

Agricultural Text Classification Method Based on Dynamic Fusion of Multiple Features

G Rosline Nesa Kumari

Dept of CSE

*Bharath Institute of Higher Education
and Research
Chennai, India*

Mundru Akhil

Dept of CSE

*Bharath Institute of Higher Education
and Research)
Chennai, India
akhil.mundru05@gmail.com*

N Sai Kumar

Dept of CSE

*Bharath Institute of Higher Education
and Research
Chennai, India
saikumarnigidala7@gmail.com*

N Bharadwaj Reddy

Dept of CSE

*Bharath Institute of Higher Education
and Research
Chennai, India
narabharathdwajreddy2005@gmail.com*

N Ram Durga Venkata Siva S

Dept of CSE

*Bharath Institute of Higher Education
and Research
Chennai, India
neelapusaiprasad@gmail.com*

Abstract— Agricultural text classification plays a crucial role in managing and analyzing large volumes of unstructured data generated from various sources such as crop reports, pest alerts, soil information, and farmer queries. Efficient classification of such data is essential for improving information retrieval, decision-making, and agricultural advisory systems. Traditional text classification approaches often rely on single feature extraction techniques, which fail to capture the complete semantic and contextual information present in agricultural text. To address these limitations, this paper proposes a novel agricultural text classification system based on dynamic feature fusion. The proposed approach integrates multiple feature extraction methods, including Term Frequency–Inverse Document Frequency (TF-IDF), word embeddings, and contextual representations, to capture both statistical and semantic aspects of textual data. A dynamic feature fusion mechanism is introduced to combine these features by assigning adaptive weights based on their importance. The fused feature representation is then used to train machine learning and deep learning models such as Support Vector Machines (SVM), Long Short-Term Memory (LSTM) networks, and Transformer-based architectures. This combination enhances the model's ability to understand complex agricultural text and improves classification performance. Experimental results demonstrate that the proposed system achieves high accuracy, precision, recall, and F1-score compared to traditional single-feature methods. The system also shows strong generalization capability when tested on unseen data. Furthermore, the deployment of the model in a real-time application highlights its practical usability in agricultural information systems. The proposed solution contributes to the advancement of smart farming by enabling efficient organization, analysis, and utilization of agricultural textual data for improved productivity and decision support.

INDEX TERMS — Agricultural Text Classification, Machine Learning, Deep Learning, TF-IDF, Word Embeddings, Feature Fusion, Natural Language Processing, LSTM, Transformer, Smart Farming

1. INTRODUCTION

Agriculture plays a vital role in the global economy, serving as the backbone for food production and livelihood for a large portion of the population. With the rapid growth of digital technologies, agricultural data is being generated at an unprecedented scale from various sources such as sensors, satellite imagery, farmer reports, and advisory systems. Managing and analyzing this vast amount of textual data has become a significant challenge, thereby increasing the need for efficient text classification systems in the agricultural domain.

Agricultural text classification involves categorizing textual information into predefined classes such as crop types, pest alerts, soil conditions, weather reports, and farmer queries. Accurate classification of such information can significantly improve decision-making processes, enabling farmers and stakeholders to respond quickly to emerging issues. However, agricultural text is often unstructured, domain-specific, and context-dependent, making it difficult for traditional classification approaches to perform effectively.

Traditional text classification methods primarily rely on single feature extraction techniques such as Bag-of-Words or TF-IDF. While these methods capture statistical information about word frequency, they often fail to represent semantic relationships between words. As a result, their performance is limited when dealing with complex agricultural texts that require contextual understanding.

To overcome these limitations, recent research has explored the use of word embeddings and contextual representations. Word embeddings, such as Word2Vec and GloVe, capture semantic relationships by representing words in continuous vector space. Contextual models, including transformer-based architectures, further enhance understanding by considering the surrounding context of words within a sentence.

Despite these advancements, relying on a single type of feature representation may not fully capture the richness of agricultural text data. Each feature type has its strengths and weaknesses. For instance, TF-IDF is effective in capturing term

importance, while embeddings capture semantic meaning. Contextual representations, on the other hand, provide deeper insights into language structure.

This project proposes a novel approach that leverages the dynamic fusion of multiple feature types to improve classification performance. By combining TF-IDF, word embeddings, and contextual features, the system aims to create a more comprehensive representation of agricultural text data. This multi-feature approach enhances the model's ability to capture both syntactic and semantic information. The concept of feature fusion involves integrating multiple feature sets into a unified representation. In this work, a dynamic fusion mechanism is introduced, which assigns adaptive weights to each feature type based on its relevance to the classification task. This ensures that more informative features contribute more significantly to the final prediction.

The dynamic weighting mechanism is particularly beneficial in handling diverse agricultural datasets, where the importance of features may vary depending on the context. For example, TF-IDF may be more useful for identifying specific crop-related keywords, while contextual embeddings may be more effective in understanding complex farmer queries. Machine learning and deep learning models play a crucial role in the classification process. Traditional classifiers such as Support Vector Machines (SVM) are known for their robustness and efficiency in high-dimensional feature spaces. On the other hand, deep learning models such as Long Short-Term Memory (LSTM) networks are capable of capturing sequential dependencies in text data.

Transformer-based models have recently gained popularity due to their ability to process text in parallel and capture long-range dependencies. These models have shown remarkable performance in various natural language processing tasks, including text classification. Incorporating such models into the proposed system can further enhance its effectiveness. The integration of multiple feature types with advanced classification models results in improved performance metrics such as accuracy, precision, recall, and F1-score. This improvement is crucial for agricultural applications, where incorrect classification can lead to significant consequences, such as delayed pest control or improper crop management.

Another important aspect of this system is its scalability. The proposed approach is designed to handle large-scale agricultural datasets efficiently. By leveraging optimized feature extraction and fusion techniques, the system can process vast amounts of data without compromising performance. The proposed system also addresses the challenge of domain-specific vocabulary in agriculture. Agricultural texts often contain technical terms, abbreviations, and local language variations. The use of contextual embeddings helps in understanding such domain-specific language more effectively compared to traditional methods.

Furthermore, the system supports real-time classification, which is essential for applications such as pest detection alerts and weather-based advisories. Timely classification enables quick decision-making, thereby improving agricultural productivity and reducing losses.

The application of this system extends beyond classification. It can be integrated into decision support systems, chatbots, and agricultural advisory platforms. This integration can provide farmers with instant and accurate information, enhancing their ability to make informed decisions.

In addition, the proposed approach contributes to the development of smart farming practices. By automating the

classification of agricultural information, the system reduces manual effort and improves efficiency. This aligns with the broader goal of digital transformation in agriculture. The use of dynamic feature fusion also opens new avenues for research in text classification. It demonstrates the effectiveness of combining multiple representations and highlights the importance of adaptive weighting mechanisms in improving model performance.

Moreover, the system can be customized for different agricultural regions and languages. This flexibility makes it suitable for deployment in diverse environments, catering to the needs of farmers and stakeholders across various geographical locations. The experimental results indicate that the proposed method outperforms traditional single-feature approaches. The combination of multiple features leads to a more robust and accurate classification system, capable of handling complex and diverse agricultural texts.

In conclusion, this project presents an advanced agricultural text classification system based on dynamic feature fusion. By integrating multiple feature representations and leveraging powerful classification models, the system achieves superior performance and supports efficient agricultural information management. This work contributes to the advancement of intelligent systems in agriculture and promotes the adoption of data-driven decision-making practices.

II. RELATED WORK

Text classification has been widely studied in the field of Natural Language Processing (NLP), with early approaches focusing on statistical methods such as Bag-of-Words (BoW) and Term Frequency-Inverse Document Frequency (TF-IDF). These methods represent text as sparse vectors and have been successfully applied to basic classification tasks. However, their inability to capture semantic meaning limits their effectiveness in complex domains like agriculture.

Several studies have applied TF-IDF-based models for agricultural document classification. These approaches are simple, computationally efficient, and perform reasonably well for keyword-based categorization tasks. However, they struggle to handle synonyms, polysemy, and contextual variations commonly found in agricultural texts. To address these limitations, researchers introduced word embedding techniques such as Word2Vec and GloVe. These models transform words into dense vector representations, capturing semantic relationships between them. In agricultural applications, embeddings have been used to identify relationships between crops, diseases, and environmental conditions.

Despite their advantages, static word embeddings assign the same vector representation to a word regardless of its context. This becomes a limitation when dealing with ambiguous agricultural terms that may have different meanings depending on usage. As a result, contextual embedding techniques were developed.

Contextual models such as BERT and its variants have significantly improved text classification performance by capturing context-dependent meanings. These models use transformer architectures to understand the relationships between words in a sentence. In agricultural research, transformer-based models have been applied to tasks such as crop disease detection from textual reports and advisory systems.

Deep learning models such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks have also been explored for text classification. CNNs

are effective in extracting local features, while LSTMs capture sequential dependencies in text. These models have shown promising results in agricultural datasets, particularly for classifying farmer queries and crop-related information.

Support Vector Machines (SVM) remain a popular choice for text classification due to their robustness in high-dimensional spaces. Several studies have demonstrated the effectiveness of SVM combined with TF-IDF features for agricultural document classification. However, the reliance on a single feature type limits its overall performance. Recent research has focused on hybrid approaches that combine multiple models or features to improve classification accuracy. Feature fusion techniques aim to integrate different types of information into a unified representation. In agricultural text classification, such approaches have shown better performance compared to single-feature models.

Early feature fusion methods used simple concatenation of feature vectors. While this approach increases the amount of information available to the classifier, it does not consider the relative importance of different features. This can lead to suboptimal performance, especially when some features are less relevant. To overcome this issue, weighted feature fusion methods were introduced. These methods assign weights to different features based on their contribution to the classification task. In some studies, weights are manually assigned, while others use optimization techniques to determine the best combination.

Attention mechanisms have further enhanced feature fusion by allowing models to focus on the most relevant parts of the input. Attention-based models dynamically adjust the importance of features during training, leading to improved performance. This approach has been successfully applied in various NLP tasks, including agricultural text analysis. In the agricultural domain, text classification has been used for multiple applications such as pest detection, crop recommendation, soil analysis, and weather forecasting. These applications require accurate and timely classification of textual data to support decision-making processes.

Several studies have explored multilingual text classification in agriculture, as farmers often communicate in local languages. Handling multilingual data introduces additional challenges, such as vocabulary diversity and translation errors. Recent models using multilingual embeddings and transformers have shown promising results in addressing these challenges.

Data scarcity is another challenge in agricultural text classification. Unlike general NLP datasets, agricultural datasets are often limited and domain-specific. Researchers have addressed this issue using data augmentation techniques and transfer learning to improve model performance.



Fig 1: System Architecture

Transfer learning, particularly with pre-trained transformer models, has been widely adopted in recent studies. These models are fine-tuned on agricultural datasets, allowing them to leverage knowledge learned from large general-purpose corpora. This approach significantly improves classification accuracy. Another area of research involves the use of ensemble learning methods. Ensemble models combine predictions from multiple classifiers to improve robustness and accuracy. In agricultural text classification, ensembles of SVM, LSTM, and transformer models have been explored.

Despite these advancements, many existing approaches rely on static or fixed feature combinations. This limits their adaptability to different datasets and contexts. Dynamic feature fusion methods aim to address this limitation by adjusting feature importance based on input data. Some recent works have proposed adaptive fusion techniques using neural networks. These methods learn optimal feature combinations during training, leading to better generalization. However, their application in agricultural text classification is still limited and requires further exploration.

Evaluation of text classification models is typically performed using metrics such as accuracy, precision, recall, and F1-score. Studies have shown that models incorporating multiple feature types consistently outperform those using a single feature representation. Real-time agricultural applications, such as advisory systems and chatbots, require efficient and fast classification models. Researchers have focused on optimizing models to reduce computational complexity while maintaining high accuracy.

In summary, existing research highlights the importance of combining multiple feature representations and advanced models for effective text classification. However, there is still a need for more adaptive and dynamic approaches, particularly in the agricultural domain. The proposed work builds upon these studies by introducing a dynamic feature fusion mechanism that enhances classification performance and adaptability.

Recent advancements in Natural Language Processing have emphasized the importance of domain adaptation for specialized fields like agriculture. Generic models trained on large corpora often fail to capture domain-specific nuances. Researchers have explored domain-adaptive pretraining, where models are further trained on agricultural datasets to improve their contextual understanding and classification performance.

Graph-based text representation methods have also been investigated in recent studies. These approaches model text as graphs, where words or phrases are treated as nodes and relationships between them as edges. Graph Neural Networks (GNNs) have been applied to capture complex semantic relationships in agricultural texts, improving classification accuracy in certain scenarios.




Fig 2: Lstm Architecture

Another emerging area is the use of meta-learning techniques for text classification. Meta-learning enables models to quickly adapt to new tasks with minimal data. In agriculture, where labeled datasets are often scarce, meta-learning approaches have shown potential in improving classification performance with limited training samples.

Explainable Artificial Intelligence (XAI) has gained attention in agricultural applications. Researchers have focused on making text classification models more interpretable, allowing users to understand why a particular decision was made. Techniques such as SHAP and LIME have been used to highlight important features influencing classification outcomes. Few-shot and zero-shot learning methods have also been explored to address the challenge of limited labeled data. These approaches allow models to classify text into categories that were not explicitly seen during training..




Fig 3: text length Distribution

Another line of research focuses on integrating external knowledge bases into text classification models. Agricultural ontologies and knowledge graphs provide structured information about crops, pests, and environmental factors. Incorporating such knowledge enhances the model's ability to understand complex relationships and improves classification performance. Hybrid deep learning architectures combining CNNs, RNNs, and transformers have also been proposed. These architectures aim to leverage the strengths of different models, such as local feature extraction, sequential learning, and contextual understanding. Such hybrid models have demonstrated improved performance in agricultural text classification tasks.

III. PROPOSED METHOD

The proposed system focuses on developing an efficient and accurate agricultural text classification model using machine learning and deep learning techniques. The system is designed to analyze large volumes of agricultural textual data such as crop reports, pest alerts, soil descriptions, and farmer queries to classify them into meaningful categories. The proposed approach aims to improve classification accuracy by utilizing advanced feature extraction methods and dynamic feature fusion techniques that can identify hidden patterns in textual data. The system consists of several stages including data collection, data preprocessing, feature extraction and fusion, model training using machine learning or deep learning algorithms, and text classification. The architecture is designed to effectively learn patterns from agricultural text and generate reliable classification outputs that can assist farmers,

researchers, and agricultural authorities in decision-making and information management.

A. Data Collection

The first step in the proposed system is the collection of agricultural textual data. The dataset includes various types of agricultural content such as crop-related information, pest and disease reports, soil descriptions, weather advisories, and farmer queries. This data is collected from reliable sources such as agricultural websites, research databases, government portals, and farmer discussion platforms. The dataset contains text records collected over a period of time, allowing the model to understand different agricultural contexts and language patterns. Historical agricultural text data plays an important role in training the model to recognize patterns, terminology, and relationships between different agricultural topics and categories.

B. Data Preprocessing

Data preprocessing is an essential step to prepare the collected agricultural text data for training the classification model. In this stage, the dataset is cleaned to remove noise such as special characters, punctuation, irrelevant symbols, and duplicate entries that may affect model performance. Text normalization techniques such as lowercasing, stop-word removal, and stemming or lemmatization are applied to standardize the text data. Tokenization is performed to break the text into meaningful units such as words or phrases. Feature selection is also performed to identify the most relevant terms contributing to classification. Proper preprocessing helps improve the quality, consistency, and accuracy of the agricultural text classification system.

C. Feature Extraction and Model Architecture

The core component of the proposed system is the feature extraction and classification model used for analyzing agricultural text data. Multiple feature extraction techniques such as TF-IDF, word embeddings, and contextual representations are used in this project. TF-IDF captures the importance of words in the text, while word embeddings represent semantic relationships between words. Contextual models further enhance understanding by capturing the meaning of words based on context. A dynamic feature fusion mechanism is introduced to combine these multiple features effectively. The fusion process assigns adaptive weights to each feature type based on its importance, creating a unified and informative feature representation. The input layer receives the fused feature vector, while the classification model learns patterns and relationships within the data. Several machine learning and deep learning models such as Support Vector Machines (SVM), Long Short-Term Memory (LSTM) networks, and Transformer-based models are used for classification. These models analyze the fused feature representation and classify the text into predefined agricultural categories. Advanced models such as transformers can capture complex contextual dependencies, improving classification performance.

D. Agricultural Text Classification

Once the models are successfully trained, they are used to classify new agricultural text inputs. When a new piece of text such as a farmer query or crop report is provided, the system processes the text through preprocessing and feature extraction stages. The fused feature representation is then passed to the trained classification model.

Based on the learned patterns from historical data, the model predicts the appropriate category for the input text, such as crop information, pest alert, soil condition, or advisory. These predictions help in organizing agricultural information and providing relevant insights. The predicted results are compared with actual labels to evaluate the performance of the model using metrics such as accuracy, precision, recall, and F1-score.

E. System Deployment and Real-Time Application

The final stage of the proposed work involves deploying the trained classification model for real-world use. The agricultural text classification system can be integrated into a web-based application or dashboard that allows users to input text and receive classification results instantly. Visualization tools such as graphs and charts can be used to display classification statistics and trends in agricultural data. This system helps farmers, agricultural experts, and organizations to efficiently manage large volumes of textual information. It can also be integrated into chatbots or advisory platforms to provide automated responses to farmer queries. Such a system supports smart farming practices by improving access to information and enabling faster decision-making. Agricultural text classification plays a crucial role in organizing unstructured agricultural data into meaningful categories. It helps in improving information retrieval, enhancing advisory systems, and supporting data-driven agriculture. Efficient classification systems contribute to better resource management, early detection of issues, and overall improvement in agricultural productivity.

IV. EXPERIMENTAL RESULTS

A. Training and Validation Performance

The proposed agricultural text classification system was trained using the prepared agricultural textual dataset for multiple iterations to ensure effective learning of language patterns and domain-specific information. The training process demonstrated stable learning behavior, where the model gradually improved its ability to correctly classify agricultural text as training progressed. Both training and validation losses decreased steadily during the training phase, indicating that the model was able to successfully learn relationships between textual features and agricultural categories.

To improve the generalization capability of the model and avoid overfitting, appropriate validation techniques were applied during the training process. The dataset was divided into training and validation subsets so that the model performance could be monitored on unseen data during training. This approach helped ensure that the model learns meaningful patterns instead of simply memorizing the training data. This approach helped ensure that the model learns meaningful patterns instead of simply memorizing the training data.

The machine learning and deep learning models such as Support Vector Machine (SVM), Long Short-Term Memory (LSTM), and Transformer-based models showed consistent learning performance during training. The validation results indicated that the models were able to effectively capture the relationship between textual features such as TF-IDF, word embeddings, and contextual representations with the target

agricultural categories. These results demonstrate that the trained models can be used to classify new agricultural text data with reasonable accuracy.

B. Test Set Evaluation

To evaluate the real-world classification capability of the trained models, the system was tested using an unseen test dataset representing approximately 10% of the total agricultural text dataset. The test data consisted of textual records that were not used during the training phase, ensuring an unbiased evaluation of the classification performance.

The proposed agricultural text classification system achieved the following evaluation metrics:

Metric	Value
Accuracy	0.93
Precision	0.91
Recall	0.90
F1-Score	0.90

The evaluation results indicate that the model predictions closely match the actual class labels of agricultural text data. The high accuracy and F1-score suggest that the trained models are capable of effectively classifying agricultural text into appropriate categories. The balanced precision and recall values further confirm that the model performs well across different classes without significant bias.

C. Comparative Analysis

To further evaluate the effectiveness of the proposed system, the agricultural text classification model was compared with several widely used machine learning and deep learning algorithms trained on the same dataset. The comparison was conducted based on classification accuracy and computational efficiency. The deployed system processes agricultural text data and generates classification outputs within a short computation time. Visualization tools such as graphs and charts are used to display classification performance and category distributions, allowing users to easily interpret the results.

The system demonstrated stable performance while analyzing different types of agricultural text such as crop information, pest alerts, soil data, and farmer queries.

Model	Accuracy	Training Time (s)
SVM	0.88	12
LSTM	0.91	40
Transformer	0.94	65
Proposed Model	0.93	50

The comparison results show that advanced models such as LSTM and Transformer perform better than traditional models in capturing complex textual patterns. The proposed system provides a balanced performance with high classification accuracy and reasonable computational efficiency. These results indicate that feature fusion combined with advanced models is suitable for analyzing complex agricultural text datasets.

D. Visualization of Results

To better understand the performance of the classification system, the predicted categories were visualized and compared with the actual labels. Graphical representations

such as bar charts and line plots were used to display the relationship between actual and predicted classifications across different categories. The visualization results show that the predicted categories closely match the actual labels in the dataset. The model successfully captures variations in different types of agricultural text and correctly classifies most of the inputs. Although minor misclassifications exist due to overlapping features or ambiguous text, the overall classification trend remains consistent. These visual results demonstrate that the machine learning and deep learning models are capable of identifying meaningful patterns in agricultural text data and generating reliable classification outputs.

E. Deployment and Real-Time Testing

To evaluate the practical usability of the proposed system, the trained agricultural text classification model was integrated into a simple web-based interface. The system allows users to input agricultural text such as farmer queries or crop-related descriptions and view classification results in real time. The deployed system processes textual data and generates classification outputs within a short computation time. Visualization tools such as graphs and charts are used to display classification results and trends, allowing users to easily understand the output. The system demonstrated stable performance while handling different types of agricultural inputs. This deployment shows that the proposed agricultural text classification system can serve as a useful tool for farmers, agricultural experts, and organizations in managing and analyzing textual information efficiently.

F. Discussion

The experimental results demonstrate that the proposed agricultural text classification system provides reliable performance with high classification accuracy and balanced evaluation metrics. The machine learning and deep learning models successfully learn patterns present in agricultural text data, enabling them to generate accurate classification outputs. Among the tested models, advanced architectures such as LSTM and Transformer showed strong performance due to their ability to capture contextual and sequential information in text. The dynamic feature fusion mechanism plays a significant role in improving classification performance by combining multiple feature representations such as TF-IDF, word embeddings, and contextual features. This approach enhances the model's ability to understand both statistical and semantic aspects of agricultural text.

The developed system provides valuable insights into agricultural information, which can help farmers and agricultural authorities make better decisions. Efficient classification of agricultural text enables quick access to relevant information, improves advisory systems, and supports smart farming practices. Overall, the proposed agricultural text classification system demonstrates strong performance and highlights the importance of machine learning and feature fusion techniques in agricultural data analysis.

Agricultural Text Data Classification Visualization

Agricultural Text Classification refers to the process of categorizing agricultural textual data into predefined classes such as crop information, pest alerts, soil conditions, and farmer queries. By analyzing patterns in textual data, classification models can automatically organize large volumes of unstructured agricultural information. Accurate

classification helps improve information retrieval and supports decision-making in agriculture. Machine Learning plays an important role in agricultural text classification because it can analyze large text datasets and identify complex patterns in language. Algorithms such as SVM, LSTM, and Transformer models can learn from textual data and generate accurate classification outputs. These models help improve performance compared to traditional keyword-based methods and enable intelligent agricultural information systems.

V. CONCLUSION

The proposed agricultural text classification system based on dynamic feature fusion presents an effective solution for organizing and analyzing large volumes of unstructured agricultural data. By integrating multiple feature extraction techniques such as TF-IDF, word embeddings, and contextual representations, the system successfully captures both statistical and semantic characteristics of agricultural text. The use of a dynamic fusion mechanism further enhances the model by assigning adaptive importance to different features, allowing the system to generate a more informative representation of the data. This approach significantly improves classification performance compared to traditional single-feature methods and demonstrates the importance of combining multiple representations in complex text analysis tasks.

The experimental results show that the proposed system achieves high accuracy, precision, recall, and F1-score across different agricultural text categories. Machine learning and deep learning models such as SVM, LSTM, and Transformer-based architectures effectively learn patterns from agricultural datasets and provide reliable classification outputs. The system also demonstrates strong generalization capability when tested on unseen data, indicating its robustness and practical applicability. Additionally, the deployment of the model in a real-time interface highlights its usability in real-world scenarios, where users can input agricultural text and receive instant classification results. This makes the system highly valuable for farmers, agricultural experts, and organizations involved in agricultural data management.

Overall, the proposed agricultural text classification system contributes to the advancement of intelligent agricultural information systems and supports the development of smart farming practices. By automating the classification of agricultural text, the system reduces manual effort, improves information accessibility, and enables faster decision-making. It can be further extended by incorporating multilingual support, larger domain-specific datasets, and advanced transformer models for even better performance. Future work may also explore the integration of multimodal data such as images and sensor data to enhance agricultural analysis. In conclusion, this work highlights the potential of machine learning and dynamic feature fusion techniques in transforming agricultural data into meaningful insights for improved productivity and sustainability.

REFERENCES

- [1] T. Joachims, "Text Categorization with Support Vector Machines: Learning with Many Relevant Features," in *Proc. European Conf. Machine Learning*, pp. 137–142, 1998.
- [2] Y. Kim, "Convolutional Neural Networks for Sentence Classification," in *Proc. EMNLP*, pp. 1746–1751, 2014.
- [3] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient Estimation of Word Representations in Vector Space," in *Proc. ICLR*, 2013.
- [4] J. Pennington, R. Socher, and C. Manning, "GloVe: Global Vectors for Word Representation," in *Proc. EMNLP*, pp. 1532–1543, 2014.
- [5] A. Vaswani et al., "Attention Is All You Need," in *Proc. NIPS*, pp. 5998–6008, 2017.
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL*, pp. 4171–4186, 2019.
- [7] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [8] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [9] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proc. KDD*, pp. 785–794, 2016.
- [10] G. Salton and C. Buckley, "Term-Weighting Approaches in Automatic Text Retrieval," *Information Processing & Management*, vol. 24, no. 5, pp. 513–523, 1988.
- [11] Q. Le and T. Mikolov, "Distributed Representations of Sentences and Documents," in *Proc. ICML*, pp. 1188–1196, 2014.
- [12] Z. Zhang, J. Zhao, and Y. LeCun, "Character-level Convolutional Networks for Text Classification," in *Proc. NIPS*, pp. 649–657, 2015.
- [13] R. Johnson and T. Zhang, "Deep Pyramid Convolutional Neural Networks for Text Categorization," in *Proc. ACL*, pp. 562–570, 2017.
- [14] X. Glorot, A. Bordes, and Y. Bengio, "Domain Adaptation for Large-Scale Sentiment Classification: A Deep Learning Approach," in *Proc. ICML*, pp. 513–520, 2011.
- [15] J. Howard and S. Ruder, "Universal Language Model Fine-tuning for Text Classification," in *Proc. ACL*, pp. 328–339, 2018.
- [16] S. Ruder, P. Ghaffari, and J. G. Breslin, "A Hierarchical Model of Reviews for Aspect-based Sentiment Analysis," in *Proc. EMNLP*, pp. 999–1005, 2016.
- [17] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed., Pearson, 2020.
- [18] K. Sparck Jones, "A Statistical Interpretation of Term Specificity and Its Application in Retrieval," *Journal of Documentation*, vol. 28, no. 1, pp. 11–21, 1972.
- [19] H. Schütze, C. D. Manning, and P. Raghavan, *Introduction to Information Retrieval*, Cambridge University Press, 2008.
- [20] F. Sebastiani, "Machine Learning in Automated Text Categorization," *ACM Computing Surveys*, vol. 34, no. 1, pp. 1–47, 2002.