Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 00 (2019) 000–000

Procedia Computer Science

www.elsevier.com/locate/procedia

The 22nd International Conference on Mobile Systems and Pervasive Computing (MobiSPC)
August 4-6, 2025, Leuven, Belgium

Exploring Explainable AI Techniques for Plant Disease Classification in Digital Agriculture

Rui Pedro Porfírio[a,b,*], Rui Neves Madeira[a,b], Pedro Albuquerque Santos[a,b]

[a] *NOVA LINCS, NOVA School of Science and Technology, NOVA University of Lisbon, Caparica, Portugal*
[b] *RESILIENCE, Setúbal School of Technology, Polytechnic Institute of Setúbal, Setúbal, Portugal*

Abstract

Integrating smart farming technologies in agriculture is crucial to address the pressing challenges of food security, economic stability, and environmental sustainability. Solutions based on artificial intelligence (AI) for plant disease detection play a critical role in optimizing crop health and yield. However, the complexity and opacity of these AI models can hinder their acceptance and practical application by end-users. To address this issue, our ongoing research explores applying explainable AI (XAI) techniques to enhance the explainability of vision-based plant disease classification models. Our experiments assess the transparency of vision-based models by applying XAI techniques, such as LIME, Grad-CAM, and occlusion-based attribution, to visualize the reasoning behind model predictions. Additionally, we analyze inherently transparent machine learning models, such as k-Nearest Neighbors, through custom visualization graphics and examine how explainability varies with model accuracy. These findings highlight the role of XAI techniques in enhancing the transparency of predictions for crucial vision tasks in agriculture such as plant disease classification. Building on these results, we explore the potential integration of XAI techniques into third-party applications or prototypes, emphasizing how tailored visual and textual explanations can enhance the transparency and interpretability of plant disease classification models.

Keywords: Plant Disease Detection; Digital Agriculture; Explainable AI; Machine Learning; Computer Vision

1. Introduction

The agricultural ecosystem is evolving towards a new paradigm fueled by digitalization and innovation [24]. According to the Food and Agriculture Organization of the United Nations (FAO), digital agriculture - the application of information and communication technologies (ICT) to agricultural contexts - has the potential to deliver significant benefits throughout the entire agrifood value chain and become a key driver of sustainability [8]. Among the applica-

* Corresponding author. Tel.: +351 935 478 708
E-mail address: r.porfirio@campus.fct.unl.pt

tions of digital agriculture, advances in artificial intelligence (AI) have contributed significantly to the sector, offering solutions for crop stress detection, weather predictions, and pesticide recommendations [9, 3]. In terms of food safety solutions, computer vision-based approaches aim to prevent critical farm issues such as pest and disease outbreaks. This issue substantially impacts smallholder farmers, who represent 84% of farms worldwide [14], with some farmers reporting up to more than 50% of yield loss due to pests and diseases [10].

Although computer vision-based approaches for plant disease classification have shown promising results, the complexity and opacity of these AI models can hinder their acceptance and practical application by end-users, such as farmers [13]. Lack of transparency in how these models make predictions can lead to a lack of trust, limiting their widespread adoption in digital agriculture. Explainable AI (XAI) techniques offer a solution by providing insights into the reasoning behind the models' decisions, fostering greater trust and perceived value among agricultural stakeholders [7]. By making the decision-making process of AI systems more interpretable and understandable, XAI can help bridge the gap between digital technologies and the practical needs of users in the agricultural domain. This, in turn, can ease the integration of smart farming solutions and accelerate the transition toward more sustainable and data-driven agricultural practices.

Motivated by the aforementioned, this paper explores the application of explainable AI techniques to enhance the interpretability[1] of vision-based plant disease classification models. Through a set of experiments, we evaluate multiple XAI techniques (LIME, Grad-CAM, and occlusion-based attributions) across different deep learning architectures (ViT B-32 and ResNet-50), providing insights into their comparative strengths. Additionally, we conduct a preliminary analysis of how model accuracy influences explanation reliability. To further demonstrate the practical value of these findings, we discuss the potential integration of XAI techniques into applications or prototypes designed for plant disease classification.

The remainder of this paper is structured as follows. Section 2 presents the research context for deep learning applications in plant disease detection and previous work on the application of XAI techniques in digital agriculture. Section 3 details the methodology of our experiments, including the dataset used and the design of each experiment. Section 4 presents the results of these experiments, highlighting the potential of XAI to provide transparent insights into plant disease classification. Moreover, it briefly discuss how XAI explanations could be integrated into third-party applications or prototypes. Finally, section 5 concludes the paper and outlines future directions for integrating XAI techniques in digital agriculture vision tasks.

2. Related Work

As the field of digital agriculture continues to expand, the capacity to detect diseases within agricultural ecosystems has emerged as a compelling research challenge, given its significance for ensuring effective pest management systems. Alongside biological approaches such as DNA and serological techniques [15], there have been numerous endeavours to develop automated systems capable of identifying diseases in existing crops, relying primarily on visual imagery.

In recent years, deep learning algorithms, particularly convolutional neural networks (CNNs), have significantly improved the performance of plant disease detection and classification models [4]. These data-driven approaches have demonstrated superior capabilities in extracting complex visual features and patterns from crop images, outperforming traditional machine learning methods. However, the growing novelty of applying CNNs to object recognition and image classification tasks has prompted extensive research exploring their performance for plant disease classification tasks. While some studies have proposed original CNN architectures to solve the classification task with a significant degree of confidence [18], many researchers have focused on reusing existing CNN architectures, such as those competing on the *ImageNet Large Scale Visual Recognition Challenge* (e.g., AlexNet, ResNet, GoogLeNet) [16, 19].

Despite the opportunities of the application of CNNs for plant disease classification, the literature highlights several inherent dependencies and challenges associated with this task. A key requirement for effective supervised learning models is the availability of large training datasets. However, the substantial effort needed to construct expert-verified

[1] In this paper, we interchangeably use interpretability and explainability

datasets with manually annotated images represents a significant obstacle to the widespread adoption of CNNs for disease detection, given the scarcity of robust training datasets and the potential for errors in data labeling [12].

While CNNs have been the dominant approach for image classification tasks, the emergence of transformers has introduced a promising new paradigm for a wide range of computer vision applications. Introduced in 2020 by a team of Google researchers, vision transformers (ViT) leverage the core principles of the transformer architecture, which has demonstrated remarkable success in natural language processing, and apply them to visual data [6]. Regarding the application of transformers in digital agriculture, the work of Reedha et al. claims to provide the first study on the potential of vision transformers for classifying weed and crop images from unmanned aerial vehicles (UAVs) [21]. The authors evaluated the performance of two vision transformer models, ViT B-16 and ViT B-32, against state-of-the-art CNN architectures, such as ResNet and EfficientNet, using a custom dataset with five classes: weed, beet, off-type beet, parsley, and spinach.

While computer vision-based approaches have shown promising results, the complexity and opacity of deep learning models can hinder their acceptance and practical application by end-users. To address this challenge, there has been a growing interest in the field of XAI, which aims to provide transparency and interpretability for AI-powered systems [1]. Particularly, XAI techniques can play a crucial role in increasing transparency and user trust in computer vision models. By generating explanations such as visual insights, these methods help users understand the decision-making process of the model, leading to greater confidence in the predictions. XAI techniques like LIME [22], Grad-CAM [23], and attribution methods like occlusion-based attributions provide valuable insights into the model's predictions. LIME creates locally interpretable models to identify the features that contribute significantly to a particular prediction, using superpixels to segment the input image and determine the most important regions. Grad-CAM leverages gradient information to highlight the important regions in the input image that drive the model's output. Additionally, occlusion-based attribution methods systematically occlude different parts of the input image and observe the impact on the model's output, thereby allowing the identification of regions critical for the classifier's decision.

Previous research on applying XAI techniques to plant disease classification tasks has been conducted by Wei et al., who examined how explainability can be explored in agricultural fruit disease classification [25]. Their work investigated the use of explainable techniques, such as SmoothGrad, LIME, and GradCAM, to understand if an attention-based ResNet network can enhance the feature extraction ability of CNN architectures. While all three techniques provided interesting insights, the authors concluded that GradCAM exhibited the strongest explanatory power, with LIME being the least explanatory method. Furthermore, the authors found that the attention module, when combined with the explainability methods, positively influenced the feature extraction process in disease classification tasks.

3. Materials and Methods

This section describes the dataset used in the experiments and the specific design of each conducted experiment.

3.1. Dataset

The *PlantVillage* dataset is considered one of the first robust training datasets for plant disease detection [11]. Its initial version contains over 54,309 images validated by expert pathologists and divided into 38 classes. This dataset represents an important step towards defining expert-validated datasets with substantial dimensionality and distribution. Notably, the version of the dataset used in related studies differs. While the work of Mohanty et al. [16] utilized the original format of the *PlantVillage* dataset, Brahimi et al. [2] applied a different version with 14,828 images and nine classes.

This study employs an augmented version of the dataset with 61,486 images and 39 classes. The key difference between the original dataset and this augmented version is the incorporation of pre-generated augmented data. The data augmentation process involved applying various random transformations to increase the number of training images artificially. This approach helped create a more diverse and representative dataset, mitigating the risk of overfitting. Additionally, the dataset includes a new class for non-crop images. While each class can distinguish both the plant and the presence or absence of disease, there is a significant imbalance in the data, with some classes having as few as 700 samples and others up to 3,854. Furthermore, an analysis of the data distribution among the different crop species revealed a similar imbalance, with a minimum of 700 training samples for raspberry and a maximum of 13,180 total training samples for tomato. This class and species-level data imbalance was also present in the original dataset, as noted in [11]. Despite the data imbalance in the target dataset, the experiments were conducted with a data split that preserved the overall distribution of the dataset.

3.2. Experiment Design

The core experiments of this study focus on the application of XAI techniques to vision-based plant disease classification. Nonetheless, we conducted the following series of experiments targeting the integration of both inherently transparent and post-hoc explainability approaches:

- **Visual Transparency**: In this experiment, we assessed techniques for elucidating the intrinsic decision-making process within inherently transparent models, such as the k-Nearest Neighbors (kNN). The goal was to understand how kNN classification works by examining the voting mechanism where the class of a data point is determined based on the majority class among its k closest neighbors in the feature space.
- **Vision Post-hoc Explainability**: This experiment investigated the application of advanced post-hoc explainability methods, including LIME, Grad-CAM, and occlusion-based attribution techniques, to interpret the predictions made by a vision transformer model, specifically ViT B-32, which was trained on the target dataset for the task of plant disease classification.
- **Model Accuracy and Explainability**: The third experiment evaluated how explainability varies with high and low-accurate vision models. For this purpose, using LIME, we compared the explainability characteristics of the ViT B-32 model against a less accurate configuration of a ResNet-50 model.

4. Results

This section presents an analysis of the results obtained from the conducted experiments.

4.1. Vision Transparency

The kNN algorithm is a simple yet effective classification technique that does not require an explicit model training process. Instead, it relies on the proximity of the input data to the training samples in the feature space. To understand the decision-making process of the kNN model, we examined the voting mechanism that determines the class prediction for a given input.

When presented with a plant disease image, the kNN classifier first identifies the k nearest neighbors in the training set, based on a distance metric (Euclidean distance) in the feature space generated through a feature extraction process using a ResNet-50 model. The predicted class is then assigned based on the majority vote among these k nearest neighbors.

Through a custom visualization (see Figure 1), a slider was implemented that allows a user to change the value of k and see the corresponding change in the voting distribution and the final prediction. For instance, the figure illustrates the 5 nearest neighbors of a given input instance, showcasing their respective class labels and distance metrics, with the instance being predicted as label 0 given that label 0 was the majority among the neighbors.

4.2. Visual Explanations

The second set of experiments focused on applying post-hoc explainability techniques to a vision transformer model, specifically the ViT B-32 architecture, trained on the plant disease classification task. The model decision was preceded by a cross-validation evaluation of two models: ViT B-32 and ResNet-50. Both models employed transfer learning with feature extraction, where the initial layers were frozen to leverage pre-trained features, and only the classifier layers were trained on the plant disease dataset. Given the ViT B-32 model's outperforming accuracy of 97.6% on the validation set compared to the ResNet-50 model, it was selected for the subsequent explainability analysis. Since this experiment aimed to compare multiple XAI techniques rather than evaluating model performance, we present explanations for ViT B-32 (see Figure 2).

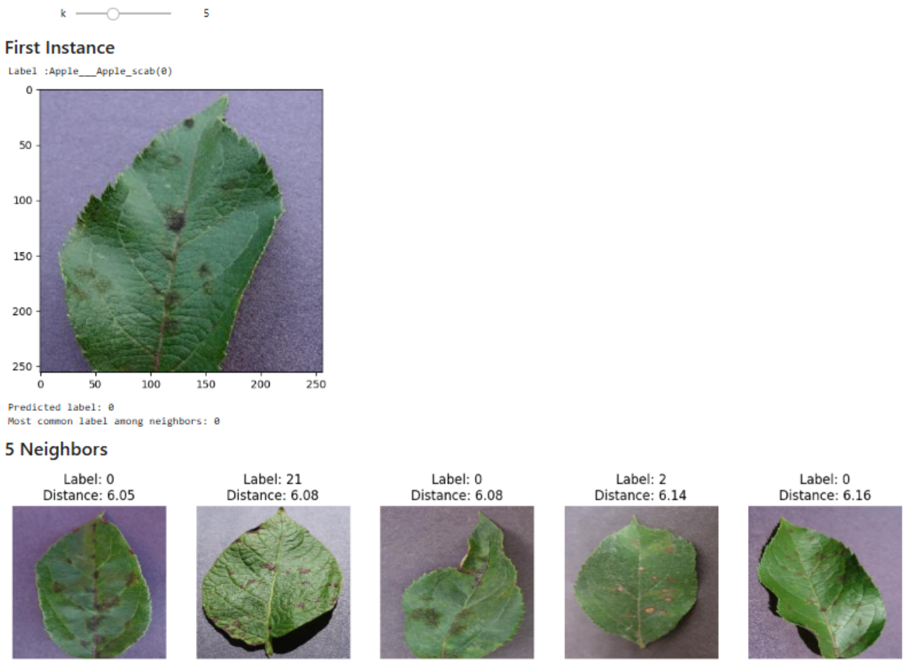

Fig. 1: kNN Slider Visualization.

As stated, the post-hoc explainability methods used in this experiment were LIME, Grad-CAM, and occlusion-based attribution. Targeting the first instance of the test set, we generated and analyzed the explanations provided by these techniques to interpret the model's decision-making process.

The visualization of the top 5 superpixels through LIME (see Figure 2b) highlighted the most salient regions within the input image that contributed the most to the model's prediction. In contrast, Grad-CAM (see Figure 2c) and the occlusion-based attribution technique (see Figure 2d) provided visual heatmaps indicating the areas of the image that the ViT B-32 model deemed crucial for the classification decision. Specifically, in occlusion-based attributions, the darker the pixels, the more important that region is for the prediction. While all visualization methods emphasized the relevance of plant foliage for the classification, there were some differences in the specific regions highlighted by each approach. Notably, the Grad-CAM and occlusion-based attributions suggested a potential discrepancy, as the Grad-CAM heatmap did not consider the center of the plant foliage as crucial for the classification, whereas the occlusion technique deemed this region as one of the most important.

Furthermore, as part of this experiment, we assessed how the visual explanations changed with different class predictions. Specifically, we applied the post-hoc explainability techniques to the same input image, but with different target class labels, focusing on the first and fifth classes among the top 5 predictions (see Figure 3). From the resulting

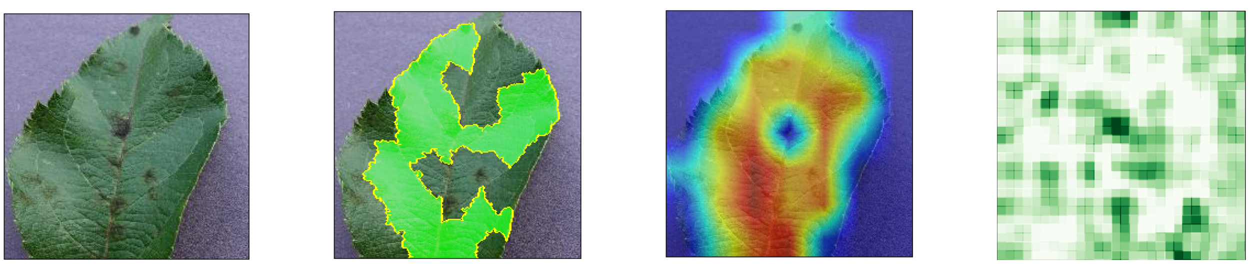

(a) Original Image (b) LIME (c) Grad-CAM (d) Occlusion

Fig. 2: Comparison of XAI techniques (LIME, Grad-CAM, occlusion-based attributions) applied to ViT B-32.

explanations, we observed that model accuracy affects the generated explanations, as seen in the differences between the first and fifth predictions. Given the risk of imperfect explanations [17], future work should validate whether the best-performing models provide the most reliable explanations by collaborating with expert agronomists.

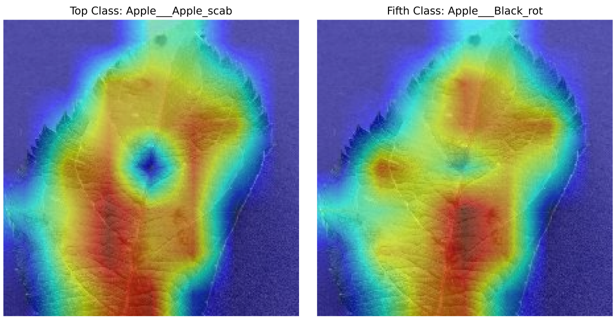

Fig. 3: Grad-CAM visualization of 1st and 5th classes of top 5 predictions.

4.3. Model Accuracy and Explainability

The third experiment compared the explainability characteristics of the high-performing ViT B-32 model and the less accurate ResNet-50 model using the LIME technique. Using the best performing ViT B-32 model configuration (97.6% validation accuracy) and a ResNet-50 model with 92.2% validation accuracy, we generated LIME explanations for the same test sample and top predicted class, then contrasted the differences in the highlighted salient regions (see Figure 4). As it can be observed, the LIME explanations for the more accurate ViT B-32 model provided a more comprehensive and intuitive interpretation of the decision-making process, with the superpixels covering only the plant foliage. However, the ResNet explanation highlighted the negative impact of plant foliage on the classification.

Therefore, this experiment highlights the relationship between model performance and the quality of generated explanations. These findings suggest a potential link between accuracy and explainability, as seen in the comparison between ViT B-32 and ResNet-50. However, this observation is based on a single sample and should not be generalized. The current literature poses the concern that the trade-off between performance and explainability is not a binary decision, but rather a complex issue influenced by multiple factors such as domain characteristics and development resources [5]. Future work should further investigate these interactions through a broader empirical analysis and expert validation.

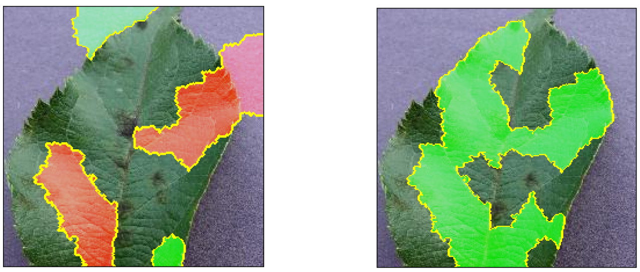

(a) ResNet-50 (b) ViT B-32

Fig. 4: Visualization of LIME implementation for worst and best performing settings.

4.4. Integration with Third Party Applications

The XAI techniques outlined in this section hold significant potential for practical integration into third-party applications or prototypes, enhancing transparency and interpretability in plant disease classification tasks. Farm digital platforms capable of generating such explanations like AgriUXE, which combine multimodal data with XAI techniques, can process plant images, classify diseases using vision-based deep learning models, and generate user-adaptive visual and textual explanations that clarify the reasoning behind disease classifications [20]. Hence, potential applications could connect with AgriUXE to retrieve these explanations, which can be tailored to meet the needs of different stakeholders, including farmers, agronomists, and machine learning experts, by adjusting the explanation technique (e.g., LIME, Grad-CAM), the explanation output modality (e.g., visual overlays or textual descriptions), and the presentation to align with their specific mental models.

This practical integration of XAI explanations can make AI-driven predictions more transparent, supporting stakeholders in understanding the decision-making process and fostering more informed, context-aware decisions in plant disease management.

5. Conclusions and Future Work

This study presents a comprehensive analysis of applying XAI techniques to a vision classification task involving plant disease recognition. Through a series of experiments, we focused on assessing not only visual techniques that could convey the inherent transparency of simple models, but also post-hoc explainability methods that can shed light on the decision-making process of complex and high-performing vision models.

The kNN slider visualization highlights the intuitive nature of this simple classification model, allowing the user to directly observe the voting process and its impact on the final prediction. Furthermore, the post-hoc explainability analysis on the ViT B-32 model, using techniques such as LIME, Grad-CAM, and occlusion-based attribution, revealed intriguing insights into the model's decision-making. While the explanations provided by these methods shared some commonalities in emphasizing the importance of plant foliage and lesions, there were also notable differences in the specific regions highlighted as crucial for the classification.

Moreover, the comparison between ViT B-32 and ResNet-50 provided insights into model explainability. While both models achieved significant classification accuracy in the task at hand, LIME explanations for ViT B-32 appeared more focused on plant lesions. However, this observation is based on limited samples and requires further validation.

As a final consideration, we briefly discussed the potential for integrating XAI explanations into third-party applications or prototypes, emphasizing their practical value in enhancing the interpretability of plant disease classifications.

Future work will focus on extending these experiments by exploring additional XAI techniques, such as Layer-wise Relevance Propagation (LRP) or similarity-based explanations, to further analyze and compare their performance in plant disease classification tasks. In addition, we highlight the potential for developing applications that incorporate explanation components to further study the transparency of such systems and their impact on user experience in real-world agricultural scenarios.

Acknowledgements

This work was supported by **Fundação para a Ciência e a Tecnologia (FCT)**, I.P., through the Ph.D. Studentship grant with project reference PRT/BD/154548/2023 and DOI identifier 10.54499/PRT/BD/154548/2023. This project has indirectly received funding from the European Union's Horizon Europe research and innovation programme under project **AGRARIAN** (grant agreement No 101134128).

References

[1] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F., 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion 58, 82–115. URL: https://www.sciencedirect.com/science/article/pii/S1566253519308103, doi:10.1016/j.inffus.2019.12.012.

[2] Brahimi, M., Boukhalfa, K., Moussaoui, A., 2017. Deep Learning for Tomato Diseases: Classification and Symptoms Visualization. Appl. Artif. Intell. 31, 299–315. URL: https://doi.org/10.1080/08839514.2017.1315516, doi:10.1080/08839514.2017.1315516.

[3] Chandra, R., Collis, S., 2021. Digital agriculture for small-scale producers: challenges and opportunities. Communications of the ACM 64, 75–84. URL: https://dl.acm.org/doi/10.1145/3454008, doi:10.1145/3454008.

[4] Coulibaly, S., Kamsu-Foguem, B., Kamissoko, D., Traore, D., 2022. Deep learning for precision agriculture: A bibliometric analysis. Intelligent Systems with Applications 16, 200102. URL: https://www.sciencedirect.com/science/article/pii/S2667305322000400, doi:10.1016/j.iswa.2022.200102.

[5] Crook, B., Schlüter, M., Speith, T., 2023. Revisiting the Performance-Explainability Trade-Off in Explainable Artificial Intelligence (XAI). URL: http://arxiv.org/abs/2307.14239, doi:10.48550/arXiv.2307.14239. arXiv:2307.14239 [cs].

[6] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N., 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. URL: http://arxiv.org/abs/2010.11929, doi:10.48550/arXiv.2010.11929. arXiv:2010.11929 [cs].

[7] Dwivedi, R., Dave, D., Naik, H., Singhal, S., Omer, R., Patel, P., Qian, B., Wen, Z., Shah, T., Morgan, G., Ranjan, R., 2023. Explainable AI (XAI): Core Ideas, Techniques, and Solutions. ACM Computing Surveys 55, 194:1–194:33. URL: https://dl.acm.org/doi/10.1145/3561048, doi:10.1145/3561048.

[8] FAO High Level Expert Forum-How to Feed the World in 2050, 2009. Global agriculture towards 2050. Technical Report.

[9] Fountas, S., Espejo-García, B., Kasimati, A., Mylonas, N., Darra, N., 2020. The Future of Digital Agriculture: Technologies and Opportunities. IT Professional 22, 24–28. URL: https://ieeexplore.ieee.org/document/8994135, doi:10.1109/MITP.2019.2963412. conference Name: IT Professional.

[10] Harvey, C.A., Rakotobe, Z.L., Rao, N.S., Dave, R., Razafimahatratra, H., Rabarijohn, R.H., Rajaofara, H., MacKinnon, J.L., 2014. Extreme vulnerability of smallholder farmers to agricultural risks and climate change in Madagascar. Philosophical Transactions of the Royal Society B: Biological Sciences 369, 20130089. URL: https://royalsocietypublishing.org/doi/10.1098/rstb.2013.0089, doi:10.1098/rstb.2013.0089. publisher: Royal Society.

[11] Hughes, D.P., Salathe, M., 2016. An open access repository of images on plant health to enable the development of mobile disease diagnostics. URL: http://arxiv.org/abs/1511.08060, doi:10.48550/arXiv.1511.08060. arXiv:1511.08060 [cs].

[12] Kamilaris, A., Prenafeta-Boldú, F.X., 2018. Deep learning in agriculture: A survey. Computers and Electronics in Agriculture 147, 70–90. URL: https://www.sciencedirect.com/science/article/pii/S0168169917308803, doi:10.1016/j.compag.2018.02.016.

[13] Kernecker, M., Knierim, A., Wurbs, A., Kraus, T., Borges, F., 2020. Experience versus expectation: farmers' perceptions of smart farming technologies for cropping systems across Europe. Precision Agriculture 21, 34–50. URL: https://doi.org/10.1007/s11119-019-09651-z, doi:10.1007/s11119-019-09651-z.

[14] Lowder, S.K., Skoet, J., Raney, T., 2016. The Number, Size, and Distribution of Farms, Smallholder Farms, and Family Farms Worldwide. World Development 87, 16–29. URL: https://www.sciencedirect.com/science/article/pii/S0305750X15002703, doi:10.1016/j.worlddev.2015.10.041.

[15] Martinelli, F., Scalenghe, R., Davino, S., Panno, S., Scuderi, G., Ruisi, P., Villa, P., Stroppiana, D., Boschetti, M., Goulart, L.R., Davis, C.E., Dandekar, A.M., 2014. Advanced methods of plant disease detection. A review. Agronomy for Sustainable Development 35, 1–25. URL: https://link.springer.com/article/10.1007/s13593-014-0246-1, doi:10.1007/s13593-014-0246-1. number: 1 Publisher: Springer.

[16] Mohanty, S.P., Hughes, D.P., Salathé, M., 2016. Using Deep Learning for Image-Based Plant Disease Detection. Frontiers in Plant Science 7. URL: https://www.frontiersin.org/articles/10.3389/fpls.2016.01419.

[17] Morrison, K., Spitzer, P., Turri, V., Feng, M., Kühl, N., Perer, A., 2024. The Impact of Imperfect XAI on Human-AI Decision-Making. Proc. ACM Hum.-Comput. Interact. 8, 183:1–183:39. URL: https://dl.acm.org/doi/10.1145/3641022, doi:10.1145/3641022.

[18] Mrisho, L.M., Mbilinyi, N.A., Ndalahwa, M., Ramcharan, A.M., Kehs, A.K., McCloskey, P.C., Murithi, H., Hughes, D.P., Legg, J.P., 2020. Accuracy of a Smartphone-Based Object Detection Model, PlantVillage Nuru, in Identifying the Foliar Symptoms of the Viral Diseases of Cassava–CMD and CBSD URL: https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2020.590889/full, doi:10.3389/fpls.2020.590889.

[19] Mukti, I.Z., Biswas, D., 2019. Transfer Learning Based Plant Diseases Detection Using ResNet50, in: 2019 4th International Conference on Electrical Information and Communication Technology (EICT), pp. 1–6. URL: https://ieeexplore.ieee.org/document/9068805, doi:10.1109/EICT48899.2019.9068805.

[20] Porfírio, R.P., Madeira, R.N., Santos, P.A., 2024. AgriUXE: Integrating Explainable AI and Multimodal Data for Smart Agriculture, in: 2024 International Symposium on Sensing and Instrumentation in 5G and IoT Era (ISSI), pp. 1–6. URL: https://ieeexplore.ieee.org/abstract/document/10720487, doi:10.1109/ISSI63632.2024.10720487.

[21] Reedha, R., Dericquebourg, E., Canals, R., Hafiane, A., 2022. Transformer Neural Network for Weed and Crop Classification of High Resolution UAV Images. Remote Sensing 14, 592. URL: https://www.mdpi.com/2072-4292/14/3/592, doi:10.3390/rs14030592. number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

[22] Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, New York, NY, USA. pp. 1135–1144. URL: https://dl.acm.org/doi/10.1145/2939672.2939778, doi:10.1145/2939672.2939778.

[23] Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2020. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. International Journal of Computer Vision 128, 336–359. URL: http://arxiv.org/abs/1610.02391, doi:10.1007/s11263-019-01228-7. arXiv:1610.02391 [cs].

[24] Walter, A., Finger, R., Huber, R., Buchmann, N., 2017. Smart farming is key to developing sustainable agriculture. Proceedings of the National Academy of Sciences 114, 6148–6150. URL: https://www.pnas.org/doi/10.1073/pnas.1707462114, doi:10.1073/pnas.1707462114. publisher: Proceedings of the National Academy of Sciences.

[25] Wei, K., Chen, B., Zhang, J., Fan, S., Wu, K., Liu, G., Chen, D., 2022. Explainable Deep Learning Study for Leaf Disease Classification. Agronomy 12, 1035. URL: https://www.mdpi.com/2073-4395/12/5/1035, doi:10.3390/agronomy12051035. number: 5 Publisher: Multidisciplinary Digital Publishing Institute.