

AgroScan*- An Interactive SOC Prediction Tool for Farmers

Margarete Breitenhuber¹, Clara Buller¹, Claudio Bussinger¹, Jerome Reiter¹,
Mennatallah El-Assady², Yannick Metz², Christina Humer²

ETH Zürich, Zurich, Switzerland




Fig. 1: The AgroScan Workflow: Schematic overview of the AgroScan web-tool showing the user interaction workflow. The workflow is enabled by an SOC prediction pipeline based on training data compiled from satellite imagery and the LUCAS soil dataset. By predicting SOC for novel locations and time-points, farmers can analyze values for individual selected fields for any time-point satellite images are available. We support the user workflow with a personalized monitoring and analysis dashboard enabling comparative analysis, as well as temporal trend analysis.

Abstract—Soil organic carbon (SOC) is a key indicator of soil health and one of the largest terrestrial carbon pools. Recent European initiatives, including the EU Soil Monitoring Law, highlight the need for scalable and cost-effective approaches for monitoring agricultural soils. This study investigates how visual analytics can support farmers in the long-term monitoring and management of soil carbon stocks. User requirements were elicited through exploratory interviews with farmers and incorporated into AgroScan, a non-commercial open-source web platform that combines remote-sensing-based SOC prediction with visual analytics. The platform provides field-level SOC estimates and derived soil organic matter (SOM) values together with contextual information such as crop rotation history to facilitate the interpretation of long-term soil health trends. By increasing the transparency and contextualizing the model outputs, AgroScan enables farmers to explore relationships between management practices and soil carbon dynamics and supports adaptive field management decisions.

Index Terms—Remote Sensing, Sustainable Soil Management, Carbon Farming, Machine Learning.

1 INTRODUCTION

Soils represent the largest terrestrial carbon pool globally and play a critical role as carbon sinks through soil carbon sequestration, where atmospheric CO₂ is stored as soil organic matter [1]. Agricultural soils with high Soil Organic Carbon (SOC) content deliver significant environmental and agronomic co-benefits: improved long-term

productivity, increased biodiversity, and preservation of ecosystem services [2]. Consequently, monitoring and enhancing SOC stocks has become a priority within major policy frameworks, including the EU Soil Strategy for 2030 that translates into the EU Soil Monitoring Law [2,3]. Despite this recognized importance, farmers face substantial barriers when integrating SOC monitoring into their agricultural practices. Conventional soil sampling requires repeated field and laboratory work across extended time periods, creating significant cost, labor, and scalability constraints that limit its routine application in large agricultural areas [4]. Machine learning-based remote sensing models have emerged as a promising approach to address these limitations [5,6]. By leveraging satellite-derived data, these models offer scalable and cost-efficient SOC estimation and allow a baseline to be established using historic satellite imagery to assess the impact of Sustainable Soil Management (SSM) practices over time. Existing SOC prediction tools, however, only provide limited transparency regarding their underlying

*In this publication, 'AgroScan' is used strictly as an academic nomenclature to identify the non-commercial ML framework developed by the authors. It does not denote a commercial product or trademark.

¹ These authors contributed equally. Corresponding e-mail: mbreitenhube@ethz.ch.

² Supervisors.

Manuscript received xx xxx. 202x; accepted xx xxx. 202x. Date of Publication xx xxx. 202x; date of current version xx xxx. 202x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.202x.xxxxxx

ing methodology. Many commercial SOC monitoring platforms are designed for carbon credit generation, compliance, and monetization based on SOC changes on field- or farm-level. These solutions provide an accounting system but place little emphasis on model transparency and the development of agronomic understanding for the user. As a result, users are constrained in their ability to critically evaluate model outputs, infer drivers of soil evolution patterns, and make adaptive field management decisions.

This motivated the central research question of the present study: *How can visual analytics empower farmers to monitor and manage soil carbon stocks in agricultural soils over the long term?* AgroScan addresses this question by providing a novel open-source platform (non-commercial web deployment accessible [here](#)) utilizing visual analytics principles as a means of increasing the transparency and interpretability of the model to empower the farmer in their decision-making. Our main contributions are: (1) A visual analytics tool based on user requirements derived from exploratory interviews; (2) End-to-end integration of an ML-based SOC prediction model into an open-source deployed platform; (3) Feedback collection for AgroScan with farmers.

2 RELATED WORK

The use of remote sensing combined with machine learning for SOC prediction in agricultural soils has become more established in academic research. Nguyen et al. [5] demonstrated a multi-sensor data fusion model combining multispectral image data from Sentinel-2 and synthetic aperture radar data from Sentinel-1, achieving high prediction accuracy ($R^2 = 0.870$, Root Mean Square Error (RMSE) = 1.818 t/ha). Similar results supporting the use of remote sensing for carbon prediction were demonstrated by Yuzugullu et al. [6], Beisekenov et al. [7], Zhang et al. [8], Guo et al. [9]. Together, these studies confirm that the technical foundation for remote sensing-based prediction SOC in agricultural contexts is well established.

Based on the academic foundation, several commercial platforms for remote-sensing-based carbon monitoring exist, highlighting the demand for such solutions. Examples are Smart Cloud Farming [10], Spacenus [11], Soyle [12], Agreea [13], and Farmonaut [14]. Common applications include soil sampling location algorithms, variable-rate fertilization, yield productivity monitoring, and measurement, reporting, and verification (MRV) for carbon credit schemes, enabling farmers to optimize inputs and generate additional income. In contrast, AgroScan targets long-term decision support by providing insights into the effects of crop rotations on soil organic carbon (SOC) and humus levels, tasks that have not been addressed in existing work. This empowers the farmer to make long-term sustainable decisions. Through visualizations of the temporal evolution of SOC and derived soil organic matter (SOM) content based on the van Bemmelen factor, AgroScan facilitates the identification of long-term trends related to the crop rotation and farming practices. Furthermore, whereas many existing tools lack transparency regarding their underlying models, AgroScan provides methodological explanations tailored to non-technical users, detailed reports, and information contextualizing the SOC prediction.

Several visual analytics applications have investigated the use of interactive visualizations to support long-term agricultural decision-making. These systems generally aim to enable users to explore temporal developments in agricultural and environmental indicators while relating them to spatial patterns and management interventions. GeoVisage, a web-based agricultural decision support platform co-developed by Nipissing University and agricultural producers in Ontario, Canada, integrated horizon charts and small-multiple visualizations to facilitate the exploration of spatiotemporal datasets [15]. Similarly, Htun et al. [16] proposed a dashboard for identifying water stress and irrigation requirements by combining map-based visualizations with temporal line charts. Their user evaluation, however, revealed high mental demand and effort, highlighting the importance of intuitive workflows and guided analytical processes in agricultural visual analytics systems [16]. Dhaliwal et al. also demonstrated the applicability of line charts for analyzing agricultural time-series data, although their implementation lacked a coordinated dashboard environment [17]. Collectively, these studies indicate the potential of integrating spatial and temporal visual-

izations for analyzing long-term agricultural developments. However, their limitations underscore the need for designing a cohesive interface based on user feedback that effectively aids users in their decision-making process and development of agronomic understanding.

3 USER REQUIREMENTS

AgroScan's target users are farmers who use software-based farm management tools, and are interested in carbon farming and SSM practices. To identify user requirements, we conducted exploratory interviews with two farmers in the municipality Eichstätt in Bavaria, Germany. As expert consultations with farmers acting as informants in their professional capacity, these interviews are exempt from ethics approval under ETH Zurich's guidelines [18]. Oral consent was obtained from both interviewees. The idea of an interactive remote-sensing SOC prediction tool was derived from literature research and pitched to the farmers to understand its relevance in their work. Subsequently, questions about the tool were asked to identify key user requirements.

- J. (28), raised on an organic farm with crops and cattle, the interviewee combines practical farming experience with his current work in agricultural policy as a civil servant in the Bavarian department of agriculture.
- A. (54), is a full-time certified organic farmer managing 145 ha of arable land. He is a professionally trained farmer with formal agricultural education and specializes in organic farming.

A typical use case of a farmer for the AgroScan tool is the routine monitoring of SOC content of their agricultural fields to identify long-term trends in SOC and derived SOM level of their fields in dependence on crop rotation and other targeted SSM practices.

Based on the interviews and insights from related work, the following user requirements were defined to inform the website design:

R1. Tangible decision support

- The implementation and understanding of the SOC analysis process requires minimal technical expertise.
- SOC predictions are presented with contextual information to aid understanding of the results and support the learning process about crop management practices.

R2. Integrated field data management system

- All information relevant to monitoring and assessing crop management practices is summarized and displayed in a comprehensive and clear manner.
- Identification of long-term trends in dependence on SSM practices is facilitated.

R3. Transparent model-user interaction

- Transparent communication of model assumptions, underlying data, and limitations.
- Agency in interacting with and developing the underlying model.

4 SYSTEM DESIGN

The identified user requirements are addressed through several complementary design elements. R1 is operationalized through a structured, step-by-step analysis workflow, as well as spatial and temporal visualizations of SOC predictions within the analysis interface and personal dashboard. R2 is realized through the design of the personalized analytical dashboard, while R3 is supported by transparent communication of the model's limitations, open-source code, publicly available data sources, and accessible documentation. The AgroScan platform can be accessed via agroscan.ivia.ch hosted on an ETH Zurich server.

4.1 Interaction Workflow

Analysis Core User Workflow: The analysis user interface employs a step-by-step navigation model. Several interaction design principles were applied to streamline the user experience and minimize the required technical expertise. Pattern affordance is established by consistently coloring every workflow-advancing button in AgroScan's

signature green to help users easily orient themselves. To prevent misguidance, these action buttons remain greyed out until they become active in the sequence. Established icons and interface conventions (e.g., a magnifier for location search) build on familiarity with widely used services such as Google Maps. To support first-time users, instructions guide users in each step of the analysis, with only the current step visible, with the aim of preventing unintentional errors and user confusion (R1). We explain the limitations of the ML model used to avoid overreliance or misconceptions of its capabilities (R3).



Fig. 2: Screenshot of the personal dashboard of the AgroScan tool, including KPI summary at the top, an overview map, and a field card.

Personal Dashboard: The analytical personal dashboard (R2) implements a hierarchical information architecture (see Figure 2). A map gives an overview of all fields monitored by the farmer. A KPI (key performance indicator) strip provides a high-level aggregation of field information (total fields, analyses, etc). Content cards provide detailed information for each field. These cards can be sorted and filtered, allowing for easy and scalable display of information. Each content card includes field metadata, a visualization of SOC content over time annotated with custom events, and field-specific KPIs. Farmers can view the temporal development of selected fields in an additional line chart to compare them. Details of the field-level analyses can be accessed via the “Click to see all analyses and event” button, alongside a dedicated button to initiate a new analysis for the same field to support long-term SOC monitoring. To support the farmer in identifying the impact of their crop rotation on the SOC level, the planted crop can be added as an event to the line chart. Additionally, the SOC value is converted to the more familiar SOM value via the van Bemmelen factor [19]. Combined with contextualizing reference information such as SOC variation within the field, and numerical historical changes, the farmer is supported in developing an agronomic understanding of the impact of their crop rotation and derive adaptive field management decisions (R1).

4.2 Visual Design and User Guidance

Spatial predictions of soil organic carbon (SOC) are communicated through an interactive map that highlights intra-field variability. This was done to facilitate the identification of low-SOC zones and support targeted interventions (R1). Predicted SOC values are represented using the color-blind-friendly Plasma color palette [20], which is applied consistently throughout the map visualizations to ensure visual coherence. The corresponding color classes are derived from the SOC value distribution of the LUCAS soil dataset and are explained via a legend. Grounding the classification in a well-established reference dataset was done to increase the transparency and interpretability of the displayed results (R1, R3).

In addition to the spatial representation that functions as an overview, temporal developments in SOC are visualized using a line chart. The chart is complemented by contextual information, such as crop rotation practices to support perceptual inference by allowing users to visually associate changes in predicted SOC levels to management decisions and to identify long-term trends. Presenting multiple visualization types aims to facilitate the interpretation of both spatial and temporal patterns and supports informed decision-making (R1, R2).

The technical agency (R3) is further promoted through several mechanisms. To accommodate users with varying levels of technical expertise, a dedicated user guide explains the underlying machine learning methodology in accessible language, while a separate technical summary provides additional details for users seeking a deeper understanding. Furthermore, the publicly available [GitHub repository](#) offers an additional channel for interaction. Farmers without programming experience can submit feature requests or report issues, whereas technically proficient users can inspect the implementation, reproduce results, and contribute to further development.

4.3 Machine Learning Model

Data Selection

Two publicly available datasets were used to train the AgroScan baseline model:

Input data: Sentinel-2 multispectral imagery was selected as the primary input data source. It is part of the European Space Agency’s Copernicus program, and provides near-global coverage with a site revisit frequency of five days. From its 13 spectral bands, which span the visible, near-infrared, and short-wave infrared regions of the electromagnetic spectrum [21], 10 were selected as model inputs (B2-B8, B8A, B11, and B12). These bands enable inference of key biophysical properties relevant to SOC estimation, such as soil moisture and crop growth status [5]. Level-2A products were used, which provide atmospherically corrected Bottom-of-Atmosphere reflectance values.

Target variable: The LUCAS Topsoil dataset from the European Union (EU) was selected as the ground-truth SOC measurements. It provides georeferenced soil samples across EU member states from the topsoil layer (0–20 cm), which is most relevant for SOC dynamics in agricultural systems [22]. The dataset includes land-use classification, enabling filtering to cropland samples to ensure consistency between the target variable and the modeling domain.

Case study region: Germany was selected as the case study region due to its significant agricultural sector, strong policy relevance, and high soil data availability. Germany is one of the EU’s largest agricultural producers [23]. It must comply with the EU Soil Monitoring Law by 2030, creating a demand for scalable and cost-effective soil monitoring methods [3], and the LUCAS 2019 soil dataset has 332 agricultural sampling locations in Germany, of which 270 could be matched to Sentinel-2 data [24].

Data Processing

The data collection and processing pipeline is based on AgroLens [25]. We updated the pipeline to the current Copernicus OData communication protocol [26], added a feature transformation, and modified it to be compatible with our interactive use case and remotely deployable.

Feature transformation: Based on the 10 raw S-2 band values selected, 13 additional predictor features were derived: 7 vegetation

indices, 4 soil indices, and 2 temporal features. This feature engineering step follows the feature selection of [5], who demonstrated that derived indices correlate more strongly with agricultural SOC than raw spectral bands alone. Additionally, a seasonal indicator was encoded using the survey month as sine and cosine components. This captures seasonal variation without an artificial discontinuity between December and January.

Interactive use case: A fixed reference min-max normalization was implemented for all predictor features to make the data processing pipeline fully reusable between training and individual farmer queries, where the size of the datasets differs significantly. Hence, instead of applying dataset-dependent normalization, a fixed approach using minimum and maximum values from a representative dataset – here, the German training set – was defined to normalize the S-2 band values and derived indices. Details can be found in [Appx. A](#).

ML Model Architecture

As the primary objective of this work was the design of an interactive web-based platform for supporting farmers' decision-making, we developed a prototype SOC prediction model rather than adopting a pre-trained model. Although this choice resulted in lower predictive accuracy, it enabled the investigation of use-case-specific interaction requirements and facilitated the assessment of practical limitations.

The model architecture and implementation are based on the FCNN architecture presented by Kammerlander in [27]. For more details about the model architecture, see our [code base](#). Training the model resulted in a RMSE of 6.536 OC g/kg. This is a relatively high error, resulting from limited computational resources. However, optimizing model performance was not the scope of this project and we deem it sufficient for a preliminary model to test our system.

5 USER FEEDBACK COLLECTION

An exploratory user feedback interview for the analysis workflow and dashboard was conducted with A., who was involved in the initial requirement ideation, in a 1.5-hour remote Zoom session. The semi-structured interview used task-based interaction to collect feedback on usability, relevance, and missing functionality and identify next steps for the tool development. During the Zoom session, three team members were present, of whom one primarily conducted the interview while the others took hand-written notes. The notes were digitized and cross-checked by the other team members. The final notes were shared with A., who confirmed their completeness and correctness. The transcribed notes are available as supplementary material ([link](#)). As in 3, this interview is exempted from an ethics approval [18].

General remarks

- A. reaffirmed the relevance of the use case, emphasizing that soil health and carbon management are increasingly important topics in both agricultural practice and policy discussions. The participant highlighted the importance of involving practitioners throughout the development process to ensure practical relevance and noted that the tool could encourage the adoption of more sustainable farming practices.
- Reliable predictions and sufficient model validation were identified as critical requirements. The participant expressed skepticism regarding the suitability of remote sensing for SOC estimation. AgroScan addresses these concerns by providing methodological transparency, and communicating model limitations.
- Existing reporting obligations already impose a significant administrative burden on farmers. Integration with existing farm management and reporting systems was identified as an important future requirement to minimize additional data management effort.

Analysis workflow

- The entry point through the home page and the overall analysis workflow were perceived positively.
- Field selection and modification using the polygon tool were found to be unintuitive. The participant suggested interaction patterns similar to existing farm management software, where fields are selected directly via a double-click and modified only if required.

- The initial download of satellite imagery required approximately five minutes. This delay was considered acceptable provided that the expected waiting time was communicated transparently. Repeated analyses benefit from the implemented data cache.
- Saving analyses and providing additional contextual information were considered straightforward. The participant emphasized the importance of displaying the acquisition date of the most recent satellite image to support correct crop assignment.

Dashboard

- The overview map and the 60 m × 60 m grid representation were considered useful for identifying intra-field variability.
- The conversion of SOC to soil organic matter using the van Bemmelen factor was perceived as beneficial, as soil organic matter is the metric most commonly used by farmers and policymakers.
- Aggregate KPI summaries received little attention. Instead, the participant expressed greater interest in farm-level and crop-specific balances to understand the long-term development of soil organic matter and the influence of individual crops.
- The primary use case of the dashboard was identified as the analysis of long-term trends at field and farm level. The temporal line charts and event annotations were perceived as supporting this objective.
- Comparison between fields was considered useful for investigating the effects of management practices, although less important than monitoring temporal trends.
- Additional agronomic indicators, such as nitrogen and phosphorus levels, were identified as potentially increasing the practical value of the platform.

6 FUTURE WORK AND LIMITATIONS

This work demonstrates how visual analytics can translate SOC predictions into actionable insights for farmers by supporting long-term monitoring and management of soil health. By combining an intuitive analytical workflow, contextual visualizations, and an accessible, transparent system design, the proposed approach enables informed decision-making while conferring technical agency to the user.

However, several limitations remain. First, the interviewed farmers shared similar backgrounds and farming contexts, which may have introduced bias into the derivation of the user requirements and limits the generalizability of the findings. Second, while the user interviews informed the next iteration of the tool development, the proposed design still requires validation with a larger and more representative group of users. Finally, the underlying machine learning model remains a prototype with limited predictive accuracy and lacks a quantification of the prediction uncertainty, which constrains the system's current applicability in real-world agricultural decision-making.

In the future, based on user feedback, we want to increase real-world applicability by improving the SOC prediction model by extending the selection of predictor features and expanding the training data. We want to integrate additional prediction targets and uncertainty quantification. Additionally, we want to increase the spatial resolution (currently 60×60 m) to support the analysis of within-field variability. Furthermore, we want to integrate the geospatial context and the characteristics of the soil, allowing definitions of thresholds relevant to the region and the management. A prognosis model to identify crops with positive and negative effects on soil organic matter can help estimate the long-term impact of a farm's crop rotation on SOC and SOM levels. Finally, we want to integrate AgroScan with existing farm reporting systems, such as government-mandated fertilization reporting, to automate data retrieval and minimize data management efforts for farmers.

AI STATEMENT

Generative AI tools were used to support the implementation of AgroScan, primarily by assisting with coding and technical development. The project team remains fully responsible for the implementation and all content presented in the report. All AI-assisted outputs were carefully reviewed, validated, and, where necessary, revised by the authors.

ACKNOWLEDGMENTS

AgroScan was developed as part of the course “Human-Centered AI for Social Good” at ETH Zürich by the Interactive Visualization and Automation (IVIA) Laboratory. Yannick Metz is supported by the ETH World Food System Center project ‘Artificial Intelligence for Sustainable Soil Microbiomes’ with funding from fenaco and the ETH Zurich foundation. This work was supported by the ETH AI Center postdoctoral fellowship awarded to Christina Humer.

REFERENCES

- [1] K. Paustian, E. Larson, J. Kent, E. Marx, and A. Swan. Soil C sequestration as a biological negative emission strategy. *Frontiers in Climate*, 1:8, 2019. doi:10.3389/fclim.2019.00008. 1
- [2] European Commission. EU Soil Strategy for 2030: Reaping the Benefits of Healthy Soils for People, Food, Nature and Climate. Technical Report COM(2021) 699 final, European Commission, 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021DC0699>. 1
- [3] European Parliament and Council of the European Union. Directive (eu) 2025/2360 on soil monitoring and resilience (soil monitoring law). Official Journal of the European Union, L 2025/2360, 2025. URL: <https://eur-lex.europa.eu/eli/dir/2025/2360/oj>. 1, 3
- [4] A. Ihasusta, A. Al Bitar, N. H. Batjes, F. van Egmond, R. Cardinael, S. Karunaratne, and E. Ceschia. Choosing the appropriate methodology to monitor soil organic carbon (soc) in croplands: Aligning methods with evolving monitoring reporting verification (mrsv) frameworks. *Carbon Management*, 17(1), 2026. doi:10.1080/17583004.2026.2638317. 1
- [5] T. T. Nguyen, T. D. Pham, C. T. Nguyen, J. Delfos, R. Archibald, K. B. Dang, N. B. Hoang, W. Guo, and H. H. Ngo. A novel intelligence approach based active and ensemble learning for agricultural soil organic carbon prediction using multispectral and sar data fusion. *Science of the Total Environment*, 804:150187, 2022. doi:10.1016/j.scitotenv.2021.150187. 1, 2, 3, 4
- [6] O. Yuzugullu, N. Fajraoui, A. Don, and F. Liebisch. Satellite-based soil organic carbon mapping on european soils using available datasets and support sampling. *Science of Remote Sensing*, 9:100118, 2024. doi:10.1016/j.srs.2024.100118. 1, 2
- [7] N. Beisekenov, W. Banakinaou, A. D. Ajayi, H. Hasegawa, and A. Tadao. Remote sensing-based soil organic carbon monitoring using advanced machine learning techniques under conservation agriculture systems. *Smart Agricultural Technology*, 11:101036, 2025. doi:10.1016/j.atech.2025.101036. 2
- [8] Y. Zhang, L. Guo, Y. Chen, T. Shi, M. Luo, Q. Ju, H. Zhang, and S. Wang. Prediction of soil organic carbon based on landsat 8 monthly ndvi data for the jiangnan plain in hubei province, china. *Remote Sensing*, 11(14):1683, 2019. doi:10.3390/rs11141683. 2
- [9] L. Guo, P. Fu, T. Shi, Y. Chen, C. Zeng, H. Zhang, and S. Wang. Exploring influence factors in mapping soil organic carbon on low-relief agricultural lands using time series of remote sensing data. *Soil and Tillage Research*, 210:104982, 2021. doi:10.1016/j.still.2021.104982. 2
- [10] Smart Cloud Farming. Smart cloud farming – about. <https://smartcloudfarming.com/about>. Accessed: 2026-06-23. 2
- [11] Spacenus. Spacenus – nitrogen rate recommendations (precision farming). <https://www.spacenus.com/#nitrogen-rate-recommendations>. Accessed: 2026-06-23. 2
- [12] Soyle. Soyle – how it works. https://www.soyle.io/?utm_source=chatgpt.com#how. Accessed: 2026-06-23. 2
- [13] Agreeena. Agreeena – regenerative agriculture platform. <https://agreeena.com>. Accessed: 2026-06-23. 2
- [14] Farmonaut. Carbon footprinting. URL: <https://farmonaut.com/carbon-footprinting>. 2
- [15] Mark P. Wachowiak, Daniel F. Walters, John M. Kovacs, Renata Wachowiak-Smolíková, and April L. James. Visual analytics and remote sensing imagery to support community-based research for precision agriculture in emerging areas. *Computers and Electronics in Agriculture*, 143:149–164, 2017. doi:10.1016/j.compag.2017.09.035. 2
- [16] Nyi-Nyi Htun, Diego Rojo, Jeroen Ooge, Robin De Croon, Aikaterini Kasimati, and Katrien Verbert. Developing visual-assisted decision support systems across diverse agricultural use cases. *Agriculture*, 12(7):1027, 2022. doi:10.3390/agriculture12071027. 2
- [17] J. K. Dhaliwal, M. E. Galbraith, C. K. Leung, and D. Tan. A data discovery and visualization tool for visual analytics of time series in digital agriculture. In *2023 27th International Conference Information Visualisation (IV)*, pages 268–271. IEEE, 2023. doi:10.1109/IV60283.2023.00053. 2
- [18] ETH Zurich Ethics Commission. Responsibility of the ethics commission. <https://ethz.ch/en/research/ethics/ethics-commission/responsibility.html>. Accessed: 2026-08-05. 2, 4
- [19] Food and Agriculture Organization of the United Nations (FAO). Fao knowledge repository, n.d. Accessed: 2026-06-24. URL: <https://openknowledge.fao.org/server/api/core/bitstreams/1b2e9dac-65e6-426e-8f7f-a111e6d305e5/content>. 3
- [20] Garnier, Simon, Ross, Noam, Rudis, Robert, Camargo, Antônio Pedro, Sciaini, Marco, Scherer, and Cédric. *viridis(Lite) - Colorblind-Friendly Color Maps for R*, 2024. viridis package version 0.6.5. URL: <https://sjmgarnier.github.io/viridis/>, doi:10.5281/zenodo.4679423. 3
- [21] Copernicus Sentinel. Sentinel-2. SentiWiki, n.d. Accessed: 2026-05-21. URL: <https://sentiwiki.copernicus.eu/web/sentinel-2>. 3
- [22] European Commission, Joint Research Centre. Lucas 2018 topsoil data. European Soil Data Centre (ESDAC), 2021. URL: <https://esdac.jrc.ec.europa.eu/content/lucas-2018-topsoil-data>. 3
- [23] Federal Ministry of Food and Agriculture. Understanding farming in germany: Facts and figures about german farming, 2023. URL: <https://www.bmleh.de/SharedDocs/Downloads/EN/Publications/UnderstandingFarming.pdf>. 3
- [24] O. Fernández-Ugalde, S. Scarpa, A. Orgiazzi, P. Panagos, M. Van Liedekerke, A. Marechal, and A. Jones. Lucas 2018 soil module: Presentation of dataset and results. Technical Report EUR 31144 EN, Publications Office of the European Union, 2022. doi:10.2760/215013. 3
- [25] CVIMS. Agrolens. GitHub Repository, 2025. URL: <https://github.com/cvims/AgroLens>. 3
- [26] Copernicus Data Space Ecosystem. Odata documentation, 2026. URL: <https://documentation.dataspace.copernicus.eu/APIs/OData.html>. 3
- [27] C. Kammerlander, V. Kolb, M. Luegmair, L. Scheermann, M. Schmailzl, M. Seufert, J. Zhang, D. Dalic, and T. Schön. Machine learning models for soil parameter prediction based on satellite, weather, clay and yield data. *arXiv*, 2025. URL: <https://arxiv.org/abs/2503.22276>, arXiv: 2503.22276. 4

A FIXED REFERENCE NORMALIZATION

The fixed reference minimum and maximum values are computed once from a representative dataset using the Python scripts `calculate_band_ranges.py` and `calculate_index_ranges.py` and saved to a respective `.json` file.

For the normalization, each raw value x is transformed to:

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}}$$

and then clipped to the unit interval if necessary:

$$x' \leftarrow \min(1, \max(0, x'))$$

If $x_{\max} = x_{\min}$, the normalization is undefined and the implementation returns 0 to avoid division by zero. NaN and infinite values are replaced with 0 afterwards. The fixed reference approach ensures that the normalization applied to the farmer’s input is identical to that used during training, preserving the statistical relationship between features that the model learned.