


Towards the actual deployment of robust, adaptable, and maintainable AI models for sustainable agriculture

Giacomo Ignesti


0000-0003-2389-3086 


Institute of Information Science
and Technologies (ISTI)

National Research Council of Italy (CNR)
Via Giuseppe Moruzzi, 1, 56124 Pisa, Italy
University of Pisa

Second Floor, Largo Bruno Pontecorvo, 3, 56127 Pisa, Italy
Email: giacomo.ignesti@isti.cnr.it

Davide Moroni, Massimo Martinelli

0000-0002-5175-5126 

0000-0001-7419-5099 

Institute of Information Science
and Technologies (ISTI)

National Research Council of Italy (CNR)
Via Giuseppe Moruzzi, 1, 56124 Pisa, Italy

Email: {davide.moroni, massimo.martinelli}@isti.cnr.it

Abstract—In the past two decades, computer vision and artificial intelligence (AI) have made significant strides in delivering practical solutions to aid farmers directly in the fields, thereby contributing to the integration of advanced technology in precision agriculture. However, extending these methods to diverse crops and broader applications, including low-resource situations, raises several concerns. Indeed, the adaptability of AI methods to new cases and domains is not always straightforward. Moreover, the dynamic global panorama requires a continuous adaptation and refinement of artificial intelligence models. In this position paper, we examine the current opportunities and challenges, and propose a new approach to address these issues, currently in the implementation phase at CNR-ISTI.

Index Terms—Sustainable Agriculture; Artificial Intelligence; Deep Learning; Crowd-sensing; Citizen science

I. INTRODUCTION

IN RECENT years, the emergence of deep learning, combined with the increasingly widespread use of visual monitoring technologies for crops, has significantly contributed to the advancement of precision agriculture [1]. Uncrewed Aerial Vehicles (UAVs) equipped with colour or multispectral/hyperspectral cameras, as well as other robotic platforms designed for close-range operations with crops, have paved the way for the introduction of AI-assisted, data-driven approaches in agriculture [2]. This has permitted the implementation of precise monitoring, treatment, and harvesting techniques. However, these advancements have primarily impacted a narrow range of cultivated crops, particularly specialized ones yielding high revenues, such as high-end wine production [3].

It is clear that Artificial Intelligence (AI) and Machine Learning (ML), including Deep Learning (DL) methods, are versatile methodologies capable of being applied to disparate fields, including agriculture, where their potential impact has yet to be fully realized. However, the transfer from specific domains to new ones is not always feasible or cost-effective due

to the associated efforts required for designing and developing new models.

Numerous research and academic initiatives focus on a wide range of crops, encompassing intensive cultivation practices that have a significant impact on the global food supply [4]. In these contexts, DL models have demonstrated unparalleled performance on standardized datasets [5].

Concerning such topics, the state of research on AI applications in agriculture is wide. There is still no standardized approach, but the literature encompasses a lot of strategies that are all focused on improving crop quality and production. While modern DL models excel in image analysis for product quality enhancement, other critical agricultural domains, like water control [6], soil management, and production chain optimization, primarily rely on tabular data or emerging multi-modal approaches. Real-time object detection is a prominent AI application in agriculture, though classification algorithms often demonstrate superior performance [7] in specific contexts.

Recent works, as [8], try to employ knowledge-distillation techniques to improve weed mapping, adapting complex transformer architecture to the agricultural domain. At the same time, other studies [9] analyse various detection algorithms and design possible edge computing solutions for their real-time applications in precision agriculture. Image acquisition modality plays a pivotal role in plant analysis studies [10], as shown by the advancements in multi-modal imaging techniques that enhance the accuracy of trait estimation and facilitate the analysis of plant morphology and development. For instance, integrating visible light, fluorescence, and near-infrared imaging allows for a comprehensive assessment of plant structures, improving the segmentation and quantification of traits critical for phenotyping. These diverse imaging modalities not only provide complementary information. But also address challenges related to variable illumination and

plant colouration, ultimately leading to more robust phenotypic data extraction and analysis. Object detection and segmentation algorithms are usually more complex than their classification counterpart; therefore, translating these models and approaches into practical use for corporate farmers of all scales presents challenges, as real-world variability differs from the conditions in static benchmark datasets. To date, while there is a right to benchmarking agricultural datasets, no foundational models have been trained in this domain, making only possible transfer learning strategies and training from scratch solutions. Non-technological factors, including user acceptability, also hinder the widespread adoption of the latest research findings [7].

In this context, there is a growing need for developing new methodologies to overcome the current limitations of AI-assisted technologies. Specifically, these necessitate broadening their application to new crops and different scales of cultivation to support niche, small-scale, local, and organic productions while preserving biodiversity and environments through sustainable resource management. These demands come from various stakeholders, including farmers and policymakers (such as the European Community [11]). At the same time, they also originate from the Sustainable Development Goals set by the United Nations [12], particularly Goal 2 “Zero Hunger”. This goal includes targets such as doubling agricultural productivity (Target 2.3), ensuring sustainable food production systems, implementing resilient farming practices, and improving land and soil quality (Target 2.4), as well as maintaining genetic diversity through well-managed seed and plant banks (Target 2.5).

This *position paper* intends to present prospective ideas that might contribute to achieving the Sustainable Development Goals and fulfilling the requirements for the widespread adoption and implementation of practical artificial intelligence. While AI has potential applications across various domains, we focus specifically on using image-based intelligent systems to support farmers in their day-to-day operations. These systems can act as effective assistants, enabling informed decision-making and promoting the best practices for increased yet sustainable production.

The paper is organized as follows. In Section II, we critically review previous experiences, including ours, and highlight their limitations. In Section III, we enumerate a set of challenges and research questions that should be addressed to reach the scope described in this introduction. In Section IV, we analyze the current opportunities provided by technological advances and then explain the proposed approach rationale. Section V concludes the paper with remarks for further analysis and prospective implementation.

II. PREVIOUS EXPERIENCES

In light of advancements in image processing, computer vision, and machine learning, considerable research efforts have been directed toward developing intelligent systems to support agriculture. These efforts include the creation of algorithms for detecting, classifying, and quantifying crops and various

potential threats such as weeds, diseases, insects, and other stressors that could impact successful harvesting. The focus has been on analyzing remote sensing images captured by UAVs and close-range photography obtained through handheld devices or robot platforms.

The curation of benchmark datasets, particularly those released as open data, has played a pivotal role in enhancing the reproducibility and extensibility of research across different domains. Surveys on existing datasets, as documented in the work by Lu et al. [13], have become readily available. For instance, the PlantVillage dataset [14] has emerged as a de facto benchmark for leaf disease classification even though images, while numerous, may not fully represent the entirety of natural variability. Consequently, the performance of deep learning models on such datasets has been exceptional, with approaches achieving maximal accuracy levels [15].

Significant endeavours have been put forth within the AGROSAT+ project, sponsored by Barilla, to address detecting and classifying weeds. Under this initiative, collaborative efforts between CNR-ISTI and CNR-IBE have led to curating a dataset specific to cereal crop weeds [16]. This dataset might be valuable for weed detection and classification problems through close-range imaging or high-resolution UAV surveys. Additionally, its suitability for machine learning methods has been demonstrated in [17], where again the top performance was obtained. While intriguing and of great importance for advancing research, the current approaches have limitations regarding practical applications. The models’ ability to generalize when processing uncontrolled, real-world images is unsatisfactory, with a significant performance degradation of over 20%. This lack of reliability and inconsistent performance may be unacceptable to users in real-world deployments, leading to distrust in artificial intelligence and overall dissatisfaction, ultimately resulting in the technology’s failure to be adopted.

In the context of the AGROSAT+ project, an additional initiative was undertaken to address these challenges, leading to the development of an app called “GranoScan”. This app is designed to serve as an expert system that can be used directly in the field to identify plant diseases and stress, as well as detect weeds, insects, and other potential threats simply by using pictures captured through the smartphone camera. The app’s backend is driven by deep learning models that handle various visual recognition tasks [18]. One notable aspect of the app is its approach, which is somewhat independent of the specific computational models employed. In more detail, following an intensive period of initial data collection to train the machine learning models, GranoScan has now entered the production stage. Since then, a continuous stream of images from diverse users has been processed, with user consent, and stored to augment the dataset. This data has provided a wealth of information that can be leveraged to enhance and refine the models developed over the years using semi-assisted and semi-supervised methods. The experience is still ongoing.

III. CHALLENGES

Based on the previous experiences reported in Section II, a critical gap in the current AI technology for sustainable agriculture is the absence of a well-established methodology for the rapid deployment of models, namely of trained deep learning architecture for solving visual tasks related to agronomical problems. These AI models must satisfy various requirements, including robustness, adaptability, and maintainability, while being versatile enough to address various crops. Notably, the methodology should also ensure that the models can be easily transferred across different domains while maintaining their effectiveness and accuracy. For example, the models should be capable of adapting from one crop variety to another, regardless of similarities or differences in cultivation practices based on geographical location, climate, and other environmental conditions such as soil quality, water availability, and farming methods (e.g. organic, with biological or natural pest control, traditional). Developing such a methodology involves confronting several key challenges outlined below.

One of the primary challenges in deploying AI models in agriculture is the *limited availability of comprehensive and high-quality datasets*. Indeed, as shown in the survey [13], agricultural datasets, particularly those related to specific crops or regions, are often sparse, fragmented, or inconsistent (see, for instance, the dataset proposed for the challenge [19]). As we have seen, thanks to data augmentation strategies and the definition of ad hoc architectures, such a scarcity has not prevented the realization of performant AI models on static benchmark datasets. However, the generalization capabilities observed in practice have been, in our experience, somewhat disappointing.

Additionally, the agricultural environment is highly dynamic and is influenced by seasonal variations, pest outbreaks, and other temporal factors. To ensure that AI models can effectively generalize, it is crucial to train them using data collected over multiple growing seasons in order to capture these variations accurately. *Longitudinal studies* that span several agricultural cycles can provide valuable insights into long-term trends and enhance the model's ability to generalize across different conditions and time periods. Such longitudinal assessment is feasible when analyzing routine remote sensing images captured by satellite-borne sensors. However, when considering the smaller scale of details (e.g., airborne sensors and close-range images), there are currently no relevant and accessible datasets that span multiple harvest seasons.

Climate change introduces significant unpredictability into agricultural systems, affecting crop yields, pest prevalence, and overall farm productivity. For this reason, AI models need to be capable of not only interpolating within the known data but also *extrapolating* to predict the impact of unprecedented climate scenarios. This is a feature that should be taken into account when selecting the deep learning architecture or other machine learning paradigm to be used in a classification or regression task. Indeed, some methods are only suited to analyze data within the convex hull of the training set, producing

in output something within the convex hull of the labels in the training data. Although most of the classification and object detection tasks are not apparently conditioned by these issues, in general, reasoning about crop status, these issues should be taken into account. In particular, this might require integrating climate models with agricultural data to create AI systems that can adapt to changing climatic conditions and provide reliable recommendations for farmers.

The ultimate objective of utilizing AI-based systems in agriculture is to convert predictive insights into *actionable knowledge* that farmers can easily put into practice. This involves not only creating user-friendly interfaces and providing effective training for farmers but also ensuring that the AI recommendations are reliable, practical, and economically feasible. In addition, there is a need for processes that facilitate continuous feedback from the field to refine and update the models, ensuring that their relevance and accuracy in real-world applications remain stable without being affected by potential non-stationary conditions.

IV. PROPOSED APPROACH

Having discussed the challenges towards the implementation and actual deployment of robust, adaptable, and manageable AI models for tackling agronomic tasks, it is important to note that several opportunities are linked to technological advances that can ease the identification of possible solutions.

From one side, indeed, there has been a flourishing of research towards identifying highly efficient and robust AI models with improved insensitivity to data variability [20].

Secondly, methods have also been analyzed from the point of view of carbon footprint, [21] taking into account not only the training and inference costs but also the overhead linked, for instance, to data transfer. This is an aspect in deciding where to collocate computationally intensive tasks over the computational continuum, determining whether to process directly near the node where the data has been captured (i.e. directly on the smartphone capturing the image or on a robotic platform) with no transfer overhead or, conversely, on the cloud (with variable transfer costs). In such a context, progress in hardware also allows for more freedom in such design choices, given the general availability of computational resources, including GPU resources, along the computational continuum.

Finally, a third opportunity arises from the successful implementation of crowd-sensing that can be attributed to two key aspects: - the first aspect is technical, in which modern accessible devices, such as smartphones, now offer enhanced sensing capabilities, including LiDAR technology, multiple camera lenses, and advanced geolocation features; - the second aspect relates to the growing awareness and willingness of individuals to participate in citizen science initiatives.

In this section, we propose the envisaged rationale and then discuss in detail the three main points it leverages.

A. Rationale

The rationale of the approaches is based on the use of three main levers that are considered to be able to effectively

contribute to fast and efficient deployments, respecting the requirements discussed in the previous section. The first aspect is based on the provision, not only of statistic classification or number produced by ML/DL models, but also on integrating these methods with Decision Support Services (Section IV-B). This is envisaged to respond to the need to translate insights into actionable knowledge. Indeed, not only the output of the image processing will be produced, but it is necessary to accompany this output with an explanation (in an explainability effort) and suggestions on how this output may be used in practice to optimize treatment, for instance, by devising an adaptive treatment plan. Secondly, a better tradeoff between performance and generalization capabilities should be sought (Section IV-C). This attains research efforts in ML/DL where new methods that have already proved promising, based on ensembling, can achieve improved generalization capabilities and allow for a faster domain transfer. A third ingredient is represented by a more strategic approach to filtering crowd-sensed information, considering uncertainty in their evaluation since they originate from non-authoritative sources (Section IV-D). In this case, new methodologies can be enlisted to determine data quality and define the confidence level the new data has to enter into the decisional processes.

In the current envisaged activities such rationale is going to be validated (see Figure 1) in a variety of cases addressing i) plant position detection, ii) plant count, iii) control of the growing phase (e.g. pre/post-germination, developed, budding, pre-flowering, pre-fruiting, ripening depending on the cultivation) and iv) anomaly detection (abnormal growth compared to market standards, sufficient/insufficient gems,...) and v) plant threats (weeds, pests, and diseases). In addition, vi) time (of budding, flowering, fruiting, ripening,...) and vii) and volume predictions (number of plants/flowers/fruits/biomass) as well as viii) quality of the final product will be considered.

B. Integration with DSS

A DSS must consider several factors depending on the plant species, including sprout number, flowering time, loss of first flowering, and other variables.

The proposed model envisages the DSS's intervention point as at least twofold. Indeed, the DSS intervenes before and after the AI models, (a) first to decide which ones to run based on historical data, context conditions, seasons, situations, and others, and (b) then to provide suggestions based on the results.

A hybrid DSS integrates different technologies and information types in order to provide greater flexibility, scalability and efficiency in helping make the right decision, the correct application, and the proper treatment in the right place and at the right time: knowledge-based modules allow a semantic representation of data to extract and infer helpful information and can include Data Mining and predictive analytics to identify hidden patterns and relationships between data, providing high quality and a clear explanation of decisions; model-based modules allow optimizing the internal decision processes by analyzing specific issues, such as the irrigation scheduling or the crop prediction processing data, when the

target audience/stakeholder is not interested in understanding the decision-making process but only in the results produced.

DSS can also be utilized to communicate and present information as needed. For example, AI tools that predict future outcomes based on historical data and trends (e.g. forecasting flowering or fruiting times or spreading a disease) can be activated proactively in response to specific events. The resulting output can then be promptly presented to the user, allowing for optimal treatment and harvesting planning.

C. Model adaptation, generalization capabilities and continuous learning

A key factor is the ability of AI models to adapt to new data, to generalize their knowledge, to apply them to new contexts, to ensure that models are able to function properly in different situations and to improve their accuracy over time. At the same time, we need a model capable of learning fast without relying on an extensive corpus of knowledge, represented in this specific case by a dataset annotated with ground truth. In DL, the capabilities of transfer learning are well known: deep models trained on a dataset belonging to a certain domain, often general purpose such as in the case of ImageNet, are then capable of adapting more quickly and with better performance to new domains with respect to the same architectures initialized in a random way. In addition, zero-shot and few-shot learning have been considered in several contexts, achieving classification with minimal training data [22].

In our view, we aim to address these elements by exploiting *adaptive ensembling* and *continuous learning*.

More specifically, in adaptive ensembling [5], a few weak models are trained in parallel, resulting in a set of specialized modules. Such weak models are based on DL models and, specifically, in architectures belonging to the EfficientNet family [23]. As such, they comprise a first set of layers, performing feature extraction and a final layer-producing classification. In our approach, such weak models are combined together to produce a strong classifier at the deep feature level. Namely, the original classification layer of each weak model is neglected, and a new global classification layer, taking into input the concatenation of the feature vectors provided by all the weak models, is introduced and trained to obtain the desired *ensemble*. Such an approach has been proven to give promising results in domain adaptation, as in the case of olive diseases [7], but extended analysis and diverse dataset partition methodology should be studied to assess the added value in robustness.

Ensembling is also suitable to support continuous learning. As already introduced based on the yearly campaign or on a steady stream of data coming from the field, the concept of a static dataset has to be surpassed. The data flow indeed offers the opportunity to update models based on deep learning to provide increasingly accurate answers by taking advantage of the expansion of the available case studies. To this end, it is neither practical nor convenient to retrain the models from

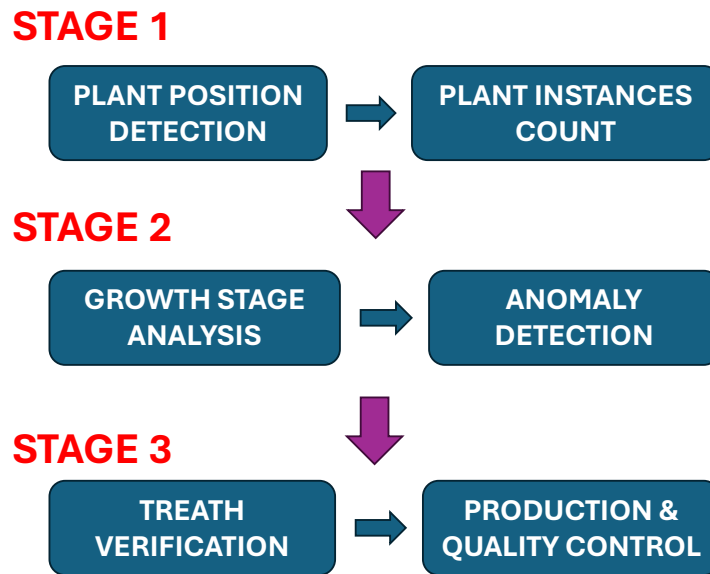


Fig. 1. Key steps diagram of a possible chain of activities as a rationale for plant monitoring and analysis

scratch at each update, but it is advisable to use a continuous learning approach.

The possibility of shifting toward a continual learning paradigm has significant potential: beyond providing constant retraining, it also enables enhanced models through continuous updates, making the system more resilient to unseen threats. This approach is more accurate and trustworthy than considering all boundary conditions simultaneously. While classical supervised deep learning algorithms can detect seasonal patterns, they often fail to accurately predict anomalous conditions. Moreover, they often fail to detect points of instability, which can adversely impact the evolution of the studied environment and potentially lead to catastrophic consequences.

From a technical perspective outside of research contexts, the use of ensemble methods is often not aligned with company objectives and means because it requires continuous resources. Other strategies, such as using state-of-the-art machine learning models with a priori studies of data distribution, can effectively produce one-shot models with an initial better performance. Ultimately, the support from advanced techniques demonstrates that moving beyond conventional methods can lead to developing more effective models, such as those achieved through ensemble approaches.

In the main studies of the AGROSAT+ project, it has become clear that transferring technology and know-how from the public to the private sector plays an important role. Even though large companies have the possibility and the means to sustain the production of high levels, in AGROSAT+, the

resources employed in the developed DL model are far lower than the computational necessity of Large Language Models (LLMs). Indeed, the training of a state-of-the-art model [15] required only a mid-range workstation (equipped with two RTX QUADRO 5000 GPUs, which have now become an example of affordable accelerators), and the inference of the trained model worked on the CPU of this machine. This API solution lets users control their production directly with their phone (basic technologies approach). The proposed ensemble model was also used successfully in other scientific fields [24], showing the potential of open-access research.

Accuracy, Precision, Recall, F1 & R1 score and any other method largely treated in the statistical literature are the main methods to evaluate the goodness of a DL model. Still, the black-box nature of these algorithms hinders trust in their performance. The public is sceptical of their benefits since it is impossible to fully understand their inner working. For the same reason, the scientific community, with their government counterparts, is questioning the danger, limits, and rightfulness of the DL models. Good practices, such as strict control of no train-test data contamination, augmentation strategies, and eXplainable Artificial Intelligence (XAI), are common methods to ensure that the systems are accurate but also trustworthy and plausible. Knowledge-based DL algorithms are other possible solutions; in genomic and molecular biology, AlphaFold [25] is a good example of how to evaluate the quality of a model. AlphaFold architecture combines the transformer attention mechanism in pairs with the Evoformer

module; this processes correctly evaluate the data of the biological sequence and the pair representation to output a new possible structure. Another possible solution is Physics-Informed Neural Networks (PiNNs) [26] that guide the systems' output towards valid output thanks to the incorporation of the boundary conditions of the described problem. The listed procedures suggest that leveraging information from crop traits could provide an intrinsic validation method for the model, as the proposed approach aligns with natural observations. Last, it is worth mentioning the possible benefits of incorporating continual learning strategies to validate the model over time. Continual learning enhances the adaptability of DL algorithms by enabling them to incrementally acquire information from new data while retaining the old ones used for the previous state. This approach not only mitigates the risk of catastrophic forgetting but also allows for dynamic updates, thereby outputting an unbiased overall accuracy of real-world phenomena. Consequently, the ability to control and fine-tune the model's performance across diverse tasks and datasets is significantly enhanced, ensuring that the model remains robust and effective over time.

D. Filtering and analysis of crowd-sensed data

In our previous experience with the AGROSAT+ project, researchers dealt with the quality of data collected from voluntary users. While the information provided, including new images to enhance the datasets, was effective in meeting the need for more varied spatial and temporal data, it is essential to implement suitable filtering to avoid errors or biases due to the non-authoritative nature of the information. To this end, one of the first elements integrated into the GranoScan app is a deep learning method, achieved through supervised learning, to differentiate relevant images from those that may not be suitable for a specific computer vision task. However, this approach can be improved and expanded by: a) incorporating blind general-purpose image quality assessment methods, such as those based on deep learning (e.g., [27], [28]), and b) developing appropriate object detectors to verify that the image is relevant to the computer vision task (for example, if the visual task involves identifying leaf diseases, there should be at least one leaf in the picture, and it should occupy a significant area). After passing through the specified filters and if the user provides feedback, the processed image can be stored in an expanded version of the datasets suitable for potential model updates and fine-tuning, also according to online procedures and to the continuous learning approach described in Section IV-C. Furthermore, additional filtering should be conducted to analyze the cross-correlation between contextual and image data. This is primarily focused on identifying potential anomalies within the data, such as a disease reported in a region of the world or during a time of year when the disease is not expected. While such anomalies may indicate the nonstationarity of the observed global situation (also as an outcome of climate change), they should be carefully reviewed by additional AI agents and, ultimately, by human observers. This is somewhat related to the continuous monitoring of

the expert system in the operational phase to prevent biases and drifts and contribute to the overall maintainability of the system.

V. CONCLUSIONS

In this position paper, we have revised and enumerated challenges and opportunities for developing AI models that can tackle visual tasks relevant to agronomy. These models must exhibit high levels of robustness, adaptability, and maintainability to be considered trustworthy for deployment across various scenarios. Our proposed approach focuses on three key elements: developing technologies for model domain adaptation, utilizing crowd-sensing with awareness of uncertainty, and integrating with reasoning and recommendation systems to transform computational intelligence outputs into actionable knowledge. The synergy among these three points is also inspired by the general principles of responsibility, accountability, explainability, and trustworthiness, which collectively enhance the acceptability of our proposed solutions by addressing both technical and non-technical requirements.

Work is currently underway as part of the STRIVE project, and it will continue over the next two years. During this time, experiments will be conducted to test the proposed approach, evaluate its effectiveness, and understand its limitations. Additional measures will involve working with the community of farmers to raise awareness and encourage engagement.

The additional benefits of engaging the farming community in this precision approach to agricultural practices include building trust and improving perceptions of this tool.

An active community can guarantee a steady flow of data, allowing the continual learning implementation part of our solution, and communicate additional information, enabling real-time adjustments to the predictive component of the employed algorithm, which would otherwise not be possible.

In the future, we may consider utilizing Generative AI and LLMs to enhance communication and interaction with end users. However, we will proceed cautiously, as these technologies are not yet fully mature, language support is not consistent, and the portability and sustainability of the technology still need to be assessed. Therefore, we may need to postpone their application in scenarios where actual deployment is being pursued.

ACKNOWLEDGMENTS

This work has been partially supported by the CNR FOE Project DIT.AD022.207 "La Scienza per le TRansizioni Industriale, Verde, Energetica - STRIVE" (Science for Industrial, Green and Energy Research).

This work is part of a project that has received funding from the European Union's Horizon Europe - Topic: DIGITAL-2022-SKILLS-03-SPECIALISED-EDU - Project ID: 101123258 - Project name: Digital agriculture for sustainable development(AGRITECH EU).

REFERENCES

- [1] T. Saranya, C. Deisy, S. Sridevi, and K. S. M. Anbananthen, "A comparative study of deep learning and internet of things for precision agriculture," *Engineering Applications of Artificial Intelligence*, 2023. doi: 10.1016/j.engappai.2023.106034
- [2] I. Zualkernan, D. A. Abuhani, M. H. Hussain, J. Khan, and M. El-Mohandes, "Machine learning for precision agriculture using imagery from unmanned aerial vehicles (uavs): A survey," *Drones*, 2023. doi: 10.1016/j.compag.2020.105760
- [3] S. Koul, "Machine learning and deep learning in agriculture," *Smart Agriculture: Emerging Pedagogies of Deep Learning, Machine Learning and Internet of Things*, 2021. doi: 10.1201/b22627-1
- [4] R. Priya and D. Ramesh, "ML based sustainable precision agriculture: A future generation perspective," *Sustainable Computing: Informatics and Systems*, 2020. doi: 10.1016/j.suscom.2020.100439
- [5] B. Antonio, D. Moroni, and M. Martinelli, "Efficient adaptive ensembling for image classification," *Expert Systems*, 2022. doi: 10.1111/exsy.13424
- [6] Ł. Błaszczuk, M. Mizura, A. Płocharski, and J. Porter-Sobieraj, "Simulating large-scale topographic terrain features with reservoirs and flowing water," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F2137
- [7] A. Bruno, D. Moroni, and M. Martinelli, "Efficient deep learning approach for olive disease classification," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F4794
- [8] G. Castellano, P. De Marinis, and G. Vessio, "Applying knowledge distillation to improve weed mapping with drones," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F960
- [9] N. Iqbal, C. Manss, C. Scholz, D. König, M. Igelbrink, and A. Ruckelshausen, "Ai-based maize and weeds detection on the edge with cornweed dataset," in *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, 2023. doi: 10.15439/2023F2125
- [10] S. Kolhar and J. Jagtap, "Plant trait estimation and classification studies in plant phenotyping using machine vision—a review," *Information Processing in Agriculture*, vol. 10, no. 1, pp. 114–135, 2023.
- [11] European Commission, "Work programme 2023-2025 – 9. food, bioeconomy, natural resources, agriculture and environment," 2024. [Online]. Available: https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/wp-call/2023-2024/wp-9-food-bioeconomy-natural-resources-agriculture-and-environment_horizon-2023-2024_en.pdf
- [12] "THE 17 GOALS," <https://sdgs.un.org/goals>, 2015, [Online; accessed 11-June-2024].
- [13] Y. Lu and S. Young, "A survey of public datasets for computer vision tasks in precision agriculture," *Computers and Electronics in Agriculture*, 2020. doi: 10.1016/j.compag.2020.105760
- [14] D. Hughes, M. Salathé *et al.*, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," *arXiv preprint arXiv:1511.08060*, 2015. doi: 10.48550/arXiv.1511.08060
- [15] A. Bruno, D. Moroni, R. Dainelli, L. Rocchi, S. Morelli, E. Ferrari, P. Toscano, and M. Martinelli, "Improving plant disease classification by adaptive minimal ensembling," *Frontiers in Artificial Intelligence*, 2022. doi: 10.3389/frai.2022.868926
- [16] R. Dainelli, M. Martinelli, A. Bruno, D. Moroni, S. Morelli, M. Silvestri, E. Ferrari, L. Rocchi, and P. Toscano, "A phenotyping weeds image dataset for open scientific research," 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.7598372>
- [17] —, 49. *Recognition of weeds in cereals using AI architecture*. Wageningen Academic, 2023.
- [18] R. Dainelli, A. Bruno, M. Martinelli, D. Moroni, L. Rocchi, S. Morelli, E. Ferrari, M. Silvestri, S. Agostinelli, P. La Cava *et al.*, "Granoscan: an ai-powered mobile app for in-field identification of biotic threats of wheat," *Frontiers in Plant Science*, 2024. doi: 10.3389/fpls.2024.1298791
- [19] T. Kattenborn, "Planttraits2023," 2023. [Online]. Available: <https://kaggle.com/competitions/planttraits2023>
- [20] N. Drenkow, N. Sani, I. Shpitsler, and M. Unberath, "A systematic review of robustness in deep learning for computer vision: Mind the gap?" *arXiv preprint arXiv:2112.00639*, 2021. doi: 10.48550/arXiv.2112.00639
- [21] K. Kirkpatrick, "The carbon footprint of artificial intelligence," *Communications of the ACM*, 2023. doi: 10.1145/3603746
- [22] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM computing surveys (csur)*, 2020. doi: 10.48550/arXiv.1904.05046
- [23] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*, 2019. doi: 10.48550/arXiv.1905.11946
- [24] A. Bruno, C. Caudai, G. R. Leone, M. Martinelli, D. Moroni, and F. Crotti, "Medical waste sorting: a computer vision approach for assisted primary sorting," in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, 2023. doi: 10.48550/arXiv.2303.04720
- [25] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko *et al.*, "Highly accurate protein structure prediction with alphafold," *nature*, 2021. doi: 10.1038/s41586-021-03819-2
- [26] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, 2021. doi: 10.1038/s42254-021-00314-5
- [27] W. Hou, X. Gao, D. Tao, and X. Li, "Blind image quality assessment via deep learning," *IEEE transactions on neural networks and learning systems*, 2014. doi: 10.1109/TNNLS.2014.2336852
- [28] S. Bianco, L. Celona, P. Napoletano, and R. Schettini, "On the use of deep learning for blind image quality assessment," *Signal, Image and Video Processing*, 2018. doi: 10.1007/s11760-017-1166-8