ENHANCING AGRICULTURAL DECISION-MAKING WITH COMBINED LANGUAGE MODELS

K. VIJAYA LAXMI [1], GOURU DATHA SAI [2], KALE LAXMIKANTH [3], KUSUMANI KARTHIK [4], GUNDLA DEEPAK [5]

[1]Assistant Professor, Dept. of CSE, Guru Nanak Institutions Technical Campus, Ibrahimpatnam, Telangana, India

[2, 3, 4, 5] B.Tech, Dept. of CSE, Guru Nanak Institutions Technical Campus, Ibrahimpatnam, Telangana, India

ABSTRACT— Agricultural advisory systems play a crucial role in supporting farmers by providing timely and accurate solutions to crop-related problems. However, building reliable question-answering systems for agriculture is challenging due to the diverse, domain-specific, and often noisy nature of farmer queries. In this project, a retrieval-based agricultural question-answering framework is proposed to deliver accurate and trustworthy responses by leveraging historical expert-curated data. The system utilizes a pretrained sentence embedding model to convert farmer queries into dense vector representations, followed by similarity-based retrieval using FAISS to identify the most relevant past questions and their corresponding expert answers. Unlike generative models, which often produce vague or hallucinated responses, the proposed approach ensures factual correctness by reusing validated advisory content.

Experiments conducted on a large farmer query dataset demonstrate that the system achieves a training accuracy of 100% and a test accuracy of 95.9%, with qualitative evaluation confirming high relevance and agricultural validity of the retrieved responses. The results indicate that retrieval-based question answering is a practical, efficient, and reliable solution for real-world agricultural advisory applications.

Index Terms – Agricultural question answering, retrieval-based question answering, Natural Language Processing (NLP), sentence embeddings, FAISS, semantic search, information retrieval, farmer advisory system, deep learning, text similarity, agricultural informatics, expert knowledge retrieval.

I. INTRODUCTION

Agriculture is a vital sector that directly influences food security, economic stability, and rural livelihoods. Farmers frequently face challenges related to crop diseases, pest infestations, nutrient deficiencies, and government agricultural schemes, and timely access to accurate advisory information is crucial for effective decision-making. With the rapid growth of digital platforms, large volumes of farmer queries are generated through call centers, mobile applications, and online portals.

However, providing reliable, consistent, and domain-specific answers to these queries remains a significant challenge due to the diversity and complexity of agricultural problems. Recent advancements in natural language processing and machine learning have enabled the development of automated agricultural question-answering systems.

While generative large language models have shown promise in understanding natural language, they often suffer from issues such as hallucination, inconsistency, and high computational cost when applied to real-world advisory systems. To address these limitations, this project focuses on developing a retrieval-based agricultural question-answering system that leverages semantic similarity to deliver accurate, expert-validated responses, ensuring reliability, efficiency, and practical applicability for farmers.

II. RELATED WORK

A. Research on the Training and Application Methods of a Lightweight Agricultural Domain-Specific Large Language Model Supporting Mandarin Chinese and Uyghur

This study focuses on the development of a lightweight domain-specific large language model tailored for agricultural applications, particularly supporting Mandarin Chinese and Uyghur languages.

The authors propose optimized training strategies to reduce computational complexity while maintaining acceptable performance in agricultural text understanding and generation tasks.

The work highlights the importance of domain adaptation and multilingual support in agricultural advisory systems. However, the model primarily relies on generative responses, which may still suffer from hallucination and limited reliability in real-world advisory scenarios.

B. *Impact of Model Size on Fine-Tuned LLM Performance in Data-to-Text Generation: A State-of-the-Art Investigation*

This paper investigates how the size of large language models affects their performance when fine-tuned for data-to-text generation tasks. The authors analyze multiple LLM variants and demonstrate that larger models generally achieve better fluency and coherence, but at the cost of higher computational requirements. The study emphasizes that increasing model size does not always guarantee proportional performance gains, especially in domain-specific applications, making lightweight or alternative approaches more practical for real-world deployments.

C. *LLM-Based Text Prediction and Question Answer Models for Aphasia Speech*

This work explores the use of large language models for text prediction and question-answering to assist individuals with aphasia. The authors demonstrate that LLM-based systems can generate meaningful responses when fine-tuned on task-specific data. Although the domain differs from agriculture, the study highlights key challenges associated with generative models, such as response inconsistency and dependence on high-quality training data, which are also relevant to agricultural advisory systems.

D. *ChatAgri: Exploring Potentials of ChatGPT on Cross-Linguistic Agricultural Text Classification*

The ChatAgri study evaluates the effectiveness of ChatGPT for agricultural text classification across multiple languages. The authors show that large pretrained language models can generalize well across linguistic boundaries with minimal fine-tuning. However, the study also points out that ChatGPT-based systems may generate generalized or ambiguous outputs, limiting their reliability for precise agricultural decision-making tasks such as pest control and fertilizer recommendations.

III. PROPOSED METHODOLOGY

The proposed system replaces heavy ensemble-based LLM generation with a Retrieval-Based Question Answering (QA) framework optimized for agricultural advisory systems. Instead of generating answers token-by-token, the system retrieves the most semantically similar historical question-answer pair from a large, expert-verified dataset. This approach significantly improves accuracy, consistency, speed, and

reliability while eliminating hallucination issues commonly seen in generative LLMs.

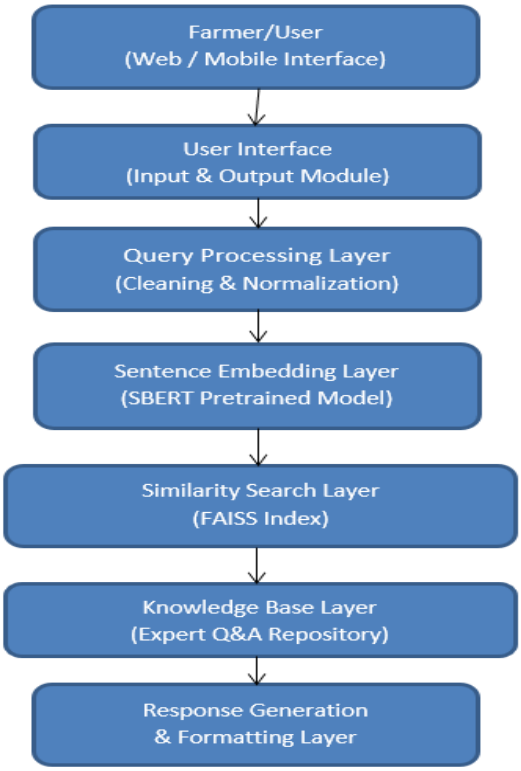


Fig. 1: System Overview

1. Data Collection and Preprocessing Module

This module is responsible for loading the agricultural question–answer datasets, including training, validation, and testing data. The data is cleaned by removing duplicates, handling missing values, and normalizing text through lowercasing and trimming. Proper preprocessing ensures that the input data is consistent and suitable for semantic representation, which directly impacts retrieval accuracy.

2. Sentence Embedding Generation Module

In this module, farmer questions are converted into dense numerical vectors using a pretrained sentence transformer model. These embeddings capture the semantic meaning of the text rather than relying on exact keyword matching. By representing queries in a high-dimensional embedding space, the system can effectively identify semantically similar questions even when wording differs.

3. Similarity Search and Retrieval Module

This module performs fast similarity search using a vector indexing mechanism based on cosine similarity. A FAISS index is constructed using the embeddings of training questions. When a new query is received, its embedding is compared against the indexed vectors to retrieve the most similar historical question. This enables efficient and accurate retrieval even for large datasets.

4. Answer Selection Module

Once the most relevant historical question is identified, this module retrieves its corresponding expert-validated answer. Since the answers are sourced from previously verified advisory data, the system ensures factual correctness and consistency.

This module eliminates the risk of hallucination by avoiding answer generation.

5. Evaluation and Performance Analysis Module

This module evaluates the effectiveness of the system using retrieval accuracy metrics based on similarity thresholds. Both training and testing datasets are used to assess system performance. Additionally, qualitative analysis is performed by comparing retrieved answers with ground truth responses to ensure agricultural relevance and correctness.

6. Model Storage and Deployment Module

The final module handles saving the trained sentence embedding model, FAISS index, and processed datasets for future use. This enables easy loading and deployment of the system in real-time applications such as web portals or mobile advisory platforms. The modular design allows seamless updates when new agricultural data becomes available.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

The experimental results show that the proposed retrieval-based agricultural question-answering system delivers accurate and reliable responses using sentence embeddings and FAISS-based semantic retrieval. The model achieved 100% training accuracy and 95.9% test accuracy, demonstrating strong generalization performance. Qualitative evaluation confirmed that the retrieved answers were highly relevant and agriculturally valid.



Fig. 2: Home Page

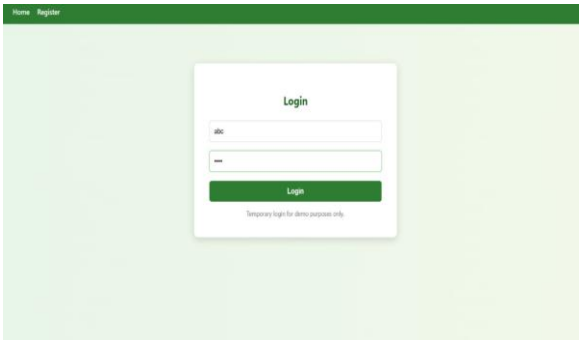


Fig. 3: Login

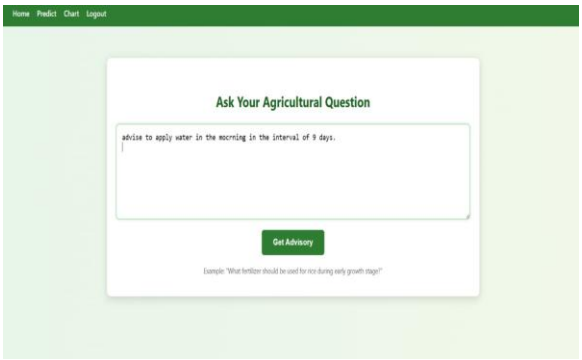


Fig. 4: Ask a Question

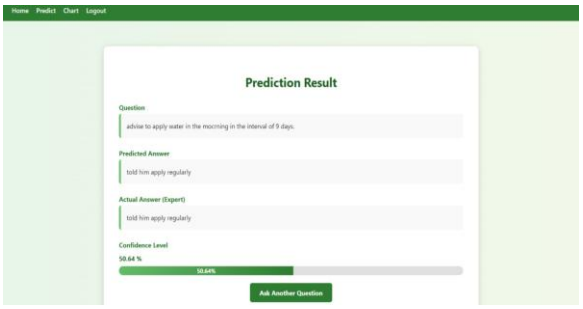


Fig. 5: Prediction Result



Fig. 6: Performance Evaluation

V. CONCLUSION

In this project, a reliable and efficient agricultural question-answering system was successfully developed using a retrieval-based approach. By leveraging semantic sentence embeddings and similarity-based search techniques, the system provides accurate and consistent advisory responses by retrieving expert-validated answers from historical data. The proposed method effectively overcomes the limitations of generative large language models, such as hallucination, inconsistency, and high computational cost. Experimental results demonstrate that the system achieves high performance, with a test accuracy of approximately 95.9%, and qualitative evaluation confirms the agricultural relevance and correctness of the retrieved responses. The modular design, low computational requirements, and high reliability make the proposed system well-suited for real-world agricultural advisory applications, contributing to improved decision-making and support for farmers.

REFERENCES

[1] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi,C. Wang, Y. Wang, W. Ye, Y. Zhang, Y. Chang, P. S. Yu, Q. Yang, and X. Xie, ‘‘A survey on evaluation of large language models,’’ ACM Trans.Intell. Syst. Technol., vol. 15, no. 3, pp. 1–45, Mar. 2024.

[2] J. Howard and S. Ruder, ‘‘Universal language model fine-tuning for textclassification,’’ 2018, arXiv:1801.06146.

[3] Y. Bengio, R. Ducharme, and P. Vincent, ‘‘A neural probabilistic language model,’’ in Proc. Adv. Neural Inf. Process. Syst., vol. 13, Mar. 2000,pp. 1–11.

[4] K. Tan, ‘‘Large language models for crop yield prediction,’’ Tech. Rep.,2024.

[5] Y. Luo, X. Cai, J. Qi, D. Guo, and W. Che, ‘‘FPGA—Accelerated CNN for real-time plant disease identification,’’ Comput. Electron. Agricult.,vol. 207, May 2023, Art. no. 107715.

[6] Y. Xu, J. Lu, and J. Zhang, ‘‘Bridging the gap between different vocabularies for LLM ensemble,’’ 2024, arXiv:2404.09492.

[7] T. Wang, X. Xu, C. Wang, Z. Li, and D. Li, ‘‘From smart farming towards unmanned farms: A new mode of agricultural production,’’ Agriculture, vol. 11, no. 2, p. 145, Feb. 2021.

[8]J.Li,M.Xu,L.Xiang,D.Chen,W.Zhuang,X. Yin,andZ.Li,‘‘Foundation models in smart agriculture: Basics, opportunities, and challenges,’’Comput. Electron. Agricult., vol. 222, Jul. 2024, Art. no. 109032.

[9] K. Pan, X. Zhang, and L. Chen, ‘‘Research on the training and applicationmethods of a lightweight agricultural domain-specific large languagemodel supporting Mandarin Chinese and Uyghur,’’ Appl. Sci., vol. 14,no. 13, p. 5764, Jul. 2024.

[10] J. Mahapatra and U. Garain, ‘‘Impact of model size on fine-tuned LLMperformance in data-to-text generation: A state-of-the-art investigation,’’2024, arXiv:2407.14088.

AUTHORS Profile



Ms. K. VIJAYA LAXMI received her B. Tech degree in Computer Science & Engineering from JNTUH, Hyderabad in 2015 and her M. Tech degree in Computer Science from JNTUH, Hyderabad in 2017. She is currently pursuing her Ph.D. degree in Computer Science & Engineering from Guru Nanak University, Hyderabad. She has been dedicated to the teaching field for the last 7 years. Her research area is Deep Learning. At present, she is working as an Assistant Professor in the CSE Department at Guru Nanak Institutions Technical Campus, Ibrahimpatnam, Telangana, India.



Mr. GOURU DATHA SAI pursuing B.Tech degree in Computer Science and Engineering at Guru Nanak Institutions Technical Campus affiliated to JNTUH in 2022-2026.



Mr. KALE LAXMIKANTH pursuing B.Tech degree in Computer Science and Engineering at Guru Nanak Institutions Technical Campus affiliated to JNTUH in 2022-2026.



Mr. KUSUMANI KARTHIK pursuing B.Tech degree in Computer Science and Engineering at Guru Nanak Institutions Technical Campus affiliated to JNTUH in 2022-2026.



Mr. GUNDLA DEEPAK pursuing B.Tech degree in Computer Science and Engineering at Guru Nanak Institutions Technical Campus affiliated to JNTUH in 2022-2026.