Enhancing Digital Agriculture with XAI: Case Studies on Tabular Data and Future Directions

Rui Pedro Porfirio
NOVA LINCS, NOVA School of Science and Technology, NOVA University of Lisbon
Portugal
RESILIENCE, Setúbal School of Technology, Polytechnic Institute of Setúbal
Portugal
r.porfirio@campus.fct.unl.pt

Pedro Albuquerque Santos
Lisbon School of Engineering, Polytechnic Institute of Lisbon
Portugal
NOVA LINCS, NOVA School of Science and Technology, NOVA University of Lisbon
Portugal
RESILIENCE, Setúbal School of Technology, Polytechnic Institute of Setúbal
Portugal
pedrosantos@deetc.isel.ipl.pt

Rui Neves Madeira
RESILIENCE, Setúbal School of Technology, Polytechnic Institute of Setúbal
Portugal
NOVA LINCS, NOVA School of Science and Technology, NOVA University of Lisbon
Portugal
rui.madeira@estsetubal.ips.pt

Abstract

Given the pivotal role of agriculture in ensuring food security, fostering economic stability, and addressing environmental sustainability, the sector has increasingly embraced smart farming solutions to respond to recent climate and societal challenges, such as rising water and food demands. These solutions provide actionable insights crucial for decision-making, enabling farm stakeholders to optimize resources, improve yields, and mitigate risks. However, the complexity of the predictive models often associated with this type of solutions results in a lack of transparency, hindering trust and adoption. To respond to such challenges, this paper explores the application of explainable AI (XAI) techniques to agriculture tabular data. Specifically, we focus on two case studies: wheat yield prediction and grapes produced for wine purposes yield prediction. Through these case studies, we propose initial contributions on how XAI techniques can be applied in the context of agriculture and how generated explanations can be adapted to the users' level of expertise. Finally, as part of ongoing and future research directions, we introduce AgriUXE (Agricultural eXperience Enhanced through eXplainability), a novel user-centered digital platform designed to augment the explainability of multimodal data and machine learning model predictions for sustainable smart farming solutions. By providing transparent, data-driven decisions and generating user-adaptive explanations, AgriUXE aims to support the optimization of the user experience within these solutions.

CCS Concepts

• **Applied computing → Agriculture**; • **Computing methodologies → Machine learning**; • **Human-centered computing → User centered design**; **Human computer interaction (HCI)**.

ICMI Companion '24, November 04–08, 2024, San Jose, Costa Rica

ACM ISBN 979-8-4007-0463-5/24/11
https://doi.org/10.1145/3686215.3689201

Keywords

Digital Agriculture, Explainable AI, Machine Learning, Multimodal Systems, Human-Computer Interaction

ACM Reference Format:
Rui Pedro Porfirio, Pedro Albuquerque Santos, and Rui Neves Madeira. 2024. Enhancing Digital Agriculture with XAI: Case Studies on Tabular Data and Future Directions. In *INTERNATIONAL CONFERENCE ON MULTIMODAL INTERACTION (ICMI Companion '24), November 04–08, 2024, San Jose, Costa Rica.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3686215.3689201

1 Introduction

Agriculture has consistently been a cornerstone industry, sustaining human evolution throughout history. In recent years, however, the industry has faced unprecedented challenges due to the impacts of climate change and the urgent need for a more efficient agrifood supply chain. Accounting for 4% of the global GDP and up to 25% of the GDP of low-income countries [1], agriculture not only finds itself among the most affected industries, with predictions of up to 30% loss on crop yields by 2050 due to climate change [4], but also attracts increasing attention as a sector where immediate action and investment can address the stated challenges, providing practical and efficient solutions [8]. Particularly, recent advancements in digital agriculture, that is the "implementation of emerging technologies and innovative services on agriculture" [5], have recently gained momentum in research and innovation schemes [8].

A prevalent number of digital solutions in agriculture involve employing machine learning models for agricultural predictive tasks. Nonetheless, the complexity of these models, often perceived as complete opaque "black boxes", especially in deep learning contributions, combined with a widespread technological illiteracy among agriculture stakeholders, poses intriguing challenges in the perceived value and complexity of smart farming solutions.

Hence, this paper presents an analysis on how explainable AI (XAI) techniques can be applied in the context of agriculture. Specifically, we delved into exploring two case studies focused on the task of yield prediction for tabular data: wheat yield and grapes for wine production yield.

Within the frame of the case studies, we propose a set of experiments on explanations for both transparent models (k-Nearest Neighbors (*kNN*) and *Decision Tree*) and opaque models (*XGBoost* and Multilayer Perceptron (*MLP*)). Additionally, we further explore the role of feature-level XAI techniques, such as LIME and SHAP, alongside feature importance techniques in crop yield prediction. By exploring these case studies, we focus not only on studying how the application of state-of-the-art XAI techniques enhances model transparency but also on how the explanations can be refined, through large language models (*GPT-3.5* and *Llama-2*), and adapted to the expertise level of different users, ultimately supporting informed decision-making in sustainable agricultural practices.

Lastly, we propose a novel user-centered digital platform *AgriUXE* (Agricultural eXperience Enhanced through eXplainability) that incorporates multimodal data and XAI techniques to offer explainable data-driven insights, therefore supporting informed decision-making in agriculture. By enhancing not only data explainability but also the interpretability of machine learning predictions, the platform intends to further study how added explainability impacts agriculture stakeholders' user experience, and ultimately, their perceived value of sustainable smart farming solutions.

In the remainder of this paper, Section 2 presents the related work and background on XAI in agriculture. Section 3 details the predictive task, case studies and and associated experimental results. Section 4 introduces the *AgriUXE* platform, outlining its features and potential impact. Finally, Section 5 concludes the paper with the research outcomes and future directions.

2 Related Work

The growing research on digital agriculture solutions brought forth the application of a comprehensive set of emerging technologies in the sector. Nevertheless, according to a survey by *McKinsey & Company* on agriculture technological adoption, as of 2022, 38% of European farmers were not using or planning to use at least one technology in their farm operations within the following 2 years [16]. Moreover, this adoption gap is more pronounced on a global scale, with 61% of farmers not being proactive in adopting technology.

To bridge this adoption gap, the motivation for farm stakeholders to embrace digitalization will increase when the perceived value of a technology becomes evident [9]. The high complexity of the solutions available coupled with the lack of data interpretability is directly linked to farmers' lack of perception of the real added value of using smart farm technologies [14]. Therefore, the perceived value of smart farm technologies can be significantly enhanced by making these technologies more transparent and easier to understand.

Since a significant proportion of digital solutions in agriculture employ machine learning models for predictive tasks, a promising approach to addressing these challenges is through the implementation of explainable AI (XAI) techniques.

While machine and deep learning models offer high prediction accuracy, there is a systematic lack of transparency attached to these models, especially as model complexity increases. Not understanding why a model reached a certain prediction X instead of a prediction Y poses missing potential new knowledge on the learning task, the correlation between feature data, and the reason why a model might fail [18]. Explainable AI (XAI) is the field responsible for tackling these issues by providing clear and interpretable explanations of models' predictions. XAI techniques achieve increased interpretability through both transparent models and post-hoc explainability techniques [2, 6, 20].

Transparent models, such as logistic regression, decision trees, and k-nearest neighbors (*kNN*), inherently provide a certain level of interpretability due to the straightforward manner in which they make predictions. These models allow users to understand the decision-making process without requiring additional explanation techniques, addressing the needs for simulability, decomposability, and algorithmic transparency.

More complex or opaque models, however, usually require the application of post-hoc explainability techniques. These techniques may include text explanations, visual explanations, explanations by simplification, local explanations, and feature relevance explanations. For example, model-agnostic techniques like *LIME* [22] and *SHAP* [15] support the explainability of the inner workings of black-box models by generating local surrogate linear models around predictions and computing additive feature importance scores, respectively. Model-specific techniques can also be employed to tailor explanations to particular model types, such as visualization methods for understanding support vector machines [25] or similarity-based explanations like gradient similarity [10], usually applied for neural networks. Regarding the latter, the output of the model is explained by identifying and presenting instances from the training data that are similar to the given input instance. This approach is grounded in the rationale that "This X is a Y because a similar instance X' is a Y".

Furthermore, given there is a lack of consensus on the definition of explainability within the computer science academia, in the context of this paper, the definition of explainability is combined with interpretability through the terminology offered by Barredo et al.: "*given an audience, an explainable Artificial Intelligence is one that produces details or reasons to make its functioning clear or easy to understand*" [2]. This implies that the model not only offers an explanation but also ensures that the explanation is comprehensible to the audience, which is a key aspect of interpretability.

Although the research community has witnessed a rise in XAI contributions in recent years, the application of explainable models within the field of agriculture is yet in the early stages of providing a significant set of contributions. Nevertheless, recent research has begun exploring how explainable techniques like *SmoothGrad*, *LIME*, and *GradCAM* can enhance the interpretability of plant disease detection *CNN* architectures [24]. Additionally, researchers have been investigating the impact of various regions in grassland imagery on model predictions [11], and developing ontology-based XAI frameworks to explain machine learning predictions in digital agriculture solutions [21].

Considering the target vision of using XAI techniques to address farm adoption barriers such as highly complex digital products, an additional concern with designing user-centered solutions should be taken into consideration. Design principles such as active self-explanation, contextual relevance, and iterative user feedback are essential to ensure that the systems are not only technically sound but also practically useful and trustworthy for farm stakeholders [19].

Considering both the technical aspects and the ethical and societal implications of XAI contributions, this emerging field has the potential to increase the transparency of farm predictive models and enhance the perceived value of smart farming applications, particularly in contributions that align with the United Nations' *Sustainable Development Goals* [2].

3 Methods and Results

This section provides an overview of the methods applied in the analysis of the case studies. Furthermore, we introduce the experimental methodology and results associated to each case study experiment.

3.1 Predictive Task

In order to assess the potential of XAI techniques in smart farming solutions, we focused on a significant predictive task: crop yield prediction. Crop yield prediction is the process of estimating the amount of crop that will be harvested from a given area. This task is critical in agricultural planning and decision-making as it helps farmers optimize resource allocation and improve overall farm management.

3.2 Datasets

For our research, we investigated the impact of various XAI techniques on agriculture tabular data.

The data collection process was constrained to publicly available statistical information related to crop production [7, 12, 13, 17, 23], weather-related features [3], and pesticide usage [7, 23]. While factors such as precipitation, temperature, and average pesticide usage are commonly recognized as influential factors in crop yield, the decision to limit the dataset's features to these factors was primarily influenced by significant constraints in data availability.

Given these constraints, we created the four following custom datasets for our analysis:

- *United States Wheat Dataset*: This tabular dataset comprises historic annual average wheat yield data at the US state-level, spanning from 1992 to 2019. In addition to wheat yield, the dataset includes annual average measurements of state-wide precipitation, temperature, and wheat pesticide usage.
- *European Union Wheat Dataset*: This custom tabular dataset captures the annual average wheat crop yield across European Union countries from 1990 to 2021. The dataset features annual average country-wide measurements of precipitation, temperature, and pesticide usage.
- *World Wheat Dataset*: This dataset provides country-level annual average wheat yield information from 120 countries spanning from 1990 to 2021. Additionally, the dataset encompasses weather-related features such as precipitation, temperature, and average pesticide usage.
- *European Grapes for Wine Dataset*: This tabular dataset contains historical data spanning from 1986 to 2022, focusing on the average annual yield of grapes intended for wine production. The dataset includes data from four major European wine-producing countries: Portugal, Spain, Italy, and France. Apart from yield information, the dataset contains average annual measurements of country-wide precipitation, relative humidity, and temperature.

3.3 Case Study 01: Wheat Yield Prediction

In this case study, we primarily focused on studying how XAI feature-level techniques (i.e., LIME, SHAP, and feature importance techniques) provide additional information related to not only the collected data but also machine learning models' crop yield predictions.

To address this evaluation, we conducted a series of experiments targeting both XAI-generic and farm-related research questions:

- How to highlight transparency in transparent models?
- How explanations differ based on model accuracy?
- How do feature-level XAI techniques enhance the interpretability of crop yield prediction?
- How can feature importance techniques assist in the process of understanding input agricultural data?
- What impact does dataset size have on the generated explanations?
- How do explanations contribute to the interpretability of metric results?
- Can large language models support the definition of user-adaptive explanations?

The following subsections focus on outlining the conducted experiments and associated results.

3.3.1 Highlighting transparency in transparent models. In this experiment, the target goal was to find mechanisms to emphasize transparency in inherently transparent models such as the *kNN* and Decision Tree models. Specifically, we applied custom visualization tools such as a custom *kNN* slider and a decision tree viewer to demonstrate the inherent transparency of these models. The *kNN* slider (see Figure 1) allowed us to interactively adjust the number of neighbors (k) and observe how predictions change, providing a clear understanding of the model's decision boundaries. Similarly, the decision tree viewer illustrated the decision paths, making it easy to follow the logic used to reach a prediction. These tools effectively highlight how transparent models allow users to grasp the decision-making process without requiring complex additional explanations.

3.3.2 Variations in Explanations Based on Model Accuracy. This experiment focused on analyzing how explanations differ based on model accuracy. Using *LIME* as the explanation generation technique, we assessed the *kNN* model's performance using both the worst (MAE = 0.61) and best (MAE = 0.45) performing k values. Subsequently, we employed a sentence transformer to measure sentence similarity between the explanations generated for each evaluation setup.

The results indicated that while the explanations diverge, they were consistent in providing valid alternative insights. For example, the sentences "If the temperature is less than or equal to 12.55 ºC then the temperature tends to have a positive impact on the crop yield." and "If the temperature is greater than 12.55 ºC and less than or equal to 14.05 ºC then the temperature tends to have a positive impact on the crop yield" had a similarity score of 0.95. Although

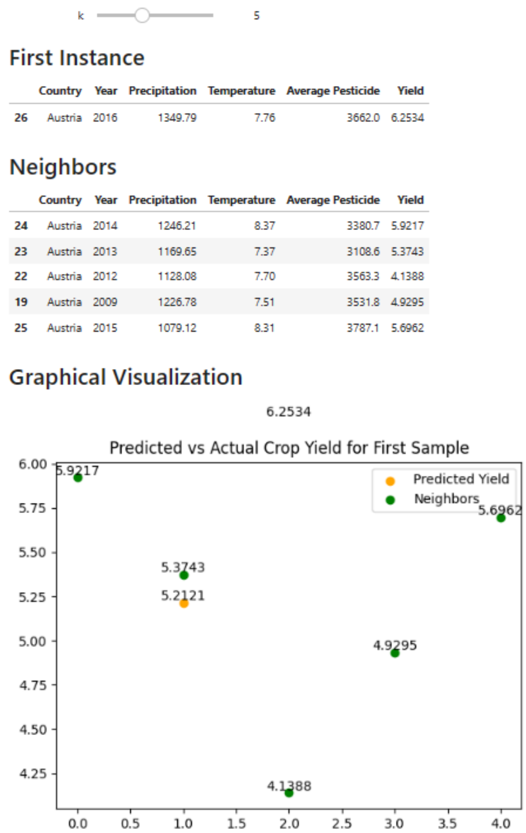

Figure 1: *kNN* Slider Visualization Tool.

they were not fully identical (i.e. sentence similarity score = 1), the explanations provide alternative non-conflicting explanations. However, further testing with models having a larger performance gap is needed to strengthen this hypothesis.

3.3.3 Enhancing Crop Yield Prediction Interpretability with Feature-Level XAI Techniques. In this experiment, we explored how feature-level XAI techniques, *LIME* and *SHAP*, contributed to the interpretability of crop yield predictions. Although *LIME* is designed for local explanations (i.e., generating explanations for individual instances), an approach was devised to also extract global explanations (i.e., explanations for the entire test set). By computing *LIME* across all data points of the test set, we aggregated local explanations into a coherent global interpretation. This method involved filtering the *LIME* explanations by key dataset features: temperature, precipitation, and pesticide usage. For each feature, we further filtered the explanations to identify the instances with the highest and lowest weights, capturing the most significant positive and negative contributions of each feature to crop yield predictions. By selecting the explanations with extreme weights, we mitigated inconsistencies and highlighted the most impactful feature values. The method for generating global *LIME* explanations proved valid given *LIME* produced a limited set of explanations across all data points in the test set, rather than generating different explanations for each individual instance.

Additionally, since *LIME* creates a local surrogate, this technique was used to assess explanations by simplification. In this context, we compared the original black-box model (i.e., in this experiment we used *XGBoost*) predictions with those of the simpler local surrogate model to determine how well the surrogate replicated the black-box model's predictions.

While *LIME* traditionally provides insights into local data points, *SHAP* offers a unified framework to understand both local and global feature importance. We used *SHAP* values to complement *LIME* explanations, combining feature importance insights with *LIME*'s highest and lowest weighted explanations.

To enhance the understandability of the explanations, we created an NLP manual template. This template transformed the filtered *LIME* and *SHAP* explanations into natural language statements, making it easier for users to understand the implications of different feature values on crop yield predictions.

Regarding experimental results, *LIME* provided both local and global explanations, illustrating the range and weight each feature had in increasing or decreasing the predictions. *SHAP* provided not only visualization mechanisms, such as waterfall plots, to interpret local predictions but also visual explanations, like the beeswap plot, that support a clear understanding of the global impact of each feature on the test set (see Figure 2).

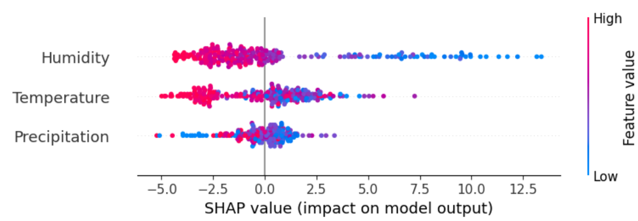

Figure 2: SHAP Beeswarm Visualization plot from Case Study 02.

By combining these insights with feature importance scores from methods such as permutation importance and gradient boosting, we developed a more robust perspective on the contributions of various input features, enhancing our understanding of what drives model predictions.

Moreover, the combined explanations of *LIME* and *SHAP*, alongside an NLP template, translated complex numerical explanations into user-friendly statements (see Figure 3). This step is crucial in making the insights accessible to non-technical stakeholders.

For the evaluation of explanations by simplification, the discrepancy between the two setups (i.e., black-box performance against local surrogate performance) resulted in a MAE of 0.45 tonnes per hectare, indicating a moderate alignment between the black-box model and the local surrogate.

3.3.4 Feature Importance Techniques Impact on Understanding Input Agricultural Data. This experiment focused on assessing the impact of feature importance techniques on understanding input agricultural data. We combined permutation importance, *SHAP*, random forest, and gradient boosting feature importance calculations to determine the degree of impact of each feature on the

The **positive** yield impacting factors are:

- If the temperature is less than or equal to 12.55 then the temperature tends to have a positive impact on the crop yield.
- If the precipitation is less than or equal to 578.30 then the precipitation tends to have a positive impact on the crop yield.
- If the humidity is less than or equal to 68.43 then the humidity tends to have a positive impact on the crop yield.

The **negative** yield impacting factors are:

- If the temperature is greater than 16.14 then the temperature tends to have a negative impact on the crop yield.
- If the precipitation is greater than 933.53 then the precipitation tends to have a negative impact on the crop yield.
- If the humidity is greater than 75.28 then the humidity tends to have a negative impact on the crop yield.

The **most important** features impacting yield prediction are:

- Humidity
- Temperature
- Precipitation

Figure 3: Textual Explanations extracted from LIME and SHAP.

predictions, providing insights into which features had the most significant influence.

Through these analyses, we uncovered the global contribution of each feature to the test set predictions, offering valuable insights for optimizing agricultural practices based on model predictions (e.g., identifying and discarding unimportant features).

3.3.5 Impact of Dataset Size on Generated Explanations. In this evaluation, we examined how dataset size influences the generated explanations. To do so, we assessed the *XGBoost* model using two datasets with different size granularity: the *World Wheat Dataset* (3746 samples) and the *European Union Wheat Dataset* (734 samples).

The findings suggested that dataset size can influence explanation consistency. For example, the explanations inconsistencies in the sentences "If the temperature is greater than 23.39 ºC then the temperature tends to have a negative impact on the crop yield." and "If the temperature is greater than 12.79 ºC then the temperature tends to have a negative impact on the crop yield." need to be taken into consideration as they pose an interpretability conflict. Nevertheless, further testing is required to validate these observations and assess their generalization across different model architectures and datasets.

3.3.6 Contributions of Explanations to Metrics' Interpretability. In this experiment, we explored how explanations contributed to the interpretability of metrics' results. Using the *GPT 3.5* model and a set of regression metrics connected to the predictive task, we asked the large language model to explain the metrics of *XGBoost* and Decision Tree models, providing insights into the models' performance.

Our analysis revealed that explanations generated by large language models may significantly increase metrics' interpretability, offering a deeper understanding of model performance and aiding in decision-making processes.

3.3.7 Applying Large Language Models for User-Adaptive Explanations. To retrieve insights on the relevance of large language models in XAI, we compared how two different models, *GPT-3.5* and *Llama-2-7b*, adapted *LIME* and *SHAP* natural language explanations for each of the following user types: farmer, agronomist, and machine learning expert. In order to support a quantitative measurement, we applied a sentence transformer to measure sentence similarity between the explanations generated for each large language model.

Our findings demonstrated that large language models offer promising potential in delivering user-adaptive explanations, with both *GPT-3.5* and *Llama-2* producing similar explanations when answering to the same prompt (i.e., the models were prompted to clarify *LIME* and *SHAP* explanations to users with different levels of expertise). Further investigation is needed to optimize prompt engineering techniques and explore alternative large language models. Additionally, research into using conversational agents, text-to-speech, and speech-to-text techniques in smart farming holds potential.

3.4 Case Study 02: Grapes for Wine Yield Prediction

In this viticulture case study, our primary aim was to evaluate how model-specific techniques, particularly similarity-based explanations, contribute to the crop yield predictive task. Additionally, we addressed supplementary research questions and revisited inquiries to validate findings from the wheat yield case study. The following experiments were conducted in this case study:

- How do similarity-based explanations provide additional insights on local yield predictions?
- How do feature-level XAI techniques enhance the interpretability of crop yield prediction?
- How explanations differ based on model accuracy?

The following subsections detail the set of experiments associated with this case study.

3.4.1 Assessing Similarity-Based Explanations for Local Yield Predictions. In this experiment, we delved into how similarity-based explanations provide additional insights into local yield predictions within the viticulture context. By applying a gradient similarity technique, we explored the potential of these explanations in supporting the interpretability of crop yield predictions. Through this technique, we focused on explaining the first instance of the test set through exposing the five most similar and least similar training set examples. While the initial results from the first test instance showcased the potential of this technique in offering similarity-based explanations, further testing is essential, specifically to assess how this type of explanations correlates with model accuracy.

3.4.2 Enhancing Crop Yield Prediction Interpretability with Feature-Level XAI Techniques. In this experiment, we applied the same goal as the analogous experiment in the wheat yield case study. This revisiting was conducted to validate the consistency and applicability of our findings across different crops. Therefore, in this scenario, we applied feature-level XAI techniques, specifically *LIME* and *SHAP*, to enhance the interpretability of viticulture yield predictions.

Concerning experimental results, when evaluating different crops, significant insights were captured in both wheat and grape yield predictions, underscoring the potential power of these model-agnostic techniques. These consistent findings across diverse datasets highlight the robustness and utility of *LIME* and *SHAP* in enhancing the interpretability of crop yield predictions.

3.4.3 Variations in Explanations Based on Model Accuracy. This experiment focused on analyzing how variations in model accuracy influence the explanations generated. By employing the *LIME* explanation generation technique, we evaluated the performance of an *MLP* model (best performing model during cross-validation with an average MAE of 3.2 tonnes per hectare) and a Decision Tree model (worst performing model during cross-validation with an average MAE of 4.2 tonnes per hectare). Lastly, similarly to the analogous experiment in the wheat yield case study, we executed a sentence transformer to perform sentence similarity between the explanations generated for both evaluation setups.

Similar to the experimental results in the wheat yield case study, the findings indicated that although the explanations varied, they consistently provided valid alternative insights.

4 AgriUXE

Concerning ongoing work and future research directions, we propose AgriUXE (Agricultural eXperience Enhanced through eXplainability), a novel user-centered digital platform designed to augment the explainability of multimodal data and machine learning model predictions in the context of sustainable smart farming solutions.

AgriUXE (see Figure 4) integrates efficient data collection from various multimodal sources, advanced data analysis techniques, and multimodal explanations (i.e, combined through a process of late fusion) on both the captured data and the machine learning models' predictions to provide actionable insights for farmers. It enables seamless integration of agricultural insights into third-party smart farming solutions through an accessible API.

By guiding agriculture stakeholders toward transparent data-driven decisions, AgriUXE was further designed to study how the combination of explainability and user-centered design can influence users' expectations and perceptions of smart farming technologies' added value.

Aligned with the principle of involving all relevant stakeholders, including farmers, agronomy experts, culture technicians, citizens, and civil society, AgriUXE adopts a collaborative approach throughout the development process. This ensures that the platform addresses the diverse needs and perspectives of all stakeholders,

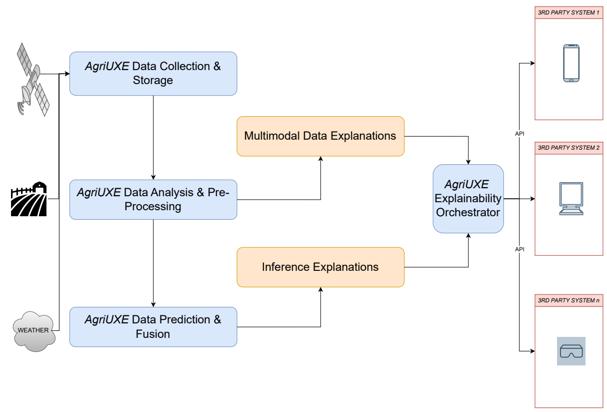

Figure 4: Simplified *AgriUXE* architecture.

fostering inclusive innovation in agricultural technology solutions. This stakeholder engagement journey has already started through workshop sessions aimed at comprehensively understanding their needs, expectations, and perceptions regarding the added value of explainable solutions in agriculture.

5 Conclusions and Future Work

XAI techniques offer valuable support in enhancing the interpretability of both opaque and transparent models in smart farming applications. Through experiments focusing on tabular data, we found that techniques such as *LIME*, *SHAP*, and state-of-the-art feature importance methods provide insightful explanations, contributing to a deeper understanding of model predictions. Furthermore, the potential of large language models in generating and adapting explanations shows promise for future research.

However, our research was limited as certain techniques, particularly visual explanations like PDP, and ACE plots, were not considered in the conducted experiments. Moving forward, our findings emphasize the necessity for continued research, particularly with larger datasets, to validate and refine our hypotheses. Moreover, while XAI techniques exhibit considerable promise in enhancing transparency within tabular data, their broader impact across diverse modalities, such as vision data, demands further investigation.

In addition to this analysis of XAI techniques, as future research directions lead to the development of the AgriUXE platform, it is imperative to adopt a holistic approach that integrates feedback from all relevant stakeholders. Through ongoing workshops and collaborative efforts, we are actively engaging with farmers, agronomy experts, and other stakeholders to understand their needs and expectations, ensuring that AgriUXE aligns with their expectations and enhances their experiences with sustainable digital agricultural solutions.

Acknowledgments

This work is funded by Fundação para a Ciência e a Tecnologia (FCT) through a Ph.D. Studentship grant (PRT/BD/154548/2023).

References

[1] World Bank. 2021. *Agriculture, forestry, and fishing, value added (% of GDP).* Technical Report. https://datacatalog.worldbank.org/indicator/e4a26c2f-c0ce-eb11-bacc-000d3a3b9510/Agriculture--forestry--and-fishing--value-added----of-GDP-

[2] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58 (June 2020), 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

[3] CGKP. 2024. World Bank Climate Change Knowledge Portal.

[4] Ranveer Chandra and Stewart Collis. 2021. Digital agriculture for small-scale producers: challenges and opportunities. *Commun. ACM* 64, 12 (Nov. 2021), 75–84. https://doi.org/10.1145/3454008

[5] Franco da Silveira, Fernando Henrique Lermen, and Fernando Gonçalves Amaral. 2021. An overview of agriculture 4.0 development: Systematic review of descriptions, technologies, barriers, advantages, and disadvantages. *Computers and Electronics in Agriculture* 189 (Oct. 2021), 106405. https://doi.org/10.1016/j.compag.2021.106405

[6] Rudresh Dwivedi, Devam Dave, Het Naik, Smiti Singhal, Rana Omer, Pankesh Patel, Bin Qian, Zhenyu Wen, Tejal Shah, Graham Morgan, and Rajiv Ranjan. 2023. Explainable AI (XAI): Core Ideas, Techniques, and Solutions. *Comput. Surveys* 55, 9 (Jan. 2023), 194:1–194:33. https://doi.org/10.1145/3561048

[7] FAO. 2024. Food and Agriculture Organization of the United Nations . Statistics.

[8] Spyros Fountas, Borja Espejo-García, Aikaterini Kasimati, Nikolaos Mylonas, and Nicoleta Darra. 2020. The Future of Digital Agriculture: Technologies and Opportunities. *IT Professional* 22, 1 (Jan. 2020), 24–28. https://doi.org/10.1109/MITP.2019.2963412 Conference Name: IT Professional.

[9] Andreas Gabriel and Markus Gandorfer. 2023. Adoption of digital technologies in agriculture—an inventory in a european small-scale farming region. *Precision Agriculture* 24, 1 (Feb. 2023), 68–91. https://doi.org/10.1007/s11119-022-09931-1

[10] Kazuaki Hanawa, Sho Yokoi, Satoshi Hara, and Kentaro Inui. 2021. Evaluation of Similarity-based Explanations. https://doi.org/10.48550/arXiv.2006.04528 arXiv:2006.04528 [cs, stat].

[11] Anna Hedström, Leander Weber, Dilyara Bareeva, Daniel Krakowczyk, Franz Motzkus, Wojciech Samek, Sebastian Lapuschkin, and Marina M.-C. Höhne. 2023. Quantus: An Explainable AI Toolkit for Responsible Evaluation of Neural Network Explanations and Beyond. https://doi.org/10.48550/arXiv.2202.06861 arXiv:2202.06861 [cs].

[12] INE. 2024. Instituto Nacional de Estadística - España.

[13] INE. 2024. Instituto Nacional de Estatística - Portugal.

[14] Maria Kernecker, Andrea Knierim, Angelika Wurbs, Teresa Kraus, and Friederike Borges. 2020. Experience versus expectation: farmers' perceptions of smart farming technologies for cropping systems across Europe. *Precision Agriculture* 21, 1 (Feb. 2020), 34–50. https://doi.org/10.1007/s11119-019-09651-z

[15] Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 4768–4777.

[16] McKinsey & Company. 2022. Agtech: Breaking down the farmer adoption dilemma | McKinsey. https://www.mckinsey.com/industries/agriculture/our-insights/agtech-breaking-down-the-farmer-adoption-dilemma

[17] Ministeral Statistical Service for Agriculture. 2024. Statistique agricole annuelle (SAA) - France.

[18] Christoph Molnar. 2022. *Interpretable Machine Learning.* https://christophm.github.io/interpretable-ml-book/

[19] Shane T. Mueller, Elizabeth S. Veinott, Robert R. Hoffman, Gary Klein, Lamia Alam, Tauseef Mamun, and William J. Clancey. 2021. Principles of Explanation in Human-AI Systems. https://doi.org/10.48550/arXiv.2102.04972 arXiv:2102.04972 [cs].

[20] W. James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. 2019. Interpretable machine learning: definitions, methods, and applications. *Proceedings of the National Academy of Sciences* 116, 44 (Oct. 2019), 22071–22080. https://doi.org/10.1073/pnas.1900654116 arXiv:1901.04592 [cs, stat].

[21] Quoc Hung Ngo, Tahar Kechadi, and Nhien-An Le-Khac. 2022. OAK4XAI: Model towards Out-Of-Box eXplainable Artificial Intelligence for Digital Agriculture. https://doi.org/10.48550/arXiv.2209.15104 arXiv:2209.15104 [cs].

[22] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 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 (KDD '16)*. Association for Computing Machinery, New York, NY, USA, 1135–1144. https://doi.org/10.1145/2939672.2939778

[23] USDA. 2024. USDA - National Agricultural Statistics.

[24] Kaihua Wei, Bojian Chen, Jingcheng Zhang, Shanhui Fan, Kaihua Wu, Guangyu Liu, and Dongmei Chen. 2022. Explainable Deep Learning Study for Leaf Disease Classification. *Agronomy* 12, 5 (May 2022), 1035. https://doi.org/10.3390/agronomy12051035 Number: 5 Publisher: Multidisciplinary Digital Publishing Institute.

[25] B. Üstün, W. J. Melssen, and L. M. C. Buydens. 2007. Visualisation and interpretation of Support Vector Regression models. *Analytica Chimica Acta* 595, 1 (July 2007), 299–309. https://doi.org/10.1016/j.aca.2007.03.023