KI 2023 – Doctoral Consortium Extended Abstract: Understanding Agricultural Landscape Dynamics with Explainable Artificial Intelligence

Stefan Stiller[1 2]

Keywords: deep learning; precision agriculture; crop yield; explainable artificial intelligence; spatial cross validation; self-supervised learning

1 Introduction

Deep learning (DL) models, particularly those utilizing computer vision techniques such as proximal and remote sensing imagery, have witnessed extensive utilization within agriculture [KPB18]. These DL applications encompass diverse areas, including land cover and crop type mapping [Ku17], crop yield estimation [KS15, NNL19], drought monitoring [Sh19], plant disease spread analysis [Te20], and overall monitoring of agricultural systems. DL applications offer significant potential to enhance agricultural practices at various scales, spanning from individual organisms, field, landscape, to regional and even continental scales [Ry22b].

2 Motivation

Ideally, DL models are highly predictive in settings beyond the training data (cf. model generalization), and the reasoning behind why a model made certain predictions are readily understandable (namely model explainability). In agriculture, for example, farmers expect the model to predict accurately in the following year, neighboring fields or across different management practices (cf. model transferability). Moreover, for crop yield prediction for instance, the models are expected to utilize features for predicting that are meaningful representations for crop yield like the number of plants.

To obtain a reliable performance estimate and assess the generalizability and transferability of prediction models, the conventional approach is to utilize random cross-validation (RCV)

[1] Research Platform “Data Analysis & Simulation”, Leibniz Centre for Agricultural Landscape Research (ZALF), Müncheberg, Germany, stefan.stiller@zalf.de

[2] Environment and Natural Sciences, Brandenburg University of Technology Cottbus-Senftenberg, Cottbus, Germany

doi:10.18420/ki2023-dc-09

[St74]. RCV provides an estimate of the model's capability to make accurate predictions on new, unseen data. However, spatial data exhibits an inherent structure known as spatial autocorrelation, wherein data points in close proximity tend to possess more similar values [To70]. Numerous studies have highlighted that spatial autocorrelation violates the assumption of data independence between the training, validation, and test sets, thereby leading to inflated inferences regarding prediction performance [Le14, Ro17, Pl20, Ka22]. Consequently, these studies propose the adoption of spatial cross-validation (SCV) as an alternative to RCV to address this issue of over-optimism. However, whether SCV is a better approach than RCV for building a DL model for agricultural prediction is largely untested. In other words, spatial dependence within DL applications has been rarely studied, particularly in agriculture for crop yield mapping.

Moreover, model generalization and transferability are affected by whether the model can learn meaningful feature representations of the studied phenomenon from the training data. For example, in agriculture a crop yield prediction model might generalize better and be more readily transferable when it understands the concepts of plants, fruiting bodies and plant health. For a model to learn generalizing features it typically needs many labeled training data [He23]. Yet, the availability of labeled training data is limited in agriculture. To facilitate learning in a limited data scenario, transfer learning and self-supervised learning (SSL) have been discussed as pre-training strategies [Ya20]. Both strategies exploit that a particular type of DL architectures, convolutional neural networks (CNNs; [Le89]), learn feature representations hierarchically [JmMhY16] by either pre-training on datasets across domains (cf. transfer learning; [BF76]) or pre-training on that same dataset (cf. SSL; [Ba23]). Transferring the acquired knowledge to agricultural tasks can improve the performance and robustness of crop yield prediction models. Yet, how to effectively train a DL model with a relatively small amount of training data is an unsolved issue in the agricultural domain. Exploring the potential of model pre-training in the context of crop yield mapping, especially, remains an underexplored area.

A good DL model should learn meaningful feature representations for a better generalizability and transferability, but how can one confirm whether training is done appropriately? DL modeling is considered as a black-box modeling, and unlike conventional parametric statistics, it is difficult to explain how the model behaves. Model explanations (cf. explainable artificial intelligence, XAI; [Ba20, Mo21]) aim to bridge the gap between the predictive power of DL models and the need for human-understandable explanations. In precision agriculture applications, such as crop disease detection or yield prediction, the interpretability offered by XAI can be crucial for understanding and validating AI model outputs. Nevertheless, the application of XAI in the agricultural domain is still rare [Ry22a]. The interpretability of AI models can offer valuable insights into the features contributing to, for example, crop yield predictions.

Addressing these research gaps in model transferability, pre-training strategies and XAI is vital for advancing the field of crop yield mapping in agriculture. It will facilitate more accurate predictions, improved decision-making, and the adoption of transparent AI technologies in the agricultural sector.

3 Research Question

The overarching aim is to study how to build a DL model for crop yield prediction using a relatively small dataset, while accounting for model generalizability, transferability, and explainability. In specific, I address the following research questions:

1. How does the RCV and SCV as model training and validation techniques affect generalization and transferability of crop yield prediction models?
2. How does model pre-training, either with cross domain transfer learning or self-supervised, affect model training and validation?
3. In which way are spatial composition and configuration, i.e. pixel arrangement and values, important for predicting?

4 Related Work

DL based crop yield prediction models that use image analysis like remote sensing are increasingly applied in precision agriculture with 70 publications in 2022. Different architectures are used, predominantly CNNs and long short-term memory (LSTMs) [BS22]. Common predictor variables that these models use are vegetation indices like Normalized Difference Vegetation Index (NDVI) that capture plant health; less often RGB imagery is used [WJD20]. Applications range through different spatial scales from country [Zh22] to field size [NNL19].

For model training and evaluation, spatial dependence has been suggested to be the cause of model overfitting by a series of studies [Le14, Ro17, Pl20]. [Ro17] suggest SCV as a means to tackle spatial autocorrelation. [Ka22] argue that CNNs are especially prone to overfit to the spatial structure, as they exploit image patterns for superior prediction power. In agriculture prediction models are for the most part trained and evaluated with hold-out validation with random sampling [Ge21] or RCV [NNL19, Va19], but only a few mapping studies employ SCV to engage spatial autocorrelation [No21].

Two model pre-training paradigms exist: cross domain transfer learning [BF76] where models are pre-trained on one dataset and are then transferred to another, and several families of SSL strategies [Ba23]. In precision agriculture, transfer learning was applied for example for semantic segmentation of crops versus weeds between crop types [Bo20]. Another study demonstrated high performance of cross domain transfer learning for plant leaf disease identification with a model that had been pre-trained on a classification challenge image dataset [Ch20a]. SSL is relatively new in agriculture and with 11 studies until the beginning of 2023 has rarely been studied. [GN21] used SSL to segment and classify agricultural images and showcased how SSL increased classification performance. Another study from [Ka21] showed how SSL improves agricultural pest classification.

XAI techniques have gained significant attention in various domains, including agriculture. They are predominantly used to understand complex processes and uncover underlying relationships from tabular datasets [Mo21]. In agriculture [Ry22a] demonstrated for example, how XAI methods can be used to study under which conditions crop yield change is positive when changing the field management practices. XAI with computer vision tasks in agriculture are at plant individual level, for example for plant stress phenotyping [Gh18] or disease identification [Na19]. Interpretable crop yield models using LSTM neural networks identified critical periods of the crop growth cycle [Ma23]. This study aims to fill the gap of spatial crop yield prediction with CNNs.

5 Approach

5.1 Research Site and Data Collection

The data was collected at the p̈atchCROPëxperimental agricultural field [Gr21] located in Brandenburg, Germany. The field spans 70 hectares and is subdivided into 30 small field arrangements, each approximately 0.5 hectares in size, the experimental site features unique treatments in terms of crop rotation, soil quality, and management practices. The patchCROP field experiment aims to study diversified agricultural landscapes, particularly the effects of patch cropping on yield stability, farming system resilience, ecosystem services, and biodiversity. I only use a subset of all 30 field arrangements because some plots had been harvested already. For this project I selected three crop types from six field arrangements in the cropping season 2020: soy, maize, and sunflowers.

The dataset comprises a high-resolution Unoccupied Aerial Vehicles (UAVs)-based RGB image from a date closest to harvest and between 100 and 300 crop yield point recordings as ground truth. RGB images were taken with a drone s̈enseFlyeBee Xẅith the camera model s̈enseFlyS.O.D.A.över the experimental field site on August 6th 2020. The overflight was performed at 84m height. The image was taken with an average ground sampling distance of 2.22 cm and recorded in the EPSG 25833 coordinate system. The crop yield points have been automatically recorded and georeferenced by the combine harvester during harvest. They have been cleaned of erroneous data [Bl99, AC02, LBO14] and spatially interpolated with ordinary kriging [Cr88] to yield maps. The datasets to train and validate the DL models are generated by sliding window sampling (224px x 224px) [Va19] on the remote sensing image and ground truth, with the average crop yield within a sliding window as image label (n=4761 per field arrangement).

5.2 Deep Learning Modeling

I use CNNs [Le89] for crop yield prediction. The models utilize RGB images captured by UAV for predicting yield from multiple crop types, including soy, sunflowers, and maize. Two architectures are tested and compared: ResNet18 [He16] and a re-implementation of a crop yield mapping model proposed by [NNL19]. Critical hyper-parameters for DL are tuned with Bayesian Optimization [FKH18] using 'flat' CV [WC21]. On the training dataset I use augmentations in order to facilitate learning and prevent overfitting to spatial characteristics like planting rows or different lighting conditions [SK19, Bu20].

5.3 Random and Spatial Cross Validation

I test and compare DL-based yield mapping models using 4-fold RCV [St74] and SCV [Ro17]. For that, the dataset is divided into four equal partitions. In each iteration, one alternating partition is assigned as the validation set, while the remaining partitions are combined as the training set. For SCV the data is partitioned by spatial blocking, whereas for RCV by random sampling. Overlapping samples between the training and validation set are removed for SCV. For SCV the combined training set comprises 3279 and the validation set 930 samples, whereas for RCV the training set size is 3570 and validation set size 1191. Model prediction performance on the same field is assessed globally over every data point predicted during CV [Me19]. To assess model generalization, I conduct external validation on another field of the same crop type (test set size n=4761). By testing the models on unseen fields, I aim to evaluate their performance in different spatial contexts and distributions.

5.4 Transfer Learning and Self-supervised Pre-Training

Pre-training models can facilitate generalization and training despite smalldata scenarios [Sa23]. I test two distinct pre-training strategies on crop yield prediction models: First, I test cross-domain transfer learning [BF76] with models that have been pre-trained on the image classification dataset ImageNetV2 [Re19]. Self-supervised pre-training methods [Ba23] use the same dataset for pre-training, in which they aim to learn generalizable features. These methods have been shown to be sensitive to image content [Ba23], with less performance on natural images. Here, I test different methods from distinct schools of self-supervised pre-training. For example contrastive learning (SimCLR; [Ch20b]) learns visual representations by encouraging similarity between two augmented views of an image or, from different schools, SimSiam [CH20c] that uses simple Siamese networks to learn meaningful representations, or VICReg [BPL22] that exploits co-variance between features.

5.4.1 Model Explanations

In order to gain an understanding for the predictive behavior of crop yield mapping models that utilize remote sensing imagery I use XAI techniques [Ba20, Mo21]. I aim to use XAI techniques to visualize important features for predicting (cf. feature attribution). Feature attribution methods are either model agnostic, in which they explain model prediction behavior by adding post-hoc explanations, or model specific, in which they leverage model characteristics for explaining. A prominent model agnostic feature attribution method is RISE [PDS18]. RISE estimates feature importance empirically by probing the model with randomly masked image versions and recording the changes to the output. Grad-CAM is a model specific gradient based approach [Se20]. Grad-CAM uses gradients of target concepts (such as high or low yield) to produce a localization map of important features. Combinedly, RISE and Grad-CAM offer visual explanations for crop yield mapping models. These explanations can facilitate the identification and understanding of underlying influences on the predictions.

6 (Planned) Evaluation

6.1 Chapter I: Spatial cross validation improves deep learning model transferability for spatial prediction in agriculture [manuscript in preparation]

Recent studies have demonstrated the impressive predictive capabilities of DL models in capturing spatial environmental patterns using remote sensing imagery. However, these studies commonly rely on RCV for training and validation, neglecting the impact of spatial autocorrelation, which can lead to overfitting. Additionally, the transferability of models across different sites has been rarely investigated. This study aims to address these gaps by investigating the effectiveness of SCV in building more robust DL models for spatial prediction. I compared the performance of DL models, utilizing two different algorithms, in predicting crop yield for three crop types using RCV and SCV. High-resolution aerial RGB imagery obtained from drone surveys was employed as input data. The findings reveal that the models trained with SCV consistently outperform those trained with RCV when tested on external fields, exhibiting an average improvement of 13% (with correlation coefficients of 0.28 for SCV and 0.14 for RCV), even though models using SCV exhibit lower performance in CV than those using RCV (with a correlation coefficient of 0.658 compared to 0.983 for RCV). The results suggest that RCV tends to overfit the spatial structure and memorize image-specific information, leading to overoptimistic outcomes. This study provides empirical evidence in the field of agriculture, highlighting the superiority of SCV over RCV in enhancing the generalizability and transferability of DL models.

6.2 Chapter II: Self-supervised learning with small data predicts fine-scale crop yield heterogeneity for smallholder agriculture

In smallholder agriculture, limited availability of labeled data poses a challenge for developing accurate crop yield prediction models. This study investigates the effectiveness of traditional transfer and SSL techniques, specifically SimCLR, SimSiam and VICReg, in predicting fine-scale crop yield heterogeneity using high-resolution UAV-based RGB images in smallholder agriculture. I investigate the impact of model pre-training, including cross-domain transfer learning and SSL, on model training and validation. By leveraging the inherent structure of unlabeled data, SSL aims to capture meaningful representations, enabling better generalization to small-scale agricultural settings with limited labeled data. I compare the performance of SSL models with traditional transfer learning approaches, assessing their ability to capture nuanced variations in crop yield. The findings contribute to improving precision agriculture and enhancing productivity and sustainability in smallholder agricultural systems.

6.3 Chapter III: Explaining concepts and features in DL based crop yield prediction

DL models have demonstrated remarkable performance in crop yield prediction, but due to their complex and opaque nature they are inherently not interpretable. In this chapter, I focus on the interpretability of DL-based crop yield prediction models. Specifically, I aim to explain and interpret identified best approaches from previous chapters. I add visual explanations and counterfactual explanations for yield prediction using RISE and Grad-CAM to understand the predictive behavior of the model. Explanations as well as counterfactual explanations are important in order to understand what influences for example high crop yield. These visual explanations are to be evaluated qualitatively and quantitatively. Explanations are quantitatively evaluated according to faithfulness [BWM20, DFM22] and localization [Ar22]. Moreover, for the final evaluation, I aim to conduct a user study, for which I invite agricultural experts and practitioners to discuss and interpret study outcomes. Such a discussion with domain experts is crucial for linking model explanations to domain knowledge. For example, the model might capture influences by plant density, plant health, nutritional competition with weeds, soil texture or management paths. Thus, model explanations and interpretations of these together with domain experts give crucial insights into crop yield's spatial heterogeneity. Model explanations provide valuable interpretability tools for precision agriculture and as a decision basis for agricultural management.

7 Acknowledgements

I thank Masahiro Ryo and Gohar Ghazaryan for supervision and support. Kathrin Grahmann for the data and encouragement. Lars Richter, Mohammad Amin Mohammadibanadaki for data collection and processing. This study was supported by the BMBF funded project 'Multi-modale Datenintegration, domänenspezifische Methoden und KI zur Stärkung der Datenkompetenz in der Agrarforschung (KIKompAG)' (16DKWN089). The maintenance of the patchCROP experimental infrastructure is supported by the Leibniz Centre for Agricultural Landscape Research.

Bibliography

[AC02] Arslan, Selcuk; Colvin, Thomas S.: Grain Yield Mapping: Yield Sensing, Yield Reconstruction, and Errors. Precision Agriculture, 3(2):135–154, 2002.

[Ar22] Arias-Duart, Anna; Parés, Ferran; Garcia-Gasulla, Dario; Gimenez-Abalos, Victor: , Focus! Rating XAI Methods and Finding Biases, 2022.

[Ba20] Barredo Arrieta, Alejandro; Díaz-Rodríguez, Natalia; Del Ser, Javier; Bennetot, Adrien; Tabik, Siham; Barbado, Alberto; Garcia, Salvador; Gil-Lopez, Sergio; Molina, Daniel; Benjamins, Richard; Chatila, Raja; Herrera, Francisco: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58:82–115, 2020.

[Ba23] Balestriero, Randall; Ibrahim, Mark; Sobal, Vlad; Morcos, Ari; Shekhar, Shashank; Goldstein, Tom; Bordes, Florian; Bardes, Adrien; Mialon, Gregoire; Tian, Yuandong; Schwarzschild, Avi; Wilson, Andrew Gordon; Geiping, Jonas; Garrido, Quentin; Fernandez, Pierre; Bar, Amir; Pirsiavash, Hamed; LeCun, Yann; Goldblum, Micah: , A Cookbook of Self-Supervised Learning, 2023.

[BF76] Bozinovski, Stevo; Fulgosi, Ante: The influence of pattern similarity and transfer learning upon training of a base perceptron b2. In: Proceedings of Symposium Informatica. volume 3, pp. 121–126, 1976.

[Bl99] Blackmore, Simon: Remedial Correction of Yield Map Data. Precision Agriculture, 1(1):53–66, 1999.

[Bo20] Bosilj, Petra; Aptoula, Erchan; Duckett, Tom; Cielniak, Grzegorz: Transfer learning between crop types for semantic segmentation of crops versus weeds in precision agriculture. Journal of Field Robotics, 37(1):7–19, 2020.

[BPL22] Bardes, Adrien; Ponce, Jean; LeCun, Yann: , VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning, 2022.

[BS22] Bali, Nishu; Singla, Anshu: Emerging Trends in Machine Learning to Predict Crop Yield and Study Its Influential Factors: A Survey. Archives of Computational Methods in Engineering, 29(1):95–112, 2022.

[Bu20] Buslaev, Alexander; Iglovikov, Vladimir I.; Khvedchenya, Eugene; Parinov, Alex; Druzhinin, Mikhail; Kalinin, Alexandr A.: Albumentations: Fast and Flexible Image Augmentations. Information, 11(2):125, 2020. Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.

[BWM20] Bhatt, Umang; Weller, Adrian; Moura, José M. F.: Evaluating and Aggregating Feature-based Model Explanations. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence Organization, pp. 3016–3022, 2020.

[Ch20a] Chen, Junde; Chen, Jinxiu; Zhang, Defu; Sun, Yuandong; Nanehkaran, Y. A.: Using deep transfer learning for image-based plant disease identification. Computers and Electronics in Agriculture, 173:105393, 2020.

[Ch20b] Chen, Ting; Kornblith, Simon; Norouzi, Mohammad; Hinton, Geoffrey: A Simple Framework for Contrastive Learning of Visual Representations. In: Proceedings of the 37th International Conference on Machine Learning. PMLR, pp. 1597–1607, 2020. ISSN: 2640-3498.

[CH20c] Chen, Xinlei; He, Kaiming: , Exploring Simple Siamese Representation Learning, 2020.

[Cr88] Cressie, Noel: Spatial prediction and ordinary kriging. Mathematical Geology, 20(4):405–421, 1988.

[DFM22] Dasgupta, Sanjoy; Frost, Nave; Moshkovitz, Michal: , Framework for Evaluating Faithfulness of Local Explanations, 2022.

[FKH18] Falkner, Stefan; Klein, Aaron; Hutter, Frank: , BOHB: Robust and Efficient Hyperparameter Optimization at Scale, 2018.

[Ge21] Ge, Shuang; Zhang, Jinshui; Pan, Yaozhong; Yang, Zhi; Zhu, Shuang: Transferable deep learning model based on the phenological matching principle for mapping crop extent. International Journal of Applied Earth Observation and Geoinformation, 102:102451, 2021.

[Gh18] Ghosal, Sambuddha; Blystone, David; Singh, Asheesh K.; Ganapathysubramanian, Baskar; Singh, Arti; Sarkar, Soumik: An explainable deep machine vision framework for plant stress phenotyping. Proceedings of the National Academy of Sciences, 115(18):4613–4618, 2018. Publisher: Proceedings of the National Academy of Sciences.

[GN21] Güldenring, Ronja; Nalpantidis, Lazaros: Self-supervised contrastive learning on agricultural images. Computers and Electronics in Agriculture, 191:106510, 2021.

[Gr21] Grahmann, Kathrin; Reckling, Moritz; Hernandez-Ochoa, Ixchel; Ewert, Frank: An agricultural diversification trial by patchy field arrangements at the landscape level: The landscape living lab "patchCROP". p. 7, 2021.

[He16] He, Kaiming; Zhang, Xiangyu; Ren, Shaoqing; Sun, Jian: Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778, 2016. ISSN: 1063-6919.

[He23] Henighan, Tom; Carter, Shan; Hume, Tristan; Elhage, Nelson; Lasenby, Robert; Fort, Stanislav; Schiefer, Nicholas; Olah, Christopher: Superposition, Memorization, and Double Descent. Transformer Circuits, 2023.

[JmMhY16] Jian-min, Liu; Min-hua, Yang; Yu, Xiang: Hierarchical features learning with convolutional neural networks based on aircraft recognition on images from remote sensing image. In: 2016 IEEE International Conference on Consumer Electronics-China (ICCE-China). pp. 1–5, 2016.

[Ka21] Kar, Soumyashree; Nagasubramanian, Koushik; Elango, Dinakaran; Nair, Ajay; Mueller, Daren S.; O'Neal, Matthew E.; Singh, Asheesh K.; Sarkar, Soumik; Ganapathysubramanian, Baskar; Singh, Arti: Self-Supervised Learning Improves Agricultural Pest Classification. In: AI for Agriculture and Food Systems. 2021.

[Ka22] Kattenborn, Teja; Schiefer, Felix; Frey, Julian; Feilhauer, Hannes; Mahecha, Miguel D.; Dormann, Carsten F.: Spatially autocorrelated training and validation samples inflate performance assessment of convolutional neural networks. ISPRS Open Journal of Photogrammetry and Remote Sensing, 5:100018, 2022.

[KPB18] Kamilaris, Andreas; Prenafeta-Boldú, Francesc X.: Deep learning in agriculture: A survey. Computers and Electronics in Agriculture, 147:70–90, 2018.

[KS15] Kuwata, Kentaro; Shibasaki, Ryosuke: Estimating crop yields with deep learning and remotely sensed data. In: 2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). IEEE, pp. 858–861, 2015.

[Ku17] Kussul, Nataliia; Lavreniuk, Mykola; Skakun, Sergii; Shelestov, Andrii: Deep Learning Classification of Land Cover and Crop Types Using Remote Sensing Data. IEEE Geoscience and Remote Sensing Letters, 14(5):778–782, 2017.

[LBO14] Lyle, G.; Bryan, B. A.; Ostendorf, B.: Post-processing methods to eliminate erroneous grain yield measurements: review and directions for future development. Precision Agriculture, 15(4):377–402, 2014.

[Le89] LeCun, Yann; Boser, Bernhard; Denker, John; Henderson, Donnie; Howard, R.; Hubbard, Wayne; Jackel, Lawrence: Handwritten Digit Recognition with a Back-Propagation Network. In: Advances in Neural Information Processing Systems. volume 2. Morgan-Kaufmann, 1989.

[Le14] Le Rest, Kévin; Pinaud, David; Monestiez, Pascal; Chadoeuf, Joël; Bretagnolle, Vincent: Spatial leave-one-out cross-validation for variable selection in the presence of spatial autocorrelation. Global Ecology and Biogeography, 23(7):811–820, 2014.

[Ma23] Mateo-Sanchis, Anna; Adsuara, Jose E.; Piles, Maria; Munoz-Marí, Jordi; Perez-Suay, Adrian; Camps-Valls, Gustau: Interpretable Long Short-Term Memory Networks for Crop Yield Estimation. IEEE Geoscience and Remote Sensing Letters, 20:1–5, 2023. Conference Name: IEEE Geoscience and Remote Sensing Letters.

[Me19] Meyer, Hanna; Reudenbach, Christoph; Wöllauer, Stephan; Nauss, Thomas: Importance of spatial predictor variable selection in machine learning applications – Moving from data reproduction to spatial prediction. Ecological Modelling, 411:108815, 2019.

[Mo21] Molnar, Christoph: Interpretable Machine Learning. 2021.

[Na19] Nagasubramanian, Koushik; Jones, Sarah; Singh, Asheesh K.; Sarkar, Soumik; Singh, Arti; Ganapathysubramanian, Baskar: Plant disease identification using explainable 3D deep learning on hyperspectral images. Plant Methods, 15(1):98, 2019.

[NNL19] Nevavuori, Petteri; Narra, Nathaniel; Lipping, Tarmo: Crop yield prediction with deep convolutional neural networks. Computers and Electronics in Agriculture, 163:104859, 2019.

[No21] Nowakowski, Artur; Mrziglod, John; Spiller, Dario; Bonifacio, Rogerio; Ferrari, Irene; Mathieu, Pierre Philippe; Garcia-Herranz, Manuel; Kim, Do-Hyung: Crop type mapping by using transfer learning. International Journal of Applied Earth Observation and Geoinformation, 98:102313, 2021.

[PDS18] Petsiuk, Vitali; Das, Abir; Saenko, Kate: , RISE: Randomized Input Sampling for Explanation of Black-box Models, 2018.

[Pl20] Ploton, Pierre; Mortier, Frédéric; Réjou-Méchain, Maxime; Barbier, Nicolas; Picard, Nicolas; Rossi, Vivien; Dormann, Carsten; Cornu, Guillaume; Viennois, Gaëlle; Bayol, Nicolas; Lyapustin, Alexei; Gourlet-Fleury, Sylvie; Pélissier, Raphaël: Spatial validation reveals poor predictive performance of large-scale ecological mapping models. Nature Communications, 11(1):4540, 2020. Number: 1 Publisher: Nature Publishing Group.

[Re19] Recht, Benjamin; Roelofs, Rebecca; Schmidt, Ludwig; Shankar, Vaishaal: , Do ImageNet Classifiers Generalize to ImageNet?, 2019.

[Ro17] Roberts, David R.; Bahn, Volker; Ciuti, Simone; Boyce, Mark S.; Elith, Jane; Guillera-Arroita, Gurutzeta; Hauenstein, Severin; Lahoz-Monfort, José J.; Schröder, Boris; Thuiller, Wilfried; Warton, David I.; Wintle, Brendan A.; Hartig, Florian; Dormann, Carsten F.: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. Ecography, 40(8):913–929, 2017.

[Ry22a] Ryo, Masahiro: Explainable artificial intelligence and interpretable machine learning for agricultural data analysis. Artificial Intelligence in Agriculture, 6:257–265, 2022.

[Ry22b] Ryo, Masahiro; Schiller, Josepha; Stiller, Stefan; Rivera Palacio, Juan Camilo; Mengsuwan, Konlavach; Safonova, Anastasiia; Wei, Yuqi: Deep learning for sustainable agriculture needs ecology and human involvement. Journal of Sustainable Agriculture and Environment, 2(1):40–44, 2022.

[Sa23] Safonova, Anastasiia; Ghazaryan, Gohar; Stiller, Stefan; Main-Knorn, Magdalena; Nendel, Claas; Ryo, Masahiro: Small data problems in deep learning applications with remote sensing: A review. 2023. Publisher: EarthArXiv.

[Se20] Selvaraju, Ramprasaath R.; Cogswell, Michael; Das, Abhishek; Vedantam, Ramakrishna; Parikh, Devi; Batra, Dhruv: Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. International Journal of Computer Vision, 128(2):336–359, 2020.

[Sh19] Shen, Runping; Huang, Anqi; Li, Bolun; Guo, Jia: Construction of a drought monitoring model using deep learning based on multi-source remote sensing data. International Journal of Applied Earth Observation and Geoinformation, 79:48–57, 2019.

[SK19] Shorten, Connor; Khoshgoftaar, Taghi M.: A survey on Image Data Augmentation for Deep Learning. Journal of Big Data, 6(1):60, 2019.

[St74] Stone, M.: Cross-Validatory Choice and Assessment of Statistical Predictions. Journal of the Royal Statistical Society: Series B (Methodological), 36(2):111–133, 1974.

[Te20] Tetila, Everton Castelao; Machado, Bruno Brandoli; Astolfi, Gilberto; de Souza Belete, Nicolas Alessandro; Amorim, Willian Paraguassu; Roel, Antonia Railda; Pistori, Hemerson: Detection and classification of soybean pests using deep learning with UAV images. COMPUTERS AND ELECTRONICS IN AGRICULTURE, 179, 2020. Place: THE BOULEVARD, LANGFORD LANE, KIDLINGTON, OXFORD OX5 1GB, OXON, ENGLAND Publisher: ELSEVIER SCI LTD Type: Article.

[To70] Tobler, W. R.: A Computer Movie Simulating Urban Growth in the Detroit Region. Economic Geography, 46:234–240, 1970.

[Va19] Valente, J.; Doldersum, M.; Roers, C.; Kooistra, L.: DETECTING *RUMEX OBTUSIFOLIUS* WEED PLANTS IN GRASSLANDS FROM UAV RGB IMAGERY USING DEEP LEARNING. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, IV-2-W5:179–185, 2019. Conference Name: ISPRS Geospatial Week 2019 (Volume IV-2/W5) - 10–14 June 2019, Enschede, The Netherlands Publisher: Copernicus GmbH.

[WC21] Wainer, Jacques; Cawley, Gavin: Nested cross-validation when selecting classifiers is overzealous for most practical applications. Expert Systems with Applications, 182:115222, 2021.

[WJD20] Weiss, M.; Jacob, F.; Duveiller, G.: Remote sensing for agricultural applications: A meta-review. Remote Sensing of Environment, 236:111402, 2020.

[Ya20] Yang, Xingyi; He, Xuehai; Liang, Yuxiao; Yang, Yue; Zhang, Shanghang; Xie, Pengtao: , Transfer Learning or Self-supervised Learning? A Tale of Two Pretraining Paradigms, 2020.

[Zh22] Zhu, Yilin; Wu, Sensen; Qin, Mengjiao; Fu, Zhiyi; Gao, Yi; Wang, Yuanyuan; Du, Zhenhong: A deep learning crop model for adaptive yield estimation in large areas. International Journal of Applied Earth Observation and Geoinformation, 110:102828, 2022.