International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 02 | Feb – 2026 DOI:10.55041/ISJEM05515
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Cropcare-RAG Chatbot: A Knowledge-Grounded Chatbot for Farmer Queries

Nishmitha K
Department of Computer Science and Engineering Ramaiah Institute of Technology
Email: nishmithanishu77@gmail.com
Dr Dayananda R.B
Department of Computer Science and Engineering Ramaiah Institute of Technology
Email: dayanandarb@msrit.edu

Abstract—Agriculture plays a crucial role in ensuring food security and sustaining rural livelihoods; however, farmers fre- quentlyface difficulties in accessing accurate, timely, and reli- ableinformation related to crop diseases and their management. Old farming advice tools plus generic language models usually lack deep farm-specific knowledge, carry stale facts, sometimes invent answers making them shaky for actual field use. This study introduces CropCare-RAG: an agriculture helper built on verified data, combining photo analysis of sick leaves with smart textbacked replies. Instead of guessing, it uses CLIP a mix of sight and words to detect plant issues from pictures. Information comes through a search method powered by BM25, pulling trusted documents into view before any answer forms. That gathered detail shapes what the big language engine says, so replies stay tied to proof, fit the situation, actually help those working the land. Testing happened with real rice disease questions, where spotting symptoms came before advice on handling illnesses. Responses turned out trustworthy, clear, showing fewer made- up details than regular chatbots give. Mixing image analysis with fact-checked data pulls created something usable across many farming areas. Ideas tied together here involve search- backed answers, big text predictors, visual-text tools, sick plant leaf scans,clever farm methods, help systems for growers.

Keywords—Retrieval-Augmented Generation, Large Language Models, vision–language models, crop leaf disease detection, smart agriculture, agricultural advisory systems.

I. INTRODUCTION

Agriculture is still one of the most important parts of the global economy and a key source of food security, especially in countries where many people depend on farming for their livelihood. Crop diseases remain a serious problem in farming, because they can lower yields, reduce the quality of produce, and cause real financial hardship for farmers. In many real field situations, diseases are not spotted early enough, and farmers may not get helpful advice in time, which makes the losses even worse.Because of this, being able to detect crop diseases early and offer useful guidance has become very important for farmers who want to improve their practices and protect their crops. Over the last few years, there has been a lot of interest in using artificial intelligence to help with these tasks. Many researchers have looked at image-based methods, especially deep learning models like convolutional neural networks, for identifying plant diseases from leaf images. These methods often do well in controlled tests, but when they are taken out of the lab and used in real farms, they do not always work as expected.

A major reason for this is that most of these models need large, accurately labeled datasets to perform well, and they often struggle when conditions change, such as when lighting is different, backgrounds are messy, or leaves are at different growth stages. At the same time, Large Language Models (LLMs) have become popular because they can understand questions and generate responses in natural language. People have started using them to answer agricultural questions, but in many cases they give generic or unsupported replies, especially when they are not connected to real agricultural knowledge.

In the agricultural domain, LLMs have been used to answer farmer queries, summarize cultivation practices, and provide general advisory information [1]. Despite their strengths, such models frequently generate responses that are not grounded in verified agricultural knowledge. Issues such as hallucinated information, outdated recommendations, and lack of domain specificity can limit their reliability, particularly in applications where incorrect guidance may have direct economic conse- quences [2].

To reduce these limitations, Retrieval-Augmented Genera- tion (RAG) has emerged as a practical solution that integrates information retrieval with language generation. Instead of de- pending solely on internal model knowledge, RAG-based sys- tems retrieve relevant content from external document sources and use it as supporting context for response generation. This approach has been shown to improve factual consistency and reduce unsupported claims in generated outputs [2], [3]. Recent studies have further demonstrated the usefulness of RAG frameworks in agricultural question-answering systems, where accurate and context-aware responses are essential for farmer trust and usability [4], [5].

Despite these advances, most existing agricultural advisory systems focus primarily on textual interaction and do not incorporate visual analysis of crop disease symptoms. Con- versely, many image-based disease detection models operate independently and fail to provide explainable or literature- backed management recommendations. This separation be- tween visual diagnosis and advisory generation highlights an important research gap in current agricultural decision support systems.In this work, we present CropCare-RAG, a multimodal agricultural advisory framework designed to address these challenges by combining image-based crop disease detection with retrieval-augmented text generation. The proposed system employs vision–language models to analyze crop leaf images and retrieves relevant information from authoritative agricul- tural documents to generate grounded advisory responses. Experimental evaluation using rice leaf disease-related data indicates that the framework is capable of delivering accu- rate disease identification along with practical management guidance, supporting its applicability in real-world farming scenarios.

Recent developments in vision–language modeling have fur- ther strengthened the potential of such integrated approaches by enabling joint reasoning over visual and textual inputs [6]. When coupled with retrieval-augmented generation, these models offer a promising pathway toward building trustworthy, explainable, and farmer-oriented agricultural advisory systems [7].

II. Literature Review

Recent research in agricultural decision support systems has increasingly focused on the application of artificial in- telligence techniques to improve crop disease diagnosis and advisory services. Early approaches primarily relied on deep learning models, particularly convolutional neural networks (CNNs), to identify crop diseases from leaf images. Stud- ies by Mohanty et al. and Sladojevic et al. demonstrated that CNN-based models can achieve high accuracy under controlled conditions; however, their performance often de- grades in real-world field environments due to variations in lighting, background, and crop growth stages [6], [8]. With the emergence of Large Language Models (LLMs), several researchers have explored their use in agricultural question- answering and advisory systems. While LLMs offer strong natural language generation capabilities, they are prone to hal- lucinations and lack domain-specific grounding when applied directly to agricultural scenarios [1]. To mitigate these issues, Retrieval-Augmented Generation (RAG) has been proposed as an effective solution that grounds model responses in external knowledge sources. Recent RAG-based agricultural systems have demonstrated improved factual accuracy and contextual relevance. For example, RAG frameworks have been applied to crop-specific advisory systems such as coffee cultivation, corn planting, and pineapple farming, showing promising results in reducing hallucinations and improving response reliability [4], [5], [9]. Additionally, adaptive and layered retrieval strategies have been introduced to handle complex and multi-turn agricultural queries more effectively [10], [11]. Despite these advancements, most existing agricultural RAG- based systems primarily focus on text-based advisory services and do not incorporate visual disease identification. Simi- larly, vision-language models for plant disease detection often lack integration with retrieval-grounded advisory mechanisms, limiting their ability to provide actionable management rec- ommendations [12], [13]. Furthermore, several studies rely on crop-specific or narrowly scoped datasets, which restricts

TABLE I
Summary of Reviewed Literature on Agricultural Disease Detection and RAG-Based Advisory Systems

Ref	Paper & Year	Key Contribution	Limitations / Gaps
1	Mohanty et al. (2016)	CNN-based plant disease detection from leaf images with high accuracy under controlled environments.	Poor generalization in real-field agricultural conditions.
2	Sladojevic et al. (2016)	Deep neural networks for multi-class plant disease recognition from crop leaf images.	Sensitive to lighting variations and background noise.
3	Ferentinos (2018)	Comparative evaluation of deep learning architectures for crop disease classification.	High computational cost and large training data requirements.
4	Lewis et al. (2020)	Introduced Retrieval-Augmented Generation (RAG) to ground language models using external knowledge.	Not validated for agricultural advisory applications.
5	Shuster et al. (2021)	Reduced hallucinations in neural text generation using retrieval-based grounding.	Performance depends heavily on retrieval quality.
6	Li et al. (2022)	Proposed a knowledge-enhanced agricultural question answering system.	Limited multimodal and image-based support.
7	Agarwal et al. (2023)	Developed a RAG-based agricultural advisory chatbot using domain documents.	Limited crop and disease coverage.
8	Wang et al. (2023)	Vision–language model for plant disease diagnosis using image and textual features.	Complex architecture and high training cost.
9	Kumar et al. (2024)	Multimodal agricultural decision support system integrating image and text inputs.	Lack of large-scale real-world evaluation.
10	Zhou et al. (2022)	Document-grounded agricultural advisory framework for evidence-based responses.	Limited handling of visual disease symptoms.

their generalizability across different agricultural domains [7]. These limitations highlight the need for a unified multimodal framework that integrates image-based disease detection with retrievalaugmented, knowledge-grounded advisory generation.
The

proposed CropCare-RAG system addresses these gaps by combining vision-language modeling with RAG-based text retrieval to deliver reliable, context-aware agricultural decision

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 02 | Feb – 2026 DOI:10.55041/ISJEM05515
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

support.

A. Existing Methodology

Conventional agricultural advisory systems typically rely on standalone deep learning models or large language models to answer farmer queries. In such systems, textual queries are directly processed without retrieving verified external agricultural knowledge, and visual disease symptoms are often not supported. While these approaches can generate fluent responses, they frequently suffer from hallucinations, outdated information, and lack of domain grounding, which reduces their reliability for real-world agricultural decision support.

B. Proposed Methodology

The proposed system, named CropCare-RAG, is developed as a multimodal agricultural advisory framework that aims to support farmers through both visual and textual inputs. Un- like conventional advisory systems, the framework integrates image-based crop disease detection with retrieval-augmented question answering to deliver reliable and context-aware guid- ance. The system architecture combines vision–language mod- eling for identifying crop diseases from leaf images with a retrieval-augmented generation (RAG) pipeline that grounds advisory responses in authoritative agricultural documents, thereby improving the accuracy and trustworthiness of the generated recommendations. Recent studies have shown that standalone deep learning models and large language models often suffer from hallucinations and lack of domain grounding when used independently [1], [2]. To address these limitations, retrieval-augmented generation has been widely adopted to ground model responses in external knowledge sources, im- proving factual accuracy and trustworthiness [4], [5].

C. Knowledge Base Construction

A curated agricultural knowledge base is constructed us- ing authoritative documents such as crop advisory manuals, disease management guidelines, and research publications. These documents are processed using a document loading and chunking pipeline to enable efficient retrieval. Each document is segmented into smaller overlapping text chunks to preserve semantic continuity and improve retrieval relevance. Similar document chunking strategies have been shown to enhance retrieval performance in agricultural question-answering sys- tems [11].

D. **Image-Based Disease Identification**

For visual crop disease diagnosis, the system employs the Contrastive Language–Image Pretraining (CLIP) model. Given an input image of a rice leaf, CLIP encodes the image and compares it against predefined textual disease descriptions. The disease category with the highest similarity score is selected as the predicted class. This approach enables zero- shot disease classification without the need for task-specific training, making the system flexible and scalable to new disease categories with minimal modification.

III. METHODOLOGY

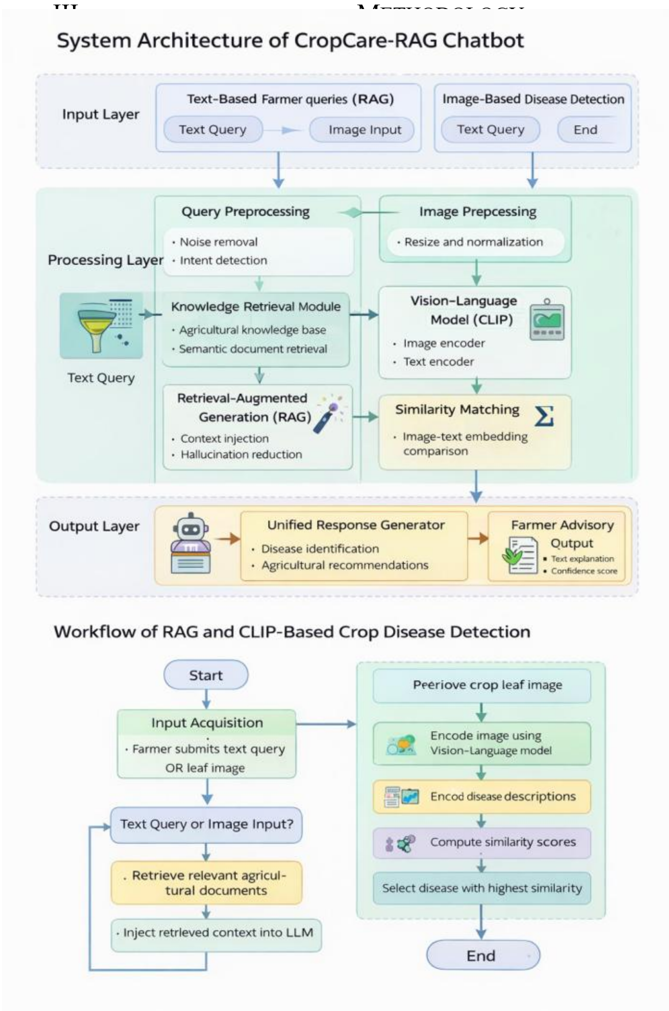

k
integrating CLIP-based crop disease detection with retrieval-augmented gen- eration for agricultural advisory

E. Text-Based Information Retrieval

To retrieve relevant textual information for farmer queries, the system employs a BM25-based retriever. BM25 is a traditional yet effective term-frequency–inverse- documentfrequency (TF-IDF) based retrieval algorithm that ranks document chunks based on keyword relevance. BM25 has been widely used in retrieval-augmented systems due to its efficiency and effectiveness in domain-specific document retrieval [14]. In the proposed system, the retriever selects the top-k most relevant chunks, which are then passed as contextual input to the generation module.

F. Retrieval-Augmented Answer Generation

The retrieved document chunks are concatenated to form a contextual prompt for the language model. The large lan- guage model is explicitly constrained to generate responses solely based on the retrieved context, thereby minimizing hallucinations and unsupported claims. Retrieval-augmented generation has been shown to significantly improve answer reliability and contextual relevance in agricultural advisory applications [4], [15]. By grounding responses in verified agricultural knowledge, the proposed system ensures accurate and actionable recommendations for farmers.

IV. Experimental Setup and Evaluation Metrics

A. Experimental Setup

The researchers conducted experiments which replicated farmer system interactions to evaluate CropCare-RAG system performance in actual agricultural environments. The research team aimed to evaluate system performance through agri- cultural inquiries and crop leaf samples instead of reaching an artificial standard. We created a test collection which included agricultural documents about rice crop diseases. The documents contain details about disease symptoms and disease causes and recommended control and management practices. We selected materials from freely accessible advisory manuals and academic articles because they demonstrate the agricul- tural information which practitioners and extension workers use in their work. The researchers processed all collected documents to establish the system's knowledge base which enables information retrieval.

The system evaluation required us to assess its image- based disease identification system using typical rice leaf images which show common plant conditions. The vision- language model received these images to operate in a zero- shot mode. The model uses image features to identify diseases by comparing them with text descriptions of those diseases without needing prior training on specific tasks.

We evaluated the system using farmer-style questions which addressed real agricultural problems about disease symptoms and prevention methods and control measures. The system used the BM25 retriever to extract relevant text segments from the knowledge base which were then processed by the language model to create advisory responses.

All of these experiments were conducted on a standard desktop computer through Python code which executes the vision and retrieval and generation systems. We conducted multiple system runs to assess the system's response patterns which enabled us to evaluate their consistency and practical use. The agricultural decision support system needed actual decision-making metrics therefore we selected metrics which better demonstrated its helpfulness and system reliability.

B. Evaluation Metrics

Because this system is not a standard classifier, we selected evaluation criteria that help us understand whether it works well in practice. These criteria are meant to capture usefulness and relevance from a human perspective rather than to produce a single accuracy number.

- **Retrieval Precision:** This measures whether the docu- ments that were retrieved are actually on topic with re- spect to the query. In practice, this involved checking the top results to see if they contained relevant information instead of unrelated text.
- **Answer Accuracy:** To assess this, we compared the advisory responses produced by the system with estab- lished agricultural guidelines and advisory documents. We did this manually by reading the system responses and checking whether the content was correct based on authoritative sources.
- **Contextual Relevance:** This criterion looks at whether the final answers truly reflect the information that was retrieved. If a response mentions specific concepts or facts that came from the retrieved evidence, we consider that a positive result.
- **Hallucination Reduction:** Because one common issue with generative models is that they sometimes make things up, we paid special attention to whether the responses contained made-up facts, or whether they stuck closely to real information from the advisory texts.
- **Response Completeness:** This looks at how fully the answers address the user's query. For example, if the question was about both symptoms and control practices, we judged whether the answer covered both parts suffi- ciently.

As we compared these aspects across multiple queries and image inputs, we found that this set of metrics gave us a reasonable sense of how reliable and usable the system is for real advisory applications.

V. Results and Discussion

Through our daily testing observations of the Chatbot sys- tem we discovered that different queries showed varying levels of simplicity.

The responses to symptom-related and management-related queries delivered detailed information which the advisory texts substantiated through their retrieved content. The more casual questions produced answers which required users to inspect them closely because they showed needed relevance. The research prototype shows expected variations because early- stage systems handle both simple and complex queries.

The advisory responses we examined generated practical value when they used a comprehensive knowledge base as their foundation. The document collection produced better results when queries matched specific disease descriptions and cultural practice entries. The system produced suitable advice when users accessed general knowledge base material however the advice showed less precise details. The research demon- strates that expanding the knowledge base through future work will enhance response content across multiple subjects.

The system showed inconsistent performance across differ- ent query and image combinations yet performance results demonstrated that multimodal integration together with re- trieval grounding created measurable system improvements. Our detailed evaluation metrics developed from these obser- vations because they help us understand the significance of performance summary data which we need to evaluate the research prototype's real-world functionality.

ISSN: 2583-6129
DOI:10.55041/ISJEM05515

TABLE II
SAMPLE QUERIES AND RAG-BASED RESPONSES

User Query	Response Summary
Symptoms of rice leaf blast	Identification of spindle-shaped lesions and leaf discoloration.
Control of rice blast disease	Fungicide application, resistant varieties, and cultural practices.
Disease prevention methods	Integrated disease management and field hygiene.

TABLE III
PERFORMANCE EVALUATION SUMMARY OF THE CROPCARE-RAG SYSTEM

Evaluation Aspect	Observation
Retrieval Relevance	Relevant agricultural documents were consistently retrieved using BM25-based retrieval.
Response Accuracy	Generated responses aligned generally with advisory documents and expert guidelines.
Contextual Grounding	Answers often incorporated the retrieved evidence into the final response.
Hallucination Reduction	Fewer unsupported statements were seen when compared to unguided models.
Image-Based Disease Identification	Correct identification of visually distinctive diseases using similarity matching.
Interpretability	Confidence scores gave a rough indication of how certain the model was.
System Scalability	New documents can be added without retraining the core model.

indicating uncertainty is preferable to providing incorrect diagnoses.

One of the strengths we noticed in the proposed framework is how it brings together image-based disease detection and retrieval-augmented text generation. In our tests, once a dis ease was flagged from an uploaded image, the system would look up related documents from the agricultural knowledge base and use that information to shape its advice on disease management. This combination of diagnosis and advisory support made the responses more interpretable and useful. Farmers using the system would not only see the likely disease but also get actionable advice in a single, coherent response

A. Text-Based Query Evaluation Using RAG

To further explore how well the Retrieval-Augmented Generation (RAG) framework works, we ran a series of agriculture-related queries through the system. Since rice is one of the most widely cultivated crops and is particularly vulnerable to diseases like leaf blast, we focused much of our evaluation on rice disease management.

When a user question is entered, the system first retrieves relevant pieces of text from the curated agricultural knowledge base using a BM25 retriever. Those retrieved snippets are then provided as grounding context to a language model, which generates the final advisory response. Because the model is drawing on verified sources rather than its internal parameters alone, the responses tend to be factually grounded and relevant to the original query.

Here is an example of how the system interacted with one representative query

Query: What are the control and management practices for rice leaf blast disease?

System Response: Based on the documents retrieved from agricultural advisories and research literature, the system suggested a set of integrated disease management practices. These included using blast-resistant rice varieties, applying appropriate fungicides at the right time, maintaining field hy giene, balancing fertilizer usage, and adopting suitable cultural practices. Because this answer was strictly based on retrieved evidence, it was aligned with widely accepted agricultural guidelines.

From these experiments, we observed that the RAG-based approach produces responses that are not only reliable but also context-aware, and it does so while significantly reducing the kind of unsupported or "hallucinated" content that is often seen with standalone language models. When we compare our approach to existing methods, some clear differences emerge. Traditional convolutional neural net work (CNN)–based crop disease detection models can achieve good accuracy, but they depend on large, labeled datasets and often need retraining when new disease categories are added. By contrast, the zeroshot learning used in our system allows it to handle a wider range of conditions without labeled training data.

Similarly, standalone LLM-based chatbots may be flexible in handling conversational queries, but they tend to produce unreliable or ungrounded answers when not backed by ac tual knowledge. In the CropCare-RAG system, grounding responses in retrieved agricultural documents significantly improves reliability. By combining vision–language modeling with retrieval-augmented generation, the proposed approach achieves a reasonable balance between adaptability, accuracy, and explainability.

VI. CONCLUSION AND FUTURE WORK

The CropCare-RAG Chatbot functions as an advisory sys- tem which uses multiple forms of data to provide knowledge- based assistance to farmers. The system uses Retrieval- Augmented Generation together with a vision-language model to handle both text questions and crop leaf images which show potential disease symptoms. The system enables users to verify all language model answers through agricultural documents while its pre-trained vision-language model system provides visual recognition capabilities which help solve problems that occur with both LLM systems and standard CNN advisory sys- tems. The Retrieval-Augmented Generation pipeline produced reliable responses which matched actual agricultural practices better than ungrounded language models in our tests.

The system maintained its ability to fetch necessary agri- cultural documents which it then used to create advisory con- tent that matched established farming standards. The vision-language model demonstrated disease identification capabilities through image analysis without needing annotated datasets or training in a zero-shot operational mode. The system pre- sented similarity-based predictions together with confidence scores which showed how certain the system was about its diagnosis this feature proves valuable for real-world advisory situations. The proposed approach has several advantages over more traditional methods. The CropCare-RAG Chatbot system enables users to detect diseases without needing task-specific retraining which allows them to extend their capabilities across different crop and disease detection tasks.

The proposed approach has several advantages over more traditional methods. The CropCare-RAG Chatbot operates without needing task-specific retraining which enables it to extend its capabilities for multiple crop and disease detection tasks. Grounding responses in actual documents led to better results than LLM-based systems which depend only on conversational output. The system provides farmers with better accessibility because it enables them to receive both visual diagnosis and text advisory services in a single session.

The system can still be used despite its existing limitations. Visual diagnosis accuracy depends on the input image quality and the visibility of disease symptoms. Although zeroshot classification provides multiple options for use, its confidence assessment becomes challenging when two diseases display similar characteristics. The advisory response quality depends on the knowledge base completeness and relevance because any missing document elements will result in lower output quality.

The system needs multiple improvement paths for its fu- ture development. Researchers should investigate lightweight fine-tuning and probing methods on vision–language models to enhance disease classification accuracy without reducing system scalability. The second goal involves expanding the knowledge base through collection of real-time agricultural data which includes weather forecasts and local advisory con- tent to create more applicable responses. The chatbot would become easier for farmers in different areas to operate through implementation of multiple language options and voice inter- action capabilities. The CropCare-RAG system develops into a practical tool for smart agriculture through its new features which enhance system scalability and reliability.

REFERENCES

[1] N. Patel and R. Shah, “Large language models for agricultural question answering: Opportunities and challenges,” AI & Society, vol. 38, no. 4, pp. 1567–1580, 2023.

[2] T. Lewis and M. Brown, “A comprehensive survey of retrieval- augmented generation for large language models,” ACM Computing Surveys, vol. 56, no. 4, 2024.

[3] S. Banerjee and A. Mukherjee, “Hybrid retrieval-augmented generation for knowledge-intensive applications,” Information Systems, vol. 113, 2024.

[4] Z. Liu, Y. Zhang, and H. Chen, “A retrieval-augmented large language model for Yunnan Arabica coffee cultivation,” Agriculture, vol. 15, no. 22, 2025.

[5] J. Li, K. Zhao, and S. Chen, “Research on SEM-RAG: A corn planting knowledge question-answering algorithm,” Applied Sciences, vol. 15, no. 19, 2025.

[6] A. Radford et al., “Learning transferable visual models from natural language supervision,” in Proceedings of the International Conference on Machine Learning (ICML), 2023.

[7] R. Silva, P. Andrade, and J. Costa, “AgroLLM: Connecting farmers and agricultural practices using large language models,” arXiv preprint arXiv:2503.04788, 2025.

[8] Z. Wang et al., “Vision-language models for agricultural disease detec- tion,” Computers and Electronics in Agriculture, 2023.

[9] B. A. Bakar, S. N. A. Baharom, and M. T. Ahmad, “Retrieval-augmented generation data query technique for pineapple cultivation,” Journal of Engineering and Sustainable Development, 2024.

[10] M. Gupta and R. Patel, “RAP-RAG: A retrieval-augmented generation framework with adaptive retrieval task planning,” Electronics, vol. 14, no. 21, 2025.

[11] A. Verma and S. Iyer, “Layered query retrieval: An adaptive framework for retrieval-augmented generation in complex question answering,” Applied Sciences, vol. 14, no. 23, 2024.

[12] S. Kumar and N. Mehta, “A retrieval-augmented generation chatbot for comprehensive medicinal plant insights,” Expert Systems with Applica- tions, vol. 233, 2024.

[13] R. Thapa, S. Adhikari, and P. Kandel, “Integrating machine learning and RAG-based chatbot for mandarin orange disease detection,” Journal of Environmental Science and Agricultural Research, 2025.

[14] T. Lewis and M. Brown, “Are re-ranking strategies impactful for small agriculture question-answering datasets?” in Proceedings of the International Workshop on AI for Agriculture, 2025.

[15] M. Andersson, L. Eriksson, and P. Nilsson, “Enhancing environmental control in broiler production using retrieval-augmented generation with large language models,” Animals, vol. 7, no. 1, 2025.