

Integrating LLM for Cotton Soil Analysis in Smart Agriculture System

Syed Hassan Ali^{*1}, Muhammad Farrukh Shahid¹, M. Hassan Tanveer², and Abdul Rauf¹

¹Department of AI & Data Science (FAST-NUCES, Karachi, Pakistan).

²Department of Robotics & Mechatronics Engineering (Kennesaw State University, Marietta, GA, USA).

***Correspondence:** syedhassanali39@gmail.com

Citation | Ali. S. H, Shahid. M. F, Tanveer. M. H, Rauf. A, “Integrating LLM for Cotton Soil Analysis in Smart Agriculture System”, IJIST, Special Issue pp 283-294, Oct 2024

DOI | <https://doi.org/10.33411/IJIST/1111>

Received | Oct 16, 2024 **Revised |** Oct 19, 2024 **Accepted |** Oct 24, 2024 **Published |** Oct 30, 2024.

Cotton is a critical crop for the agricultural economy, with its productivity closely tied to soil quality, particularly soil nutrient levels and pH. Monitoring and optimizing these properties is essential for sustainable cotton cultivation. This study proposes using fine-tuned Large Language Models (LLMs)—specifically GPT-2 and LLaMA-2—to automate soil analysis and produce detailed soil reports with actionable recommendations, addressing the limitations of traditional machine learning models in this context. A custom dataset was created by extracting key information from cotton-specific resources, focusing on soil nutrient interpretation and recommendations across different growth stages. Fine-tuning was applied to GPT-2 and LLaMA-2 models (specifically, the Nous Research version LLaMA2-7b-hf from Hugging Face), enabling them to generate data-driven reports on cotton soil health. The fine-tuned GPT-2 model achieved a training loss of 0.093 and an evaluation loss of 0.086, outperforming LLaMA-2, which had a training loss of 0.033 and an evaluation loss of 0.25. Evaluation with BERT Score showed that GPT-2 scored average Precision, Recall, and F1 scores of 0.9284, 0.9308, and 0.9296, respectively, highlighting its superior report accuracy and contextual relevance compared to LLaMA-2. The generated reports included soil properties and actionable nutrient management recommendations, effectively supporting optimized cotton growth. Implementing fine-tuned LLMs for soil report generation enhances nutrient management practices, contributing to higher yields and more sustainable cotton farming.

Keywords: Large Language Models (LLMs), Soil Health, Cotton Soil Reports, Cotton Farming, Soil Analysis.



Introduction:

Agriculture is a key economic sector for many countries, supporting the livelihoods of billions worldwide. Soil is a fundamental component of agriculture, crucial for crop productivity, nutrient availability, and ecosystem health, underscoring the need for efficient soil management as global food demand rises [1]. Soil degradation is driven by both natural factors and improper land use [2], and analyzing soil fertility is essential for effective farming [3]. Soil composition directly affects crop productivity [4]. At the start of the sowing season, farmers typically evaluate soil quality using a soil quality index, which involves laboratory testing—a process that is both costly and time-intensive, especially for those with limited budgets and strict timelines [5].

With agriculture being labor-intensive, rising population growth and increased demand for agricultural products have made automation increasingly important [6]. Cotton is a major crop economically, supporting the textile industry, which significantly influences the global economy [7]. Successful cotton cultivation depends on soil quality, making soil assessment critical. Traditionally, soil quality assessment has relied on laboratory techniques and expert analysis, providing valuable insights but often lacking speed and real-time updates. Faster, data-driven approaches are therefore essential.

Recent AI research in agriculture focuses on predicting cotton yield, detecting crop diseases, predicting soil properties, and classifying soil quality through traditional machine-learning techniques. For example, predicting cotton crop yield has been approached using Random Forest models based on soil chemical properties [8]. Some studies focus on a single crop, such as sugarcane, for which soil quality and fertilizer amounts are classified [3]. Soil classification with ML algorithms, such as in [4], identifies soil quality without detailed interpretations of soil factors. However, approaches that interpret dominant variables—i.e., the factors most affecting soil quality—are limited [9]. Much existing work lacks in-depth interpretation or actionable reports to improve soil quality.

AI, particularly LLMs, offers the potential to transform soil analysis through their ability to contextualize data. Pre-trained LLMs are now used in finance [10], biomedical fields [11], and agriculture [12], creating opportunities to revolutionize soil data analysis. LLMs can serve as trainers and explainers for digital agriculture [13]. This research seeks to integrate LLMs with soil quality analysis to produce detailed reports based on cotton soil readings. A custom cotton soil report dataset was created, comprising soil properties and improvement recommendations for cotton. For fine-tuning, GPT-2 [14] and LLaMA-2 (specifically, the Nous Research LLaMA2-7b-hf from Hugging Face) [15] were used. LLaMA-2 achieved a training loss of 0.033 and an evaluation loss of 0.25, while GPT-2 attained a training loss of 0.0932 and an evaluation loss of 0.086. BERT Score evaluations showed that GPT-2 outperformed LLaMA-2 with average Precision, Recall, and F1 scores of 0.9284, 0.9308, and 0.9296, respectively. Using LLMs for soil analysis reporting aims to bridge existing gaps, offering detailed, data-driven insights for cotton soil quality improvement and supporting optimal crop growth through advanced technology.

Novelty Statement:

This study presents a novel application of fine-tuned Large Language Models (LLMs)—specifically GPT-2 and LLaMA-2—to automate soil analysis and generate actionable reports for cotton farming. Unlike traditional machine learning models, this approach offers detailed interpretations and customized recommendations, addressing a critical gap in current soil management practices and enabling more efficient, data-driven agricultural decision-making.

Objectives:

The primary objectives and contributions of this research include:

- **Dataset Creation:** Developed a robust dataset containing soil nutrient interpretations and tailored recommendations, based on nutrient levels and different stages of cotton growth, sourced from agricultural resources [1].

- LLM Fine-Tuning for Report Generation: Fine-tuned Large Language Models (LLMs) to produce detailed, contextually accurate reports on cotton growth, customized to specific soil nutrient inputs [2].
- Agricultural Advancement: Pioneered the integration of LLMs into cotton soil nutrient management, with the potential to enhance crop yields and promote sustainability in cotton farming [1].

Literature Review:

Researchers worldwide are increasingly focused on automating agriculture [3]. Key areas of ongoing research include soil quality and health analysis, which are crucial for achieving high yields and ensuring sufficient food production. Effective soil management is essential for sustainable agriculture, with the first step being the assessment and adjustment of soil nutrient levels [16]. Recent advancements in machine learning and AI have significantly improved the prediction of soil quality and crop yield. With the rise of generative AI, particularly Large Language Models (LLMs), researchers are leveraging these models for various domain-specific automation tasks.

Traditional Machine Learning:

Traditional machine learning techniques have been widely applied in agriculture for tasks such as classification and identification related to soil. For instance, in [8], a Random Forest classifier was used to identify the most significant variables affecting cotton yield predictions. Soil parameters such as pH, Na, K, and other chemical properties were included as inputs, underscoring the importance of chemical soil properties. Another study in this field focused on predicting cotton yield using satellite remote sensing images, climate data, and soil parameters, applying Explainable Boosting Machines [7]. This research emphasized the importance of various features during the cotton growth cycle. Given the direct impact of soil quality on plant growth, it is essential to analyze soil to ensure it is suitable for cotton cultivation, which adds another dimension to soil-cotton studies.

In [3], soil quality classification for sugarcane was explored, where a classifier determined soil suitability based on data from sensors. The classification results were then used to guide nutrient recommendations, suggesting external fertilizers for soil enhancement. Similarly, in [4], soil fertility prediction was examined, utilizing parameters such as clay, sand, pH, N, P, and K. The study proposed four classification algorithms—Artificial Neural Networks, Decision Trees, K-Nearest Neighbors, and Random Forest—to determine soil fertility and predict crop yield. Weather also plays a role in agriculture, as demonstrated in [17], which provides daily rainfall forecasts and alerts for heavy rainfall up to three days in advance. Real-time soil quality data, as discussed in [5], helps farmers make timely decisions based on current soil conditions. Additionally, soil quality analysis for suitable crop selection has been performed using image processing techniques [2]. AI-driven systems for soil property prediction, irrigation, and fertilization, as well as real-time monitoring of key soil parameters like moisture and nutrient content, are presented in [1]. Similarly, [16] explored AI applications in soil management, including machine learning algorithms for soil property prediction, sensor technologies for real-time monitoring of soil moisture and pH, and UAV-driven precision agriculture. Most of these approaches focus on predicting, identifying, or classifying soil fertility or quality, but they often lack a comprehensive interpretation of the soil data. In contrast, [9] took a broader approach by not only predicting soil fertility but also providing a waterfall plot interpretation. This methodology consists of three layers: the first uses K-means clustering to categorize soil data into quality labels, the second layer applies Random Forest for classification, and the third layer explains the classification outcomes through the waterfall plot. While this approach highlights which variables most influence the classification, it still lacks a more detailed, descriptive analysis and actionable remedies.

Large Language Models:

Large language models (LLMs) have demonstrated remarkable effectiveness in interpreting and generating text that closely resembles human communication [18]. These models are typically trained on extensive corpora and are designed as general-purpose task models. A significant recent development is the ability to fine-tune these models for specific tasks, which has shown promising results in various domains [10]. For example, in the finance domain, LLMs like BERT have been fine-tuned for financial sentiment analysis, creating models such as Fin BERT [10]. Similarly, in the biomedical domain, LLMs have been fine-tuned to create Bio BERT, which outperforms the base BERT model in biomedical tasks [11]. LLMs have also found applications in the fashion industry, with models like GPT-FAR, termed FashionReGen, being used for fashion-related tasks [19]. In disaster reporting, LLMs have been adapted for flood disaster reporting, branded as Flood Brain [20].

Given the success of fine-tuning LLMs in these diverse domains, they are also being explored for agricultural applications. In [12], LLMs were evaluated for their performance on agriculture-related exams, employing two agents: an answer agent to generate responses and an evaluation agent to assess the correctness of those responses. In [21], GPT-3.5 was fine-tuned for the agricultural context in Nigeria, enabling farmers to interact with the model to receive agricultural information, with a central repository developed to store conversational data. Additionally, [18] demonstrates that LLMs, when combined with prompt engineering, outperform traditional machine learning models, particularly in enhancing the precision and utility of crop yield predictions. LLMs for text generation tasks are typically fine-tuned using prompt templates, and in [12], placeholders for questions, answers, and information within prompts were used, with GPT-4 outperforming GPT-3.5 and LLaMA.

While traditional machine learning models have been used to predict soil quality and nutrient levels, they often fall short in providing detailed descriptions and comprehensive reports. In [22], the integration of AI and the Internet of Things (IoT) in precision farming was explored, where IoT sensors monitor soil parameters and nutrient levels, and AI algorithms optimize fertilization, irrigation, and pest management practices. However, this research primarily relied on classical machine-learning techniques. In contrast, this research employs LLMs for specific soil-related tasks, generating detailed, actionable reports for cotton farming based on sensor data or predictions from machine learning models, addressing the gap in existing methods by providing comprehensive soil reports without requiring user queries.

Material and Methods:

Dataset Collection:

Data was gathered from a variety of sources related to cotton management, including cotton reports from research institutes, scholarly articles, and online resources. The documents were selected based on their relevance to soil nutrient recommendations for cotton at various growth stages, specifically in the Punjab region of Pakistan. Table 1 summarizes the documents used for data collection, detailing the reference ID, source title, and source information.

Data Preparation:

Soil nutrient values, their interpretations at different stages of cotton growth, and the corresponding recommendations for each stage were manually extracted from PDF documents. According to [31], there are six stages of cotton growth, as outlined in Table 2. However, our dataset includes only five of these stages, excluding the maturity stage, as it represents the point when the crop is fully developed and ready for harvest.

For each stage of cotton growth (excluding the maturity stage) listed in Table 2, we prepared interpretations and recommendations for Nitrogen, Potassium, Phosphorus, and pH based on their availability in the soil, categorized as either low or optimal. Consistent measurement units, specifically kg/ha, were used to quantify the nutrient levels in the soil.

According to [23], the optimal ranges for N, P, K, and pH are presented in Table 3. The frequency distribution of Nitrogen, Potassium, Phosphorus, and pH levels is depicted in Figure 1, while the overall frequency distribution of nutrient levels in the dataset is shown in Figure 2.

Table 1: Summary of Data Sources

Reference ID	Source Title	Source Information
[23]	Pakistan Cotton grower	Central Cotton Research Institute, Multan, Pakistan
[24]	Fertilizer Role in Sustainable Cotton Production	Central Cotton Research Institute, Multan, Pakistan
[25]	Cotton Fertility Management	University of Missouri
[26]	Role of proper management of nitrogen in cotton growth and development	International Journal of Biosciences
[27]	Potassium fertilization of cotton	Virginia Cooperative Extension
[28]	Inorganic nutrient management for cotton production in Mississippi	Mississippi State University
[29]	Phosphorus application strategies to improve cotton productivity under arid climatic conditions	International Journal of Cotton Research and Technology
[30]	Integrated Crop Management	Advancing Cotton Education Soil Fertility, National Cotton Council
[31]	Cotton Crop Development in Central Punjab	Regional Agromet Centre Pakistan Meteorological Department

Table 2. Growth stages of cotton

No	Growth Stage
1	Sowing
2	Vegetative
3	Budding
4	Flowering
5	Boll Development
6	Maturity

Table 3. Optimal values of N, P, K, and pH

Nutrient	Optimal Value
Nitrogen	140 kg/ha
Phosphorus	25 kg/ha
Potassium	125 kg/ha
pH	6.5 - 7.5




Figure 1. Frequency Distribution of Nutrients in Dataset




Figure 2. Frequency Distribution of Nutrient Levels in Dataset

Prompt Template:

To fine-tune a large language model, a prompt template is essential to guide the model in learning the desired structure. Figure 3 illustrates the prompt templates we developed using data from our dataset.




Figure 3. Prompt Template

The soil properties include nutrient levels along with labels indicating whether they are low or optimal, as well as information about the corresponding growth stage of cotton. The first template, designed for soil analysis, provides interpretations of how low or optimal nutrient levels affect cotton growth at each specific growth stage. The second template, focused on soil recommendations, offers guidance on nutrient management, with recommendations based on whether the nutrient level is low or optimal, considering the cotton growth stage.

Large Language Based Models:

Large language models (LLMs) have emerged as powerful tools capable of generating text and human-like responses. Typically, LLMs are trained on vast corpora and designed as general-purpose models. A significant recent advancement is the ability to fine-tune pre-trained models for specific tasks, as demonstrated in [10]. The LLMs used in this study are:

- **GPT-2:** GPT-2 [14] is a large transformer model with 1.5 billion parameters, trained on the WebText dataset, which comprises 8 million web pages. Its architecture consists of multiple transformer decoder layers, each incorporating self-attention mechanisms and feed-forward neural networks. Key components of the model include self-attention mechanisms to capture contextual relationships, feed-forward neural networks to learn complex text patterns, and layer normalization combined with residual connections to support efficient training and learning.
- **Llama-2:** LLaMA 2 7b [15] is a transformer-based model with 7 billion parameters, trained on a diverse and extensive dataset consisting of publicly available text data. This vast corpus enables the model to develop a broad understanding of language, encompassing a

wide range of topics and styles. For this study, we used the Nous Research version, LLAMA2-7b-hf, available on Hugging Face.

Methodology:

Figure 4 illustrates our methodology. First, we collected data from cotton-related documents and then divided it into training and testing sets for fine-tuning the LLM. After fine-tuning, we evaluated the model's performance using various evaluation metrics.




Figure 4. Methodology followed in this research

Report Generation:

The report generation process is iterative. Figure 5 illustrates this iterative process of generating a report using the LLM.




Figure 5. Report Generation Process

The nutrient readings are first converted into prompts for each nutrient and the corresponding growth stage of cotton. These prompts are then passed to the fine-tuned LLM for two purposes: first, to generate a nutrient analysis, and second, to provide recommendations for that nutrient. The final step involves merging all the responses generated by the LLM into a report template to produce the complete, generated report.

Results and Discussion:

Results:

In our research, we fine-tuned the GPT-2 and LLaMA-2 models to generate soil reports aimed at optimizing cotton growth through soil nutrient values. The dataset was split into two parts: 80% for training and 20% for testing.

GPT Fine Tuning:

The dataset was initially tokenized using a GPT-2 tokenizer, with the parameters specified in Table 4. We then fine-tuned the pre-trained GPT-2 model, and the parameters used during the fine-tuning process are listed in Table 5.

Table 4. GPT-2 Tokenizer Parameters

Name of Parameter	Value
Vocabulary Size	50257
Max Length	750
No repeat n-gram	3
No return sequences	1




Figure 6. GPT 2 train and test loss

Figure 6 displays the graph of training and testing loss for the GPT-2 model over 20 epochs. The training loss begins at 1.12 and decreases to 0.09 by the 20th epoch. Similarly, the evaluation loss starts at 0.18 and drops to 0.08 after 20 epochs.

Table 5. Parameters for GPT-2 and Llama2

Name of Parameter	GPT-2	Llama2
Batch Size	4	4
Learning Rate	0.0001	0.0001
Optimizer	Adam	Adam
Epochs	20	20

Llama 2 Fine Tuning:

The dataset was split into 80% for training and 20% for evaluation. The parameters used for fine-tuning LLaMA-2 are provided in Table 5. Figure 7 illustrates the training and evaluation loss for LLaMA-2. The training loss begins at 2.02 and decreases to 0.03 after 20 epochs, while the evaluation loss starts at 1.91 and drops to 0.25 by the 20th epoch.

The results indicate that GPT-2 outperforms LLaMA-2. Consequently, we used GPT-2 to generate soil reports based on nutrient values.

Generated Report by LLM:

Figure 8 shows a report generated by the LLM using input values (Nitrogen = 150 kg/ha, Potassium = 100 kg/ha, Phosphorus = 30 kg/ha, pH = 7, Growth Stage = sowing), which could

be derived from field data or predictions made by another model. These inputs follow the report generation process outlined in Figure 5, resulting in the final report.




Figure 7. Llama2 train and test loss

Reports Evaluation:

We evaluated the reports generated by our fine-tuned models, GPT-2 and Llama-2. While BLEU [32], a common metric in machine translation, counts n-gram overlaps between generated and reference texts, we did not use it because it fails to account for meaning-preserving lexical and compositional diversity [33]. Instead, we employed BERT Score [33], a more advanced language generation evaluation metric based on pre-trained BERT contextual embeddings. BERT Score measures sentence similarity by summing the cosine similarities between token embeddings. Unlike BLEU and ROUGE, which rely on exact token matching, BERT Score captures semantic similarity by leveraging BERT's contextual embeddings to assess the meaning between generated and reference texts. According to [33], BERT Scores are calculated as follows:

Recall (R_{BERT}):

$$R_{\text{BERT}} = \frac{1}{|x|} \sum_{x_i \in x} \max_{\hat{x}_j \in \hat{x}} x_i^\top \hat{x}_j \quad (1)$$

This equation calculates the recall of the generated text. For each token x_i in the reference text x , it finds the maximum similarity with any token $\hat{x}_j$ in the generated text $\hat{x}$. The recall is then averaged over all tokens in the reference text.

Precision (P_{BERT}):

$$P_{\text{BERT}} = \frac{1}{|\hat{x}|} \sum_{\hat{x}_j \in \hat{x}} \max_{x_i \in x} x_i^\top \hat{x}_j \quad (2)$$

This equation calculates the precision of the generated text. For each token $\hat{x}_j$ in the generated text $\hat{x}$, it finds the maximum similarity with any token x_i in the reference text x . The precision is then averaged over all tokens in the generated text.

F1 Score (F_{BERT}):

$$F_{\text{BERT}} = \frac{2P_{\text{BERT}}R_{\text{BERT}}}{P_{\text{BERT}} + R_{\text{BERT}}} \quad (3)$$

This equation calculates the F1 score by computing the harmonic mean between precision and recall, offering a single metric that balances both. Table 6 summarizes the precision, recall, and F1 scores for the test sets, calculated using BERT Score. The results in Table 6 show that the GPT-2 model outperformed the Llama-2 model across all BERT Score

metrics, demonstrating higher precision, recall, and F1 scores in generating soil nutrient recommendations for cotton growth.

Table 6. BERT Score Results for GPT-2 and LLama2 Models

Model	Avg. Precision	Avg. Recall	Avg. F1-Score
GPT-2	0.92846	0.9308	0.9296
LLama-2	0.8446	0.8722	0.8581




Figure 8. Generated Report by GPT-2

Discussion:

In this research, we fine-tuned the GPT-2 and LLaMA-2 models to generate soil analysis reports based on field values. After evaluating the models using BERT Score, GPT-2 achieved a higher score, indicating that its generated reports were more contextually relevant. The results section provides a detailed breakdown of the evaluation metrics. Additionally, the generated reports, as shown in Figure 8, demonstrate that LLMs are capable of producing well-structured and informative reports. Traditionally, report generation is a manual process that involves analyzing soil nutrient levels and providing recommendations based on agricultural guidelines, a task that requires expertise and can be time-consuming. By automating this process, LLMs not only reduce the time taken but also offer a solution for real-time soil monitoring and decision-making. To the best of our knowledge, this is the first application of LLMs to this specific task, representing a novel contribution to the field of smart agriculture.

Conclusion and Future Work:

This study focuses on generating soil reports for cotton growth using Large Language Models (LLMs). We collected a dataset from PDF documents related to cotton and fine-tuned the LLM using a custom prompt template. The report generation process is iterative, where the

LLM generates responses based on nutrient values and the cotton growth stage, which are then merged to form the final report. The training loss of LLaMA-2 is 0.03, with an evaluation loss of 0.25. For the fine-tuned GPT-2 model, the training loss is 0.0932, and the evaluation loss is 0.086. We evaluated the performance of the fine-tuned GPT-2 model using BERT Score, achieving average Precision, Recall, and F1 scores of 0.9284, 0.9308, and 0.9296, respectively. These results demonstrate that the fine-tuned GPT-2 model performs effectively in generating accurate reports for cotton growth. Future work will involve expanding the dataset to include other related tasks, such as irrigation management for cotton, based on soil and environmental conditions, and further fine-tuning LLMs for these applications.

Acknowledgement: This work is supported by the Department of AI and Data Science, FAST-NUCES Karachi, Pakistan, and Kennesaw State University, Marietta, GA, USA.

Author's Contribution: Conceptualization, Muhammad Shahid; Formal analysis, Muhammad Shahid and M. Hassan, Syed Hassan Ali, and Abdul Rauf; Methodology, Syed Hassan Ali; Validation, Muhammad Shahid; Writing - original draft, Syed Hassan Ali and Abdul Rauf; Writing - review editing, M. Hassan.

Conflict of Interest: The authors declare no conflict of interest.

References:

- [1] S. Lilhare, P. Sanodiya, V. K. Dwivedi, and T. Singh, "Leveraging artificial intelligence (ai) for soil management: A comprehensive overview," *International Journal of Emerging Technologies and Innovative Research*, vol. 11, no. 4, pp. 230–238, 2024.
- [2] T. Pinthong and M. Ketcham, "The soil quality analysis using k-mean technique and model color," in *2022 International Conference on Cybernetics and Innovations (ICCI)*, pp. 1–4, IEEE, 2022.
- [3] S. Patel and H. M. Kumar, "Soil quality identifying and monitoring approach for sugarcane using machine learning techniques," in *2022 Fourth International Conference on Emerging Research in Electronics, Computer Science and Technology (ICERECT)*, pp. 1–5, IEEE, 2022.
- [4] A. J. Swetha, G. Kalyani, and B. Kirananjali, "Advanced soil fertility analysis and crop recommendation using machine learning," in *2023 7th International Conference on Trends in Electronics and Informatics (ICOEI)*, pp. 1035–1039, IEEE, 2023.
- [5] A. Banerjee, S. Mahajan, A. Rathore, S. Shiture, and M. Rajput, "Digi-farming assistant for soil quality analysis," in *2023 International Conference on Applied Intelligence and Sustainable Computing (ICAISC)*, pp. 1–4, IEEE, 2023.
- [6] M. Javaid, A. Haleem, I. H. Khan, and R. Suman, "Understanding the potential applications of artificial intelligence in agriculture sector," *Advanced Agrochem*, vol. 2, no. 1, pp. 15–30, 2023.
- [7] M. F. Celik, M. S. Isik, G. Taskin, E. Erten, and G. Camps-Valls, "Explainable artificial intelligence for cotton yield prediction with multisource data," *IEEE Geoscience and Remote Sensing Letters*, 2023.
- [8] F. M. Carneiro, A. L. de Brito Filho, F. M. Ferreira, G. d. F. S. Junior, Z. N. Brandao, R. P. da Silva, and L. S. Shiratsuchi, "Soil and satellite remote sensing variables importance using machine learning to predict cotton yield," *Smart Agricultural Technology*, vol. 5, p. 100292, 2023.
- [9] H. Chandra, P. M. Pawar, R. Elakkiya, P. S. Tamizharasan, R. Muthalagu, and A. Panthakkan, "Explainable ai for soil fertility prediction," *IEEE Access*, 2023.
- [10] D. Araci, "Finbert: Financial sentiment analysis with pre-trained language models. arxiv 2019," *arXiv preprint arXiv:1908.10063*.
- [11] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "Biobert: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
- [12] B. Silva, L. Nunes, R. Estevao, and R. Chandra, "Gpt-4 as an agronomist assistant? answering agriculture exams using large language models," *arXiv preprint arXiv:2310.06225*, 2023.
- [13] M. T. Kuska, M. Wahabzada, and S. Paulus, "Ai for crop production—where can large language models (llms) provide substantial value?," *Computers and Electronics in Agriculture*, vol. 221, p. 108924, 2024.

- [14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., "Language models are unsupervised multitask learners," OpenAI blog, vol. 1, no. 8, p. 9, 2019.
- [15] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al., "Llama 2: Open foundation and fine-tuned chat models," arXiv preprint arXiv:2307.09288, 2023.
- [16] X. Zhang, P. Yang, and B. Lu, "Artificial intelligence in soil management: The new frontier of smart agriculture," *Advances in Resources Research*, vol. 4, no. 2, pp. 231–251, 2024.
- [17] R. Priscilla, R. Deepa, and A. Pandi, "Agriculture-based automation with recommendation systems based on ai models," in *2023 Third International Conference on Artificial Intelligence and Smart Energy (ICAIS)*, pp. 1582–1589, IEEE, 2023.
- [18] K. Tan, "Large language models for crop yield prediction," 2024.
- [19] Y. Ding, Y. Ma, W. Fan, Y. Yao, T.-S. Chua, and Q. Li, "Fashionregen: Llm-empowered fashion report generation," in *Companion Proceedings of the ACM on Web Conference 2024*, pp. 991–994, 2024.
- [20] G. Colverd, P. Darm, L. Silverberg, and N. Kasmanoff, "Floodbrain: Flood disaster reporting by webbased retrieval augmented generation with an llm," arXiv preprint arXiv:2311.02597, 2023.
- [21] O. S. Sowole, "Leveraging large language models for improving agricultural extension in nigeria," in *Deep Learning Indaba 2023*, 2023.
- [22] A. Sharma, A. Sharma, A. Tselykh, A. Bozhenyuk, T. Choudhury, M. A. Alomar, and M. Sanchez-Chero, "Artificial intelligence and internet of things oriented sustainable precision farming: Towards modern agriculture," *Open Life Sciences*, vol. 18, no. 1, p. 20220713, 2023.
- [23] "Pakistan cottongrower." <https://www.ccri.gov.pk/cottongrower/PCGJul-Sep2020.pdf>. Accessed: 2024-05-30.
- [24] M. Malik, "Fertilizer role in sustainable cotton production," Place Published, pp. 20–29, 1998.
- [25] G. Stevens, "Cotton fertility management," 2019.
- [26] N. Khan, Y. Han, Z. Wang, G. Wang, L. Feng, B. Yang, and Y. Li, "Role of proper management of nitrogen in cotton growth and development," *International Journal of Biosciences*, vol. 14, no. 5, p. 5, 2019.
- [27] A. O. Abaye, "Potassium fertilization of cotton," 2019.
- [28] J. L. Oldham, D. M. Dodds, W. H. McCarty, et al., "Inorganic nutrient management for cotton production in mississippi," 2010.
- [29] F. Ahmad, M. Akhtar, et al., "Phosphorus application strategies to improve cotton productivity under arid climatic conditions," *International Journal of Cotton Research and Technology*, vol. 56, no. 1-4, pp. 14– 23, 2020.
- [30] "Advancing cotton education." <https://www.cotton.org/tech/ace/soil-fertility.cfm>. Accessed: 2024-05-30.
- [31] A. Imran and K. Noureen, "Cotton crop development in central punjab (faisalabad, 2019)," 2019.
- [32] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- [33] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "Bertscore: Evaluating text generation with bert," arXiv preprint arXiv:1904.09675, 2019.

Copyright © by authors and 50Sea. This work is licensed under Creative Commons Attribution 4.0 International License.

