers can provide a rating of how accurate, helpful, and pertinent the responses provided are. Such feedback is used to fine-tune the LLM iteratively so that the system to learn from user interactions and become better over time. Through this strong architecture, the RAG-based framework provides farmers with timely, accurate, and personalized information to aid important agricultural decision-making.

E. Evaluation Metrics

In evaluating the performance of the models used in our agricultural advisory system, we applied a number of important evaluation metrics that are well-known in natural language processing. These metrics are BLEU, ROUGE, Precision, Recall, F1 score, and RAGAS. Each of these metrics sheds important light on various dimensions of the models' text generation abilities.

1. BLEU Score: BLEU (Bilingual Evaluation Understudy) score is a frequently used measure of the quality of text output generated by a model with respect to one or multiple reference texts. It measures to what extent the generated text matches the reference through examination of n-grams (sets of n consecutive words). The BLEU score is calculated based on the following formula:

$$\mathrm{BLEU} = \exp\left(\min\left(1, \frac{\text{length of generated text}}{\text{length of reference text}}\right)\right) \cdot \prod_{n=1}^{N} p_n^{w_n} \tag{1}$$

where p_n is the precision of n-grams, and w_n is a weight for each n-gram.

2) ROUGE Score: The ROUGE (Recall-Oriented Understudy for Gisting Evaluation) score is used to assess the quality of generated summaries by comparing them against reference summaries. It emphasizes recall by measuring the overlap of n-grams between the generated output and the reference texts. The ROUGE-N score is calculated as:

$$\text{ROUGE-N} = \frac{\sum_{n\text{-grams} \in \text{Recall}} \text{Count}_{\text{match}}(n\text{-gram})}{\sum_{n\text{-grams} \in \text{Reference}} \text{Count}(n\text{-gram})} \tag{2}$$

A higher ROUGE score signifies better performance in retrieving relevant information.

3. RAGAS Score: The RAGAS (Retrieval-Augmented Generation Assessment Score) is a tailored metric developed to evaluate the effectiveness of systems using Retrieval-Augmented Generation methods. It measures how accurately and coherently the system incorporates retrieved knowledge base context into its generated responses. The RAGAS score is calculated as:

$$\text{RAGAS} = \alpha \cdot \text{Relevance} + \beta \cdot \text{Fluency} + \gamma \cdot \text{Diversity} \tag{3}$$

where α, β, and γ are weights assigned to each component, and the components are defined as follows:

Relevance (α): Measures how closely the generated response aligns with the user's query and the retrieved context.
Fluency (β): Assesses the grammatical correctness and readability of the generated text.
Diversity (γ): Evaluates the variability of responses to similar queries, ensuring a range of relevant outputs.

Together, these metrics provide a comprehensive evaluation framework for assessing the performance of our models, ensuring that the responses generated by the agricultural advisory system are accurate, relevant, and valuable to farmers seeking assistance.

TABLE II
RAGAS Results Comparison

Query	GroundTruth	RAGOutput	Context Precision	Faithfulness	Answer Relevancy	Context Recall
Priceofureain India?	TheMRPofureaisstatutorilyfixedbythe Government of India and at present it is Rs.268 for a 50 Kg bag of urea/ Rs. 242 for a45 kg bag of urea.	ThecurrentpriceofUreainIndiaisa round 242per45KGbagand268per50 KGbag,as mentioned on the official website of theDepartment of Fertilizers.	0.95	1.0	0.98	0.90
Priceofonionin TamilNadu?	Mettupalayam(uzhavarSandhai)-Rs62; Mettur(uzhavarSandhai)-Rs60	ThecurrentpriceofonioninTamil Nadu todayisbetween53to76perkilogra m.	0.85	0.7	0.75	0.60

V. RESULTS AND DISCUSSION

A. Quantitative Analysis

The quantitative evaluation of the models is summarized in the table. Several performance metrics—including BLEU, ROUGE, Precision, Recall, and F1 score—were used to compare their effectiveness. Among the models, LLaMA demonstrates the strongest performance across all metrics, achieving a BLEU score of 0.560, ROUGE of 0.730, Precision of 0.850, Recall of 0.855, and an F1 score of 0.850. These results indicate that LLaMA is highly effective in generating responses that closely align with the reference text, both syntactically and semantically. In comparison, models such as Flan-T5-base and T5-base show competitive but slightly lower performance. Flan-T5- base achieves a BLEU score of 0.555, ROUGE of 0.724, Precision of 0.834, Recall of 0.846, and an F1 score of 0.837. While these models perform reasonably well, they lag behind LLaMA in overall text generation quality. Other models, including Pegasus-xsum and MT5-small, display even lower Precision and Recall, reflecting a reduced ability to capture and convey relevant information within this task.

Overall, the superior performance of LLaMA across all evaluation metrics underscores its suitability for the agricultural advisory system, ensuring that generated responses are both accurate and contextually appropriate for farmers' needs.

TABLE I: Comparison of models based on BLEU, ROUGE, Precision, Recall, and F1 scores.

Model	BLEU	ROUGE	Precision	Recall	F1
t5-small[23]	0.521	0.701	0.823	0.833	0.825
bart-base[22]	0.550	0.720	0.836	0.837	0.835
t5-base[23]	0.547	0.718	0.831	0.843	0.834
flan-t5-base[30]	0.555	0.724	0.834	0.846	0.837
flan-t5-small[30]	0.530	0.708	0.824	0.836	0.828
llama	**0.566**	**0.732**	**0.854**	**0.852**	**0.851**
mt5-small[26]	0.542	0.715	0.830	0.840	0.833
pegasus-xsum[24]	0.552	0.717	0.827	0.844	0.833

B. Qualitative Analysis

Qualitative analysis further underscores the performance ad- vantages of the LLaMA model. Its fine-tuning on agriculture- specific datasets enables it to produce responses that are both highly relevant and contextually accurate to farmer queries. The high ROUGE score of 0.730 reflects the model's strength in capturing essential, context-rich information critical for effective agricultural guidance. Similarly, the BLEU score of 0.560 indicates that LLaMA's responses are well-structured and linguistically coherent, ensuring the generated content is accessible and easy to comprehend for the intended audience.

When compared to models like Flan-T5-base and T5-base, LLaMA shows superior understanding of agricultural terminology and domain-specific contexts, resulting in greater precision and relevance in its answers. On the other hand, models such as MT5-small and Pegasus-xsum, while capable of generating fluent text, lack the domain adaptation necessary for agriculture-focused applications. This often leads to generic or less useful responses that do not adequately address the nuanced needs of farmers.

In summary, the alignment of both quantitative metrics and qualitative assessment clearly positions LLaMA as the most effective model for enhancing agricultural advisory services. Its ability to deliver accurate, relevant, and context-aware responses makes it the optimal choice for meeting the information needs of farmers in real-world scenarios.

TABLE III
RAGAS Scores for Different Models

Model	Relevance	Fluency	Diversity	Context Recall	RAGAS Score
LLaMA	0.90	0.95	0.80	0.87	0.85
Flan-T5-base	0.75	0.88	0.70	0.82	0.77
T5-base	0.73	0.85	0.68	0.80	0.74
MT5-small	0.68	0.80	0.65	0.75	0.72
Pegasus-xsum	0.65	0.75	0.60	0.70	0.67

VI. CONCLUSION

In this work, we leverage advanced Large Language Models (LLMs) for sequence-to-sequence text generation to effectively respond to farmer queries using the comprehensive Kisan Call Center (KCC) dataset, which contains real-world interactions between farmers and agricultural experts. Our approach in- volved meticulous data cleaning and preprocessing to address noise and inconsistencies in the dataset, marking one of the first implementations of LLMs in natural language generation specifically tailored for agricultural advisory services. The fine-tuning of the LLaMA model on this domain-specific corpus demonstrates strong performance in generating high- quality, context-aware responses across a diverse range of agricultural queries.

By integrating multilingual capabilities and harnessing powerful models like LLaMA, our system is capable of producing responses that are not only accurate but also linguistically tailored to a diverse audience—extending beyond local dialects to effectively engage a broader and more educated demographic. This approach significantly enhances the overall utility of the system, reducing the reliance on traditional call centers and agricultural experts while ensuring that farmers receive timely, precise, and contextually appropriate information.

Ultimately, this work paves the way for more scalable and inclusive agricultural advisory services, contributing to the development of a sustainable and technologically empowered agricultural ecosystem.

In future research, we aim to further enhance the system by incorporating metadata associated with each farmer query, enabling more precise and personalized model training. Additionally, we plan to expand the system's reach to encompass more regions across India by leveraging LLaMA's multilingual capabilities to provide localized support in various regional languages. This will allow the system to better address the unique challenges faced by farmers in different parts of the country, fostering a more responsive and adaptable agricultural support network.

REFERENCES

[1] S. Jaiswal, T. Kharade, N. Kotambe, and S. Shinde,” Collaborative Recommendation System for Agriculture Sector,” *ITM Web Conf.*, vol. 32, no. 03034, 2020. DOI: 10.1051/itmconf/20203203034.

[2] S. Kulkarni, S. Madge, T. Madhave, and S. P. Kosbatwar, ”Automatic Document Verification and Government Policy Recommendation System,” *EasyChair Preprint*, no. 5612, pp. 1-4, May 2021.

[3] T. Ahmad and R. Haneef, ”Government Agricultural Schemes in India: A review,” November 2019.

[4] A. Marla, R. Paul, A. K. Saha, N. K. Basha, and B. Anandhakrishnan, ”An AgroBot: Natural Language Processing Based Chatbot for Farmers,” *2023 4th International Conference on Smart Electronics and Communication (ICOSEC)*, Trichy, India, 2023, pp. 1235-1241.

[5] S. Godara, J. Bedi, R. Parsad, D. Singh, R. S. Bana, and S. Marwaha, ”AgriResponse: A Real-Time Agricul- tural Query-Response Generation System for Assisting Nationwide Farmers,” *IEEE Access*, vol. 12, pp. 294-311, 2024.

[6] Z. Ur Rehman, M. Raghuvanshi, and N. Kumar, ”KisanQRS: A Deep Learning-based Automated Query-Response System for Agricultural Decision-Making,” *Computers and Electronics in Agriculture*, vol. 213, 2023.

[7] G. Kalaiselvi, B. Rubavarshini, T. Thirumullai, and Dhivyadharshini, ”FarmTechBot (FTB): Empowering farmers for better harvest,” *International Journal of Science and Research Archive*, vol. 12, pp. 055-062, 2024.

[8] C. Dhavale, T. Pawar, A. Singh, S. Pole, and K. Sabat, ”Revolutionizing Farming: GAN-Enhanced Imag- ing, CNN Disease Detection, and LLM Farmer Assis- tant,” *2024 2nd International Conference on Computer, Communication and Control (IC4)*, Indore, India, 2024, pp. 1-6.

[9] Didwania, P. Seth, A. Kasliwal, and A. Agar- wal, ”AgriLLM: Harnessing Transformers for Farmer Queries,” arXiv preprint, arXiv:2407.04721, 2024.

[10] Z. Yang, Y. Yang, C. Zhao, Q. Guo, W. He, and W. Ji, ”PerLLM: Personalized Inference Scheduling with Edge-Cloud Collaboration for Diverse LLM Services,” arXiv preprint, arXiv:2405.14636, 2024.

[11] D. McDonald, R. Papadopoulos, and L. Benningfield, ”Reducing LLM Hallucination Using Knowledge Distillation: A Case Study with Mistral Large and MMLU Benchmark,” TechRxiv, May 25, 2024.

[12] H. Inan, K. Upasani, J. Chi, R. Rungta, K. Iyer, Y. Mao, M. Tontchev, Q. Hu, B. Fuller, D. Testuggine, and M. Khabsa, ”Llama Guard: LLM-based Input-Output Safeguard for Human-AI Conversations,” arXiv preprint, arXiv:2312.06674, 2023.

[13] S. Parida, S. Panwar, K. Lata, S. Mishra, and S. Sekhar, ”Building pre-train LLM Dataset for the IN- DIC Languages: a case study on Hindi,” arXiv preprint, arXiv:2407.09855, 2024.

[14] M. Besta et al., ”Multi-Head RAG: Solving Multi-Aspect Problems with LLMs,” arXiv preprint, arXiv:2406.05085,2024.

[15] G. B and A. Purwar, ”Evaluating the Efficacy of Open-Source LLMs in Enterprise-Specific RAG Systems: A Comparative Study of Performance and Scalability,” arXiv preprint, arXiv:2406.11424, 2024.

[16] Ma et al., ”LASER: LLM Agent with State- Space Exploration for Web Navigation,” arXiv preprint, arXiv:2309.08172, 2024.

[17] A. Balachandran, ”Tamil-Llama: A New Tamil Language Model Based on Llama 2,” arXiv preprint, arXiv:2311.05845, 2023.

[18] R. Zhang, J. Han, C. Liu, P. Gao, A. Zhou, X. Hu, S. Yan, P. Lu, H. Li, and Y. Qiao, ”LLaMA-Adapter: Efficient Fine-tuning of Language Models with Zero-init Attention,” arXiv preprint, arXiv:2303.16199, 2024.

[19] Z. Zhang, A. Zhang, M. Li, H. Zhao, G. Karypis, and A. Smola, ”Multimodal Chain-of-Thought Reasoning in Language Models,” arXiv preprint, arXiv:2302.00923, 2024.

[20] S. Rezayi et al., ”AgriBERT: Knowledge-Infused Agricultural Language Models for Matching Food and Nutrition,” in *IJCAI*, pp. 5150-5156, 2022.