



Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Boosting crop classification by hierarchically fusing satellite, rotational, and contextual data

Valentin Barriere ^{a,b,c,*}, Martin Claverie ^{c,1}, Maja Schneider ^d, Guido Lemoine ^c,
Raphaël d'Andrimont ^c

^a Departamento de Ciencias de Computacion, Universidad de Chile, Santiago, Chile

^b Centro Nacional de Inteligencia Artificial (CENIA), Santiago, Chile

^c European Commission, Joint Research Centre (JRC), Ispra, Italy

^d Technical University of Munich, Munich, Germany

ARTICLE INFO

Edited by Marie Weiss

Keywords:

Agriculture
Crop type mapping
Earth Observation
Geospatial Application
Long-Short-Term-Memory
Sentinel-2

ABSTRACT

Accurate early-season crop type classification is crucial for the crop production estimation and monitoring of agricultural parcels. However, the complexity of the plant growth patterns and their spatio-temporal variability present significant challenges. While current deep learning-based methods show promise in crop type classification from single- and multi-modal time series, most existing methods rely on a single modality, such as satellite optical remote sensing data or crop rotation patterns. We propose a novel approach to fuse multimodal information into a model for improved accuracy and robustness across multiple crop seasons and countries. The approach relies on three modalities used: remote sensing time series from Sentinel-2 and Landsat 8 observations, parcel crop rotation and local crop distribution. To evaluate our approach, we release a new annotated dataset of 7.4 million agricultural parcels in France (FR) and the Netherlands (NL). We associate each parcel with time-series of surface reflectance (Red and NIR) and biophysical variables (LAI, FAPAR). Additionally, we propose a new approach to automatically aggregate crop types into a hierarchical class structure for meaningful model evaluation and a novel data-augmentation technique for early-season classification. Performance of the multimodal approach was assessed at different aggregation levels in the semantic domain, yielding to various ranges of the number of classes spanning from 151 to 8 crop types or groups. It resulted in accuracy ranging from 91% to 95% for the NL dataset and from 85% to 89% for the FR dataset. Pre-training on a dataset improves transferability between countries, allowing for cross-domain and label prediction, and robustness of the performances in a few-shot setting from FR to NL, i.e., when the domain changes as per with significantly new labels. Our proposed approach outperforms comparable methods by enabling deep learning methods to use the often overlooked spatio-temporal context of parcels, resulting in increased precision and generalization capacity.

1. Introduction

Crop-type maps are an essential element used in crop production monitoring that feed into global food security assessments (Porter et al., 2014). Satellite Earth Observation (EO) systems offer a valuable data source for crop classification due to the synoptic, repeated, consistent, and timely availability of observations (Weiss et al., 2020). Since 2015, the data from the European Union (EU)'s Copernicus program, in particular those of the Sentinel-1 (S1) and S2 sensors, provide systematic and consistent EO data at a spatial resolution higher than the size of

most agricultural parcels. Meanwhile, and in agreement with the EU's Common Agricultural Policy (CAP), member states are committed to manage geospatial farmer declaration referred as GSA datasets to serve as support tools for paying agencies in charge of administrating the payments of farmers' subsidies of the CAP. GSA are datasets, delivered annually, and consist of geographic polygons representing agricultural parcels; each polygon is annotated with the cultivated crop type. When made publicly available (Schneider et al., 2023), these datasets can be used as a valuable source of extensive Ground Truth (GT) data for

* Corresponding author at: Departamento de Ciencias de Computacion, Universidad de Chile, Santiago, Chile.

E-mail address: vbarriere@dcc.uchile.cl (V. Barriere).

¹ Shared first authorship.

<https://doi.org/10.1016/j.rse.2024.114110>

Received 19 May 2023; Received in revised form 13 February 2024; Accepted 8 March 2024

Available online 14 March 2024

0034-4257/© 2024 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

List of Acronyms

ANN	Artificial Neural Network
BDAP	Big Data Analytics Platform
biLSTM	bidirectional LSTM
BoC	Bag-of-Crops
BRDF	Bidirectional Reflection Distribution Function
BV	Biophysical variables
CAP	Common Agricultural Policy
CD	Crop Distribution
CDL	Cropland Data Layer
CNN	Convolutional Neural Networks
CNN-LSTM	Convolutional Neural Network-Long-Short-Term-Memory
CR	Crop Rotation
EO	Earth Observation
EU	European Union
FAPAR	Fraction of Absorbed Photosynthetically Active Radiation
FOI	Feature Of Interest
FR	France
GSA	Geospatial Aid Application
HLS	Harmonized Landsat Sentinel-2
InterYE	Inter-Year Encoder
IntraYE	Intra-Year Encoder
L8	Landsat 8
LAI	Leaf Area Index
LSTM	Long-Short-Term-Memory
m-F1	micro-F1 score
NIR	Near Infrared bands
NL	Netherlands
PSE-LTAE	Pixel-Set Encoder with a Lightweight Temporal Attention Encoder
QA	Quality Assessment
RNN	Recurrent Neural Network
RS	Remote Sensing
S1	Sentinel-1
S2	Sentinel-2

mapping crop types. Capitalizing on the abundance of data, Satellite Image Time Series (SITS) datasets have evolved significantly over the past decade (Selea, 2023), paving the way for the adoption of data-intensive methodologies. Unlike the earlier practice of relying on GT from in situ data collections (thus limited in size), the availability of extensive datasets has facilitated the development of deep learning models.

This study introduces a novel multi-modal modeling approach that integrates EO data with historical cropping practices through the use of Crop Rotation (CR). This innovative scheme has the capability to predict crop types in the absence of GT data for the test year. While previous attempts have been made, none have endeavored to integrate these diverse modalities in a learning model. Existing methodologies primarily concentrate on assessing intra-year dynamics, often relying solely on EO data from the specific season of interest, neglecting the essential inter-year dynamics that play a crucial role in reflecting agricultural practices. By fusing the different data sources in a hierarchical way, the approach described in this paper allows one to take advantage of both the inter-year and the intra-year dynamics. To this purpose, we release an aligned multiyear SITS dataset S2-based at parcel level on the basis of the Dutch and the French GSA.

1.1. Related works

In this section, we will explore previous contributions from studies focused on crop type mapping using deep learning. Such methods have demonstrated superior results compared to classical machine learning and signal processing-based approaches, particularly on a large scale.

1.1.1. EO-based models

Over the past five years, several research teams have focused on developing EO-based crop type classifiers relying on deep learning. Taking California as study area, Zhong et al. (2019) investigated the viability of employing 1D Convolutional Artificial Neural Network (ANN) (Conv1D) and LSTM for crop type mapping. They showcased the superiority of the Conv1D algorithm compared to conventional machine learning algorithms (e.g., XGBoost, Random Forest, and Support Vector Machine). Although the LSTM did not perform as well as anticipated, it is possible that the utilization of relatively low-density time series data from Landsat may have had an impact. In Europe, Pelletier et al. (2019) was among the pioneers in developing EO-based crop type classifiers relying on deep learning. They introduced the 1D-Convolutional Neural Networks (CNN) TempCNN algorithm using the Formosat -2 data, which offers high spatial and temporal resolution time series (similar to S2) but with a simplified spectral resolution. When comparing the TempCNN algorithm, they demonstrated a slight improvement compared to random forest. Taking advantage of the extensive collection of Copernicus Sentinel data and GSA datasets, numerous studies have emerged, employing data-intensive learning methods to achieve advanced crop type classification. In the study by Rußwurm et al. (2019a), a parcel-level classification was conducted targeting 13 crop types in French Brittany using the full spectral information of S2. They assessed the performances of two models: a Transformer-Encoder (Vaswani et al., 2017), and a Recurrent Neural Network (RNN) based on LSTM architecture (Hochreiter and Schmidhuber, 1997). Their findings indicated that both the Transformer-Encoder and LSTM models exhibited similar classification performance.

Later on, Rußwurm and Körner (2020) designed a crop classifier at the parcel-level using S2 data from three regions of Germany and compare different approaches to model the signal, including among others a Transformer and an LSTM. They conclude that EO data processing (atmospheric correction, cloud filtering) provided small but consistent benefits for these types of models. A similar approach was taken by Rußwurm et al. (2019b) on 40k parcels in Central Europe, for which they proposed a new early classification mechanism to enhance a classical model with an additional stopping probability based on previously seen information. Furthermore, Rußwurm and Körner (2018) tackled the task of crop classification at the pixel level, by accounting for the spatial variation to detect parcels boundaries, using Convolutional Neural Network-Long-Short-Term-Memory (CNN-LSTM) to classify 17 types of crops in a unique German region. (Sainte Fare Garnot et al., 2019) proposed to use a CNN before a RNN to learn the aggregation of the parcel pixels instead of classically averaging them, and applied their system on 200k parcels of the south-west of France. Finally, Sainte Fare Garnot et al. (2020) proposed a innovative method to tackle parcel-level crop classification, by randomly sampling pixels of the parcels to learn expressive descriptors that are processed by a transformer.

1.1.2. Transferability and domain adaptation containing unseen labels

While most crop type classification studies have relied on current-year GT data, the exploration of models that exhibit spatio-temporal transferability is essential for certain applications. This can be seen as learning from a source data distribution and testing on a target data coming from a different domains, hence called domain adaptation.

Transferability can be extended not only across time but also across space. In this scenario, a model is trained using a dataset from one

location (e.g., country A) and assessed on a separate dataset from a different location (e.g., country B). The task is complex as dataset varies; for instance, new labels can occur in the test set and not in the train set.² The model can be therefore fine tuned with few examples from a new dataset, which is named “few-shot” model. We use “zero-shot” model naming when no examples are used for fine tuning (Peng et al., 2018). Only some works attempt few-shot classification with EO-data knowing that a majority of the systems work poorly without domain data. Nevertheless, Rußwurm et al. (2019c) and Tseng et al. (2021a) have both proposed to use Model-Agnostic Meta-Learning (MAML) in order to tackle few-shot crop or land cover classification at the pixel-level, where “the generalization is induced by a few labeled examples in the target domain”, using EO data only.

(Jia et al., 2019) proposed a method to map data in the target domain to a distribution similar to that of the source domain using a cyclic GAN to generate artificial target domain data. They evaluated their approach on cropland mapping and burned area detection, utilizing only two classes from the Cropland Data Layer (CDL) dataset for cropland mapping, without incorporating unseen labels. Wang et al. (2021), in a four-class classification scenario also without unseen labels, aimed to minimize the Mean Minimum Discrepancy between several hidden states of the network across different domains, forcing adaptation. (Nyborg et al., 2022b) introduced a clever method to align different domains by learning a temporal shift that models the temporal variations in phenology time series for distinct regions. They focused on 15 classes, leaving others unknown, and evaluated their approach on the target domain using classes from the source domain, thereby mitigating the issue of unseen labels. Finally, Lucas et al. (2023) propose a model for few-shot domain adaptation, without a significant part of new labels in the test set (maximum 0.57% of unknown test labels). Thus, these advancements in the realm of Earth Observation data demonstrate the evolving potential of few-shot and zero-shot learning models in overcoming domain-specific challenges and enhancing the adaptability of models to new, unlabeled datasets.

1.1.3. Early-season classification

Although end-of-season crop type maps are essential for agricultural land monitoring (Weiss et al., 2020), there are applications, such as early-season crop production monitoring, that demand a more rapid response, necessitating the release of crop type map before the harvest (named as early-season classification).

The early-season classification necessitates navigating a trade-off between the timeliness and the accuracy of the forecast. Rußwurm et al. (2019b) proposed to solve the problem in an elegant way, with an adapted cost function that only rewards the classifier for an early classification if the right class has been predicted with a respectable degree of accuracy. They extend this work in Rußwurm et al. (2023) by presenting end-to-end Learned early-season classification of Time Series, also classifying crops at the parcel-level in France, Germany, Ghana and South Sudan. Finally, Lin et al. (2022) proposed an original topology-based approach to automatically label instances of the test season in early-season.

Without using a special cost function, Weilandt et al. (2023) use a Pixel-Set Encoder with a Lightweight Temporal Attention Encoder (PSE-LTAE) with a data-augmentation technique initially proposed by Barriere and Claverie (2022) on hierarchical LSTM for crop-type classification at the parcel-level. They perform data-augmentation by randomly cropping the end of the EO time series during training. The data-augmentation technique boosts the performances of early-season classification. They also compared separate models trained on data cropped up to a unique certain period in the year (i.e. one model for one period), which is not efficient in terms of computation and yielded similar results.

² contrary to few-shot domain adaptation, because of new labels in the test set.

1.1.4. Crop-rotation-based models

Crop rotation is a common agronomic practice that provides valuable information for predicting the crop type in following-season. The exploration of Crop Rotation (CR) has been made easier by the availability of year-dependent GSA datasets. A good understanding and design of crop rotation is also vital for sustainability and mitigating the variability of agricultural productivity induced by climate change (Bohan et al., 2021). Therefore, the importance of CR modeling is evident, as it has the potential to enhance mapping performance, particularly for early-season estimates.

Crop rotation patterns are complex and non-stable in time, often dependent on farmer management decisions and subject to changes due to economic considerations and administrative regulations (Dogliotti et al., 2003). As a result, expert knowledge-based models have limitations in terms of accuracy and applicability over large areas and long periods. Alternative approaches, such as estimation of crop sequence probabilities using survey data and hidden Markov models have been demonstrated in France (Xiao et al., 2014), but these methods are not always feasible at large scale due to the extended size of the required sample.

Past research has focused on using machine learning techniques to predict crop rotations. In Osman et al. (2015), a Markov Logic model is used to predict the following season’s crop in France, achieving an accuracy of 60%. Other studies have utilized deep neural networks, such as (Yaramasu et al., 2020), which reaches a maximum accuracy of 88% on a 6-class portion of the US CDL dataset over 12 years.

Only three studies (Johnson and Mueller, 2021; Giordano et al., 2020; Quinton and Landrieu, 2021) have been identified that combine the use of crop rotations and satellite time-series data with deep learning. Johnson and Mueller (2021) applied this method over several seasons to derive near real-time CDL. However, this methodology is constrained to a small number of crop types and the use of a Random Forest classifier, while recent advancements in deep learning have shown significant improvements in such high-data regime problems. Giordano et al. (2020) used Conditional Random Fields to model the temporal dynamics of crop rotations. They focused on two French regions with very different climate conditions and agricultural practices, using around 9,230 and 1,902 parcels with 2 seasons of data. (Quinton and Landrieu, 2021) propose to use a PSE-LTAE (Sainte Fare Garnot et al., 2020) combined with a multi-year classification method. They represent the past crops with a one-hot encoder that they sum, without modeling the dynamics of the sequence. In our work, we not only focus on modeling the sequential aspects of crop rotations, but also incorporate the Remote Sensing (RS) signals from previous seasons.

1.1.5. Class hierarchy

As far as the authors know, the only works using Deep Learning that focus on leveraging the hierarchy between the classes in order to produce significant metrics are the ones of Sainte Fare Garnot and Landrieu (2021) and Turkoglu et al. (2021). The first one propose to use the Average Hierarchical Cost (AHC), which depends on the distance between the different classes in the hierarchical ontology. Nevertheless, they are not taking into account the distribution of the actual test set, in order to create a set of sub-classes that are pertinent for the validation of the model’s performances (i.e. equivalent of setting the distance between some classes to zero when their group is not significant). The second one, in the same way as (Sanh et al., 2018), proposes to integrate the hierarchical structure of the labels by integrating several loss functions at different depth regarding the complexity of the task.

1.2. Positioning and objectives

To the best of the authors’ knowledge, there are some gaps in the existing literature. A significant amount of research has focused on using RS to predict crop types at the pixel or parcel level using only EO and in-situ observations of the current season, treating the

signal as independent from one season to another. Other studies have used parcel crop rotations to address pre-season prediction of crop types (Osman et al., 2015; Yamasu et al., 2020), but the lack of sufficient information in the signal (i.e. short duration of the time series) limits their performance even when targeting minor classes. Although the integration of EO data with CR has been investigated in certain studies (Giordano et al., 2020; Quinton and Landrieu, 2021), no one has yet taken this analysis one step further by incorporating its dynamic modeling. The contributions of this study are 5-fold:

- (i) We release a new aligned SITS dataset of more than 7.4 million parcels with their associated crops, and S2 signals for the period 2016–2020 in France (FR) and NL. This allows for the integration of CR patterns with the EO signal.
- (ii) We create a data- and knowledge-driven technique to automatically group crops together in a meaningful way according to their similarity and importance in the region, for a fairest comparison of classifiers in any crop datasets. It leverages expert knowledge from the EuroCrops (Schneider et al., 2023) taxonomy.
- (iii) We construct a novel approach for crop type mapping from CR and S2 optical time series in a multimodal way using a hierarchical LSTM network. This approach is unique in its conception, as it fuses large amounts of temporally fine-grained EO data with crop rotation analysis in an advanced deep learning method. The crop rotations and the S2 time series are enhanced by previous-season crop distributions of the neighboring parcels.³
- (iv) We develop a data-augmentation technique for the early-season classification, by randomly cropping the end of the RS time-series data. This allows our model to classify parcels before the end of the season, a crucial feature for real-life application of crop monitoring.
- (v) We assess the cross-domain generalization potential of the framework based on a modified nomenclature of EuroCrops, without using any strategy to mitigate domain gaps, target shifts, or handle new classes.

2. Materials

This section presents a description of the study area and the EO data processing procedure.

2.1. Crop reference data, study area, and harmonization of parcel data

The GSA corresponds to the annual crop declarations made by EU farmers for CAP area-aid support measures. The electronic GSA records include a spatial delineation of the parcels. A GSA element is always a polygon of an agricultural parcel with one crop (or a single crop group with the same payment eligibility). The GSA is operated at the region or country level in the EU 27 member states, resulting in about 60 different designs and implementation schemes over the EU. Since these infrastructures are set up in each region, data are not interoperable at the moment, and the legends are not semantically harmonized. Furthermore, only few EU member states release GSA data as open data, although the overall trend is towards increasingly opening up the data for public use.

Some efforts have been made to provide harmonized GSA dataset over the EU. AI4boundaries (d'Andrimont et al., 2023) provides harmonized parcel geometries over 7 countries in the EU to benchmark method for parcel delineation. EuroCrops (Schneider et al., 2021) proposed a semantic harmonization framework to harmonize the legend of GSA across different countries. This harmonization is open source and

is maintained by the community.⁴ While EuroCrops provides a unique effort so far, this work is still in progress especially regarding the time dimension. A recent European Commission Implementing Regulation (EU) 2023/138⁵ identifies a list of specific high-value datasets and the arrangements for their publication. This should be a game changer in the opening of the GSA for public access in the future and thus foster their use for research.

The Hierarchical Crop and Agriculture Taxonomy version 2 (HCATv2) from EuroCrops offers a knowledge graph regrouping crops together in a hierarchical way that is coherent with agricultural practices. It contains 393 classes, which are defined at six hierarchical levels of which the first two are fixed due to compatibility with other taxonomies. For example 33-01-01-05-01 corresponds to the class *Summer Oats*, which is included in its parent class 33-01-01-05-00 (*Oats*) and its grand-parent class 33-01-01-00-00 (*Cereals*). Nevertheless, it is not possible to compare the labels only using the hierarchy because some branches go to a deeper level than others. For example, the class *Capsicum* is level-4 and represent 0.004% of FR, which is the same level than the class *Cereal* representing 32%. HCATv2 (Schneider et al., 2021) was used to represent a Sankey diagram linking the French GSA (left) and the Dutch GSA (right), using HCATv2 (centre) is represented in Fig. A.1 using 40 main crop types for each country.⁶

For this study, FR and the NL were selected because of (i) their open parcel data availability, (ii) their EU representativeness in covering a latitude range from 40° to 55° Northern latitude as well as four biogeographic regions (i.e. Oceanic, Continental, Alpine and Mediterranean) and (iii) the countries have different size and landscape. Parcel GSA data from seasons⁷ 2015 to 2020 over FR and from 2013 to 2020 over the NL were collected 1.

2.2. Geometric minimum common parcel extraction through time

From season-to-season, the parcel boundaries, as delivered by the yearly GSA dataset, may change. We intersected GSA data (i.e. 2013–2020 for NL and 2015–2020 for FR) in order to extract minimum common area, each with a distinct multi-annual crop sequence, named hereafter Feature Of Interest (FOI). Since FOI are the cross-section of varying parcel bounds, their overall size is smaller than the annual GSA parcels. We discarded any FOI with an area of less than 0.1 ha and 0.5 ha for NL and FR, respectively. The total FOI area cover 85% and 93% of the average GSA area for NL and FR respectively (the “stack” entries in Table 1). For each FOI, a crop type sequence was extracted, as well as the remote sensing time series 1 (see Fig. 1).

2.3. Earth observation processing

Remote sensing data were extracted from S2 MSI and L8 OLI sensors. While the GSA data spans from 2013 and 2015, for NL (It starts from 2009 but we only took data from 2013) and FR respectively, the remote sensing data were used starting 2016 cropping season (i.e., from October-2015), corresponding to the first cropping season with Sentinel-2 A in-orbit.

S2 MSI products with Level-2 A surface reflectance data were downloaded from the Copernicus Open Access Hub. L8 OLI surface reflectance data were downloaded from the L30 products of the Harmonized Landsat Sentinel-2 (HLS) data set. For both products, L2 A and L30 Quality Assessment (QA) layers were used to mask

⁴ <https://github.com/maja601/EuroCrops>

⁵ https://eur-lex.europa.eu/eli/reg_impl/2023/138

⁶ An interactive version of the diagram without class limitation is available on https://jeodpp.jrc.ec.europa.eu/ftp/jrc-opendata/DRLL/CropDeepTrans/data/sankey_All_crops.html.

⁷ season n means from October n to October n+1.

³ Our method is completely feature-independent and could be used with other bands.

Table 1
Original Geospatial Aid Application (GSA) parcel numbers and area per season used in the study for France (FR) and Netherlands (NL). The number of distinct crop types are provided using original GSA (4th column) and harmonized using EuroCrops (5th column). The “stack” lines correspond to the Feature Of Interest (FOI, see Section 2.2).

Country	Season	RS	# distinct crop types		Number of polygons	Total area (1000 ha)
			original	harmonized		
NL	2013	✗	76	41	762,725	1855
NL	2014	✗	75	41	765,006	1859
NL	2015	✗	260	117	790,930	1873
NL	2016	✓	296	133	786,572	1871
NL	2017	✓	300	136	785,710	1882
NL	2018	✓	312	135	774,822	1871
NL	2019	✓	317	139	772,565	1868
NL	2020	✓	326	141	767,034	1872
NL	stack		401	148	596,762	1407
FR	2015	✗	261	150	9,434,672	27,856
FR	2016	✓	261	147	9,334,043	27,876
FR	2017	✓	280	148	9,393,747	27,889
FR	2018	✓	282	149	9,517,878	27,917
FR	2019	✓	241	149	9,604,463	27,960
FR	2020	✓	239	148	9,778,397	27,998
FR	stack		319	151	7,051,683	25,495

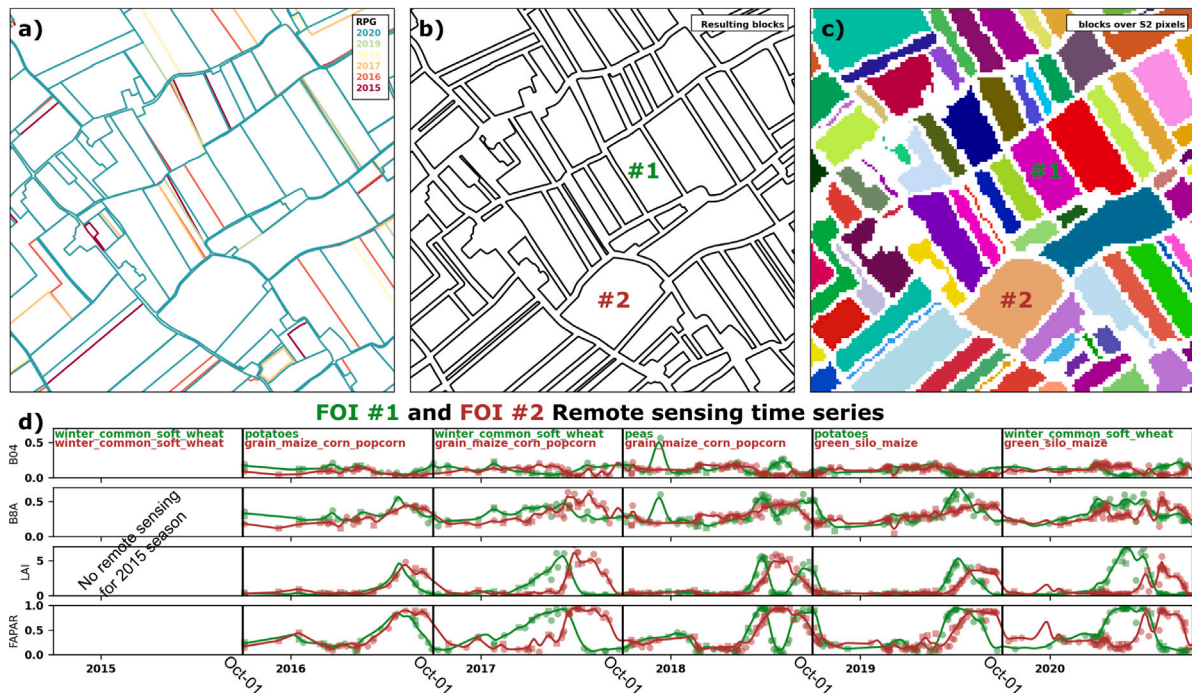


Fig. 1. Feature Of Interest (FOI) extraction, time series extraction and smoothing. (a) the map shows an overlap of the six GSA layers; (b) Resulting blocks corresponding to the intersection of the six GSA layers, reduced by an inner buffer; (c) rasterized version of the blocks used for extracting the S2 data; (d) full S2 and crop types time series of two selected FOI (shown in panels (a–c)). Yearly crop types are displayed on the top sub-panel. Input variables time series are displayed using daily observations (circles and squares correspond to S2 and L8 data, respectively) and smoothed signal (used as LSTM inputs).

non-surface-related information. We masked all pixels flagged as cloud, cloud-shadow, cirrus and snow.

Leaf Area Index (LAI) and Fraction of Absorbed Photosynthetically Active Radiation (FAPAR) Biophysical variables (BV) maps were derived from the S2 L2 A (20 m spatial resolution) and L8 L30 (30 m spatial resolution) products, using the BV-NET algorithm developed by Weiss and Baret (1999). It aims to retrieve the two BV from multispectral reflectance using the inversion of the radiative transfer model PROSAIL and a back-propagation ANN. Following the configuration of Delloye et al. (2018), the architecture of the ANN consists of two layers: (i) one layer with five tangent sigmoid transfer functions neurons and (ii) one layer with one linear transfer functions neuron. This configuration allows for greater dynamics in the output variables (Claverie et al., 2013). The HLS products are normalized

using the Bidirectional Reflection Distribution Function (BRDF) with a nadir view zenith angle and a variable sun angle (Claverie et al., 2018), while the S2 L2 A products are unadjusted with BRDF. We retained these data specifications and configured two BV-NET models to account for them. For both product types, the cosine of the solar zenith angle was included in the BV-NET input set; for S2 L2 A, the view zenith and relative azimuth angles were also included.

Only the Red and Near Infrared bands (NIR) were kept for further analysis; the remaining spectral bands were discarded. Four variables (LAI, FAPAR, Red band and NIR band) pixel-based maps were thus used to derive time series per FOI. Pixels whose centres fell within the FOI boundaries, reduced by a 15 m inner buffer (to prevent from using mixed pixels and reduce impact of the geometric precision), were averaged using a zonal statistics technique; flagged values (cloud,

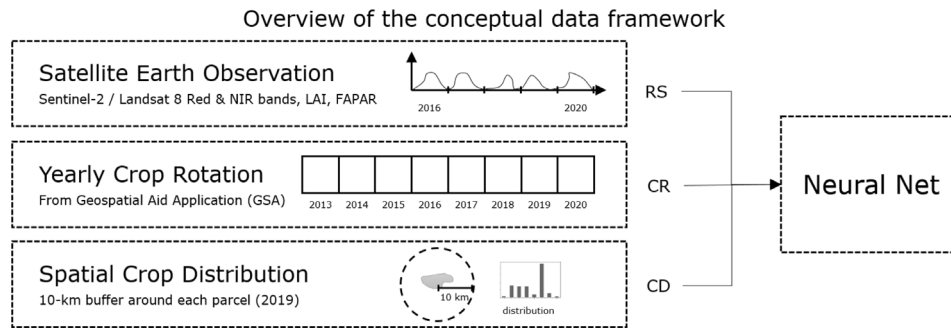


Fig. 2. Overview of the conceptual data framework for crop classification to leverage satellite optical time series, yearly crop rotation history, and spatial local crop distributions.

cloud shadow, cirrus or snow) from QA layers were not included in the averaging. FOI values were only considered valid if more than 75% of the LAI pixels were valid.

Despite filtering the data using relevant QA layers, the resulting FOI-based time series are still contaminated by missed cloud, cloud shadow, haze or dense atmosphere. To remove these remaining outliers, we applied a Hampel filter using red and NIR bands to discard cloud and cloud-shadow in the time series respectively; the parameters of the filter follow (Claverie et al., 2018).

Finally, filtered time series of four variables aggregated at FOI level were smoothed individually using a Whittaker filter.⁸ The time series are first gap-filled in time using a linear interpolation and a time step of 2 days. We applied the Whittaker configuration with V-curve optimization of the smoothing parameter and expectile smoothing using asymmetric weight, with an “Envelope” value of 0.9 and a tested lambda range between -1 and 1 (Eilers et al., 2017). This results in a smoothed time series with a time step of 4 days and no interruption between the seasons.

3. Methods

The feature extraction and the model architecture are first described in Section 3.1), followed by a description of the learning model and the integration of features as observations (Section 3.2). In Section 3.3, we delve into the early-season data augmentation technique. Then in Section 3.4, we explore the training process and the application of models in various countries. Finally, Section 3.7 outlines the processing facilities utilized in the study and includes links to the data and code.

3.1. Models description

A series of models were developed involving various configuration and input modalities. The three modalities are RS, CR and Crop Distribution (CD) (see the conceptual framework in Fig. 2). This section describes the model and the integration of the data as features.

Crop rotation, defined as sequences of crops throughout the seasons, has been modeled in a manner similar to a sequences of words within a language model (Mikolov et al., 2010). This modeling process is further enhanced by adding S2 time-series data, which is treated as analogous to the prosody of a speaker (Wöllmer et al., 2013a,b; Schuller et al., 2016), i.e. the pattern of intonation, stress and rhythm in a speech.⁹ Ultimately, the high-level spatial crop distribution features we add on the last layer of the network can be seen as the distribution of the words generally used by the speaker.

3.1.1. Features extraction

Crop Rotation

The crop types labels were extracted from the respective GSA and remapped using EuroCrops. This yields to a total of $V_{NL} = 141$ and $V_{FR} = 151$ classes, for NL and FR respectively, corresponding to $V = 225$ unique classes. We modeled the crop by a one-hot vector of size V and used it as an input to an embedding layer. For each FOI, we extracted the crop sequence which corresponds to the CR feature.

Remote sensing temporal integration

We integrated the RS time series into features using a sliding window of size $W = 1$ month with a step of size $s_w = 0.5W$, obtaining a sequence of $t_w = \frac{12}{s_w} = 24$ inputs windows for the 12 months of the season, from 1st of October to 31st of September. By utilizing this setup, we obtained some overlap between the windows, which should prevent loss of information by breaking the signal dynamics, albeit with a slight trade-off of redundancy in the features. On every time window, the four RS signals (LAI, FAPAR, Red band and NIR band) were integrated for each window using seven statistical functionals (Schuller et al., 2016) representing the signal as a fixed vector: average mean, standard deviation, min, max, median, first quartile, third quartile. As a total, we obtained $4 \times 24 \times 7 = 672$ features per FOI per season. Finally, we normalized each of the 24 features.

Crop distribution

We computed the total area of each crops in a circle of radius $r = 10$ km from each FOI and turned it to percentage of the total cropped area. The crop type distribution around the FOI accounts spatial for variation in terms of agricultural practices in relation with local agro-meteorological conditions, economic and historical factors. In the absence of major shocks, the distribution of the crop types in a region is expected to be stable over the seasons (Merlos and Hijmans, 2020), which determines the *a priori* probability of local crop occurrence. We integrate this local information by adding a vector representing the CD over the surrounding crop types centered around each FOI centroid. The spatial CD was always derived from the same season for computational reasons. We used the 2019 validation set season, i.e. not for the test season 2020. We rounded the probability at 10^{-4} , leading to some values being 0 when not null. Despite the EuroCrops harmonization, the crop lists of the two datasets (FR and NL) are not identical. We used the union of the two crop lists for both datasets.

3.1.2. Architecture of the models

Eight models were developed for the study. Their architectures are presented hereafter and summarized in Table 2.

⁸ as developed by <https://github.com/WFP-VAM/vam.whittaker>

⁹ The RS encoding could use any other encoder type, like a state-of-the-art model such as PSE-LTAE (Sainte Fare Garnot et al., 2020).

Table 2

Summary of the different models used in this paper, using Crop Rotations (CR), Remote Sensing (RS), and Crop Distribution (CD).

Models	CR	RS	CD	Modelisation-level		Hierarchical
				Within season	Between seasons	
IntraYE _{RS}	✗	✓	✗	✓	✗	✗
IntraYE _{MM}	✓	✓	✗	✓	✗	✗
InterYE _{Crop}	✓	✗	✗	✗	✓	✗
InterYE _{RS}	✗	✓	✗	✗	✓	✗
InterYE _{MM}	✓	✓	✗	✗	✓	✗
HierE _{RS}	✗	✓	✗	✓	✓	✓
HierE _{MM}	✓	✓	✗	✓	✓	✓
HierE _{final}	✓	✓	✓	✓	✓	✓

Baselines using year-independent models. We used two baselines models that treat the RS time series signal in a classical way without using hierarchical networks and without modeling the dynamics between the seasons. One unimodal model is using only RS data and another one is multimodal using RS and CR, based respectively on the works of Rußwurm et al. (2019a) and Quinton and Landrieu (2021). **These models are referred to as IntraYE_{RS} and IntraYE_{MM}, respectively.**

As stated in Rußwurm et al. (2019a), they only consider the time series of a single season, without incorporating a multi-season modeling approach for the RS data which is a key aspect of our proposed approach.¹⁰ This unimodal network InterYE_{RS} is the identical component utilized for encoding the RS signal at the season-level (one green Intra-Year Encoder (IntraYE) in Fig. 3). This provides a strong RS unimodal baseline. The second baseline that we add comes from the work of Quinton and Landrieu (2021), which is the strongest baseline among the three works incorporating CR as per with RS data. It integrates the CR modality by using a one-hot encoder vector of the past crop sequence. In our case, as we model the crops as words, this would mean a well-known simple vector representation for text called Bag-of-Words (Harris, 1954), hence we will call it Bag-of-Crops BoC. Although this type of representation is known to not work well for short texts like tweets or speech turns (Benamara et al., 2016; Barriere et al., 2018; Barriere, 2017; Barriere et al., 2017), we can expect better results with crop sequences which are considerably less complex than natural language.

Multi-year non-hierarchical models. Following the introduction, a set of novel model architectures is suggested hereafter, and their performance is evaluated in comparison to the existing baseline models. We first aimed to model the sequence of seasons with a recurrent encoder. These models use season-level features (that can be CR or RS) and are called Inter-Year Encoder (InterYE). This corresponds to the orange top Encoder in Fig. 3, modeling the sequence of seasons.

We modeled the multi-annual crop rotations in a language model fashion by representing the crops as tokens and learning to predict the next one. This model takes the past sequence of crops (c_1, \dots, c_t) as inputs and output the new crop c_{t+1} , modeling the crop rotation dynamics through the seasons. This corresponds to the orange InterYE in Fig. 3, if only using crop embeddings. It does not use the blue local crop distribution vector. **This unimodal model is denoted InterYE_{Crop},** corresponding to an unimodal Crop Rotations model.

Using solely previous rotations to forecast future crop yields results in inadequate performance due to the limited amount of information provided. Therefore, we decided to enhance the model's robustness by incorporating satellite data, leveraging either the consensus principle or the complementary principle (Xu et al., 2013). We enhanced the unimodal model InterYE_{Crop} by adding season-level information from RS. This corresponds to the orange InterYE in Fig. 3, with a green vector being the season-level concatenation of the RS signal (without being processed by the green IntraYE). It does not include the blue

local crop distribution vector. **This model is denoted as InterYE_{MM},** corresponding to a non-hierarchical multimodal model with RS. If the model only uses the RS modality, then it is denoted as InterYE_{RS}.

Multi-year hierarchical models. We chose to model jointly the inter-year and intra-year dynamics with a hierarchical model composed of one network modeling the RS dynamics within a season underneath another network modeling the rotation dynamics between the seasons. We processed the RS signal beforehand using another RNN, before concatenating this unimodal RS vector obtained with the crop embedding, in a hierarchical way.

We incorporated the sequential aspect of the RS time-series by processing the RS features at the season level with a first sequential encoder before adding their yearly representation into the second neural network modeling the crop types, leading to a hierarchical network with one top network modeling the sequence outputs of another bottom network (Serban et al., 2015).

Compared to InterYE_{MM} and InterYE_{RS}, there is another network modeling the RS signal at the bottom. This corresponds to the orange color top InterYE and the green color IntraYE in Fig. 3, modeling between the seasons as well as within a season. It does not use the blue local crop distribution vector. **This model is denoted as HierE_{MM},** corresponding thus to a hierarchical iIntraYE and InterYE to model both crop rotation and RS time-series.

We enhanced the model by adding the CD vector after the IntraYE because it is a high-level feature regarding the task we are tackling and the deeper you go into the layers the higher-level the representations are w.r.t. the task (Sanh et al., 2018). This corresponds to the full network presented in Fig. 3, including the blue local crop distribution vector. **This model is denoted as HierE_{final},** corresponding thus to a hierarchical IntraYE and InterYE to model the three modalities.

Encoders. We compared several type of models using different architectures and different modalities. Because our work mainly focuses on how to integrate multimodal data, we opted to use RNN-LSTM backbones, proven competitive for this kind of task (Rußwurm and Körner, 2020). Our method is also applicable using other encoders such as transformers (Vaswani et al., 2017) or Gated Recurrent Units (Chung et al., 2015).

For the Inter-year encoder, we first add an embedding layer to transform the crop type c_t at season t into a vector $\mathbf{emb}_t = f_e(c_t)$. This embedding vector \mathbf{emb}_t is used as input of the LSTM to produce a hidden state h_t at season t as seen in Eq. (1), which will be used to predict the next crop c_{t+1} in Eq. (2).

$$\mathbf{h}_t = \text{LSTM}_y(\mathbf{emb}_t, \mathbf{h}_{t-1}) \quad (1)$$

$$P(c_{t+1}|c_t, \dots, c_1) = f_c(\mathbf{h}_t) \quad (2)$$

The RS features were integrated at the season-level into a feature vector \mathbf{RS}_t before the modeling of the crop types by the LSTM. We feed the season t feature vector \mathbf{RS}_t into a neural network layer f_{rs} to reduce its size and then concatenate it with the crop embeddings before the LSTM (see Eq. (3)), using \mathbf{emb}_{MM_t} instead of \mathbf{emb}_t in Eq. (1).

¹⁰ This will be shown in next paragraph.

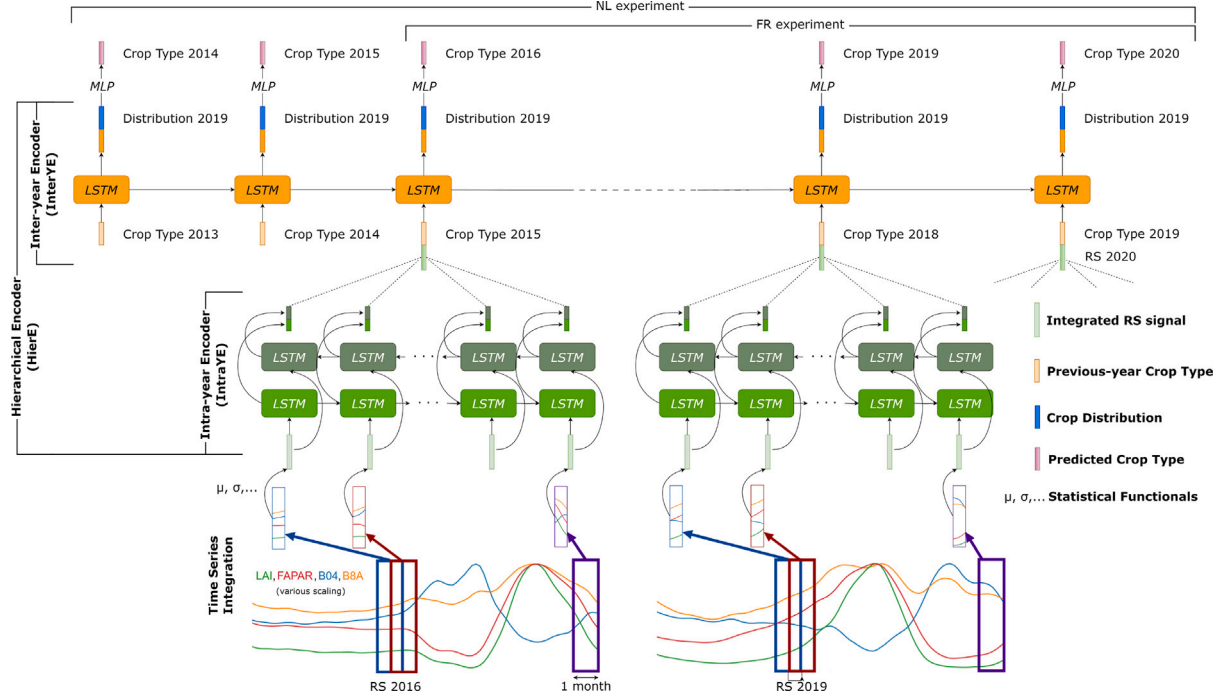


Fig. 3. Hierarchical Multimodal Model Conceptual Diagram. The two experiments, NL and FR, are represented with actual seasons used, 2013–2020 and 2015–2020 for NL and FR, respectively. Crop type of 2020 season is the final predicted label used for the test set.

$$\text{emb}_{M_{M_t}} = [\text{emb}_t, f_{rs}(\text{RS}_t)] \quad (3)$$

For the IntraYE, we chose to use a bidirectional LSTM (biLSTM) with a self-attention mechanism (Bahdanau et al., 2016) following the assumption that some parts of the season are more important than others to discriminate the crop type. The biLSTM is composed of two LSTM, one of which reads the sequence forward and the other reads it backward. The final hidden states are a concatenation of the forward and backward hidden states. For a sequence of inputs $[\text{RS}_{t_1}, \dots, \text{RS}_{t_w}]$ it outputs w hidden states $[\mathbf{h}_{RS_{t_1}}, \dots, \mathbf{h}_{RS_{t_w}}]$. The self-attention layer¹¹ will compute the scalar weights u_{t_w} for each of the $\mathbf{h}_{RS_{t_w}}$ (see Eq. (4)) in order to aggregate them to obtain the final state \mathbf{h}_{RS_t} (see Eq. (5)).

$$u_{t_w} = \text{att}(\mathbf{h}_{RS_{t_w}}) \quad (4)$$

$$\mathbf{h}_{RS_t} = \sum_w u_{t_w} \mathbf{h}_{RS_{t_w}} \quad (5)$$

For the crop distribution, we concatenated the hidden state \mathbf{h}_t of the LSTM with the crop distribution vector \mathbf{d} and mixed them using two fully connected layers f_{fc1} and f_{fc2} (see Eq. (6)). Hence, we obtain \mathbf{h}_{d_t} instead of \mathbf{h}_t before the final fully connected layer f_{fc} from Eq. (2).

$$\mathbf{h}_{d_t} = f_{fc2}(f_{fc1}([\mathbf{h}_t, \mathbf{d}])) \quad (6)$$

3.2. Automatic data- and knowledge-driven label aggregation

3.2.1. Rationale

Training and evaluating a model at large scale, on regions that contain different agro-climatic zones is complex due to the heterogeneity of the temporal and spectral representations of the crops and variability of the climate and agricultural practices. Indeed, if the labels distribution is highly variable between two datasets, one label that was representative in one domain would become not representative in the other. In this work we propose an aggregation of

the labels, that would be on the one hand representative of the dataset, and on the other hand thematically pertinent. In this way, it should be possible to evaluate the performances of the classification model at different scales: using all the labels from the region, even the ones with very few examples, then using an aggregation of labels that is representative of the region. This also offers the advantage to evaluate a model on two different datasets with a relevant evaluation on each dataset. We discuss this method in Section 5.1.

3.2.2. A hierarchical method to group labels

The labels obtained with the Hierarchical Crop and Agriculture Taxonomy from EuroCrops (see Section 2.1) have heterogeneous distribution and level of interest because of geographically constrained occurrence. We propose a method to merge non-representative crops together to only keep the most relevant in a region of interest by using its label distribution. The method is applied for the evaluation only. We take the best of both worlds by fusing expert knowledge and data-driven method. The method is applicable to any dataset, at any geographic scale and fully automatic.

Benefiting on the hierarchical structure of EuroCrops, we selected crops for which the number of FOI is above a given threshold (th). The crops with number of FOI below th are merged together with the other EuroCrops-sibling crops toward their parent-class. The remaining number of FOI falling down into the parent-class is compared with the threshold for an iterative process.

For both countries, we set the threshold th at 0.3% of the dataset size, which roughly corresponds to 2k samples for the NL and 20k samples for FR. We regrouped all the classes falling under the subclass “Permanent Crops” together as one for both datasets, as they are the simplest examples to classify when considering rotations. This was done to mitigate its impact by creating multiple labels for permanent crops. The automatic aggregation over FR and NL is shown in Fig. 4.

3.3. Early-season data augmentation

Applying the end-of-season models (i.e. trained with the data from the whole season) for early-season (i.e. using incomplete time-series)

¹¹ composed of a feedforward, a relu, another feedforward and a softmax layers.

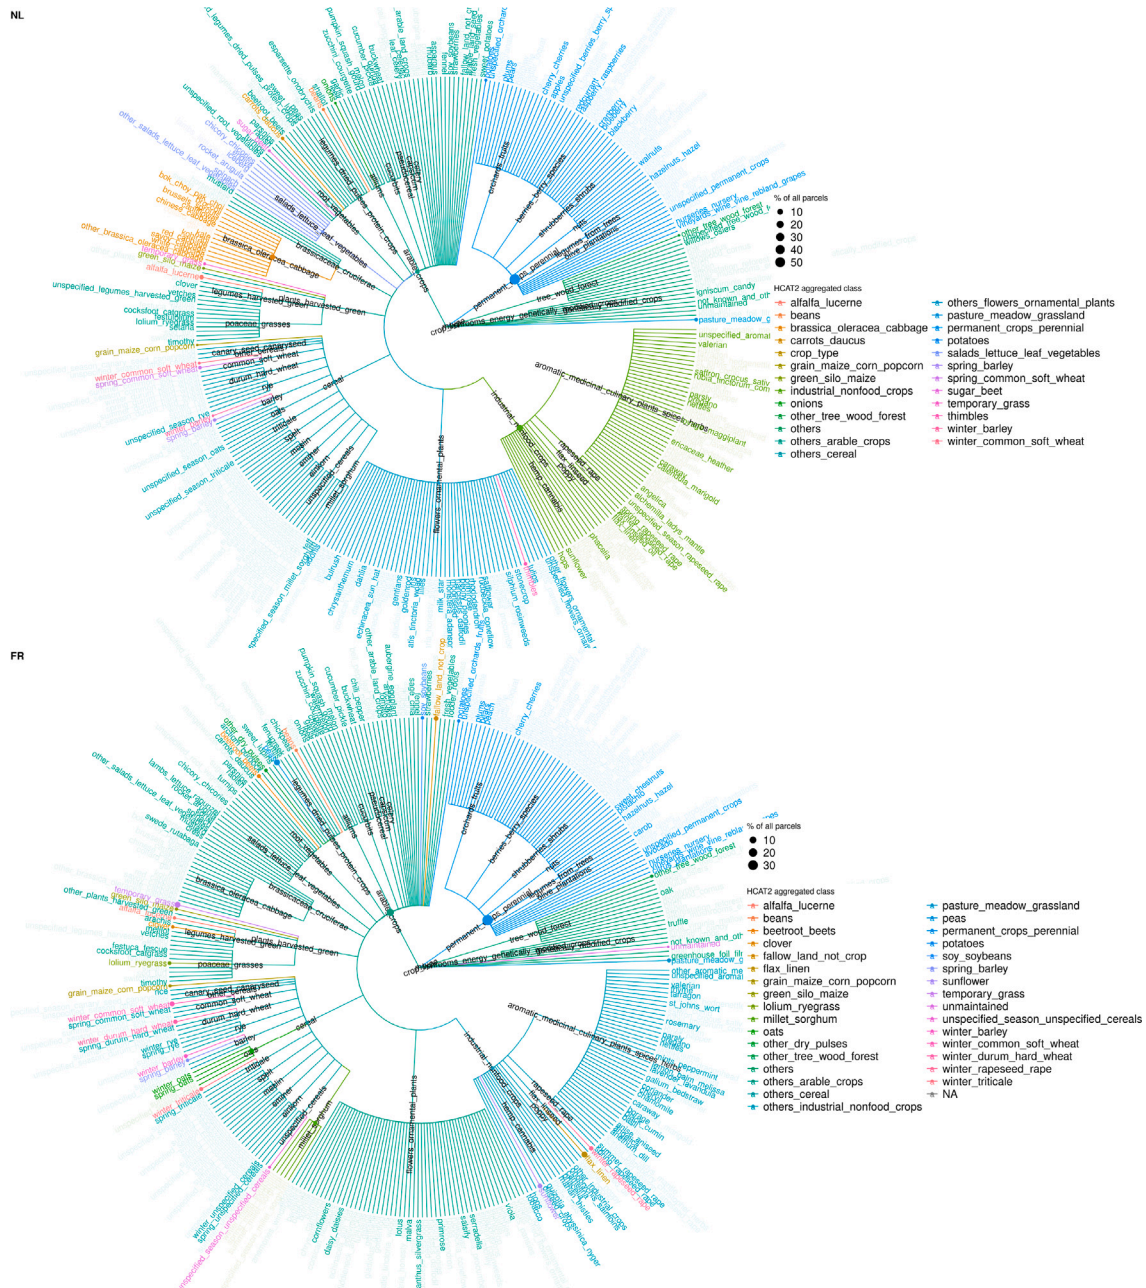


Fig. 4. Aggregated classes selected for validation in each country for FR (A) and NL (B) along with the distribution. The color highlight the crop type or crop group that were assessed for both country in red and only for one respective country in cyan.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Settings of the different experiments. $N \in \{4, 6, 8, 10\}$ for the few-shot experiments.

Name	Pre-Training	Training	Testing	# data from target	Models
few-shot-NL	\emptyset	NL	NL	2^N	HierE _{final}
few-shot-FR	\emptyset	FR	FR	2^N	HierE _{final}
Vanilla-FR	\emptyset	FR	FR	100%	All models
Vanilla-NL	\emptyset	NL	NL	100%	All models
0-shot-NL	\emptyset	FR	NL	0	HierE _{final}
Transfer-few-shot-NL	FR	NL	NL	2^N	HierE _{final}
transfer-NL	FR	NL	NL	100%	HierE _{final}
Transfer-FR	NL	FR	FR	100%	HierE _{final}

is not suitable. We propose a data-augmentation technique in order to help the model to classify a sample even without getting the full time series of the season. The idea behind this method is to force the model to predict the right class even though it does not observe the full time-series, as it would do at end-of-season.

We follow the approach of Barriere and Claverie (2022) by randomly cropping the end of the vector feature of the RS data $RS = (RS_1, \dots, RS_{t_w})$, t_w being the last bi-monthly date of the time-series varying between 10 and 24. By setting the minimum number of steps to 10, we ensure that each sample contains sufficient information (at least 5 months) to facilitate the training phase. Knowing the start of the time-series is October, it means we do not crop the end of the time series up to 1st of March. We used the same cropping size t among all the samples of the same mini-batch. For the hierarchical models, we used the same cropping size t among all the seasons.

3.4. Transfer learning between countries

We ran several experiments in order to take advantage of the normalized taxonomy that we used for both countries, by investigating the potential of transferring knowledge between different domains. For these purposes, we compared the performances of a model trained from scratch (i.e. Vanilla) and a model pre-trained over one country before being fine-tuned over another one. This pre-training allows to transfer knowledge from a source task and domain to a target task and domain.

We tested this approach in (i) cross-domain zero-shot setting, for which the pre-trained model does not see any samples of the new domain, and (ii) cross-domain few-shot setting,¹² for which the pre-trained model sees limited number samples of the new domain. For the cross-domain zero-shot setting, we used a network that was trained over one country on the other country, without fine-tuning it. For the cross-domain few-shot setting, we fine-tuned the network only on a subset of the target dataset, taking few examples representative of this dataset. We generated the few-shot subsets by randomly sampling 2^N (with $N \in \{4, 6, 8, 10\}$) examples of each of the aggregated class (see Section 4.1). For this, we sampled using the global crop distribution of 2019, which was used as validation cropping season. We think that this setup is realistic as we only sample from the aggregated classes that are the prominent ones in each of the datasets, and the ones to calculate the metrics to validate the models. We added more and more data increasingly so that all the examples from 2^{N_1} are comprised in 2^{N_2} , with $N_1 < N_2$. A summary of the different experiments can be seen in Table 3. We did not freeze any layer during the fine-tuning.

3.5. Dataset segmentation and validation

The datasets are generally split regarding time or space. Weilandt et al. (2023) proposed to compare the results of models trained with or without training data from the same cropping season as the test season, in end-of-season and early-season settings. They found out that the difference in performance was minimal.

In this work, we are interested to apply early-season setting and therefore segmenting the train and test sets among seasons. This provides two key advantages: we are in a real-life setting without in-situ data from the test season, and prohibit training with data from the end of the season in early-season setting. We trained our networks as for a sequence classification task, always with several seasons of data. The labels up to 2018 were used as training set, while the labels from 2019 were used as validation set and the labels from 2020 as test set (see Fig. 3). All results presented hereafter refer to the analysis of 2020 crop types, which are based on models trained with the period 2013–2019 for NL and 2015–2019 for FR, thus independent from the

2020 crop types observations. We zero-padded when no RS data was available (before 2016).

In order to confirm (Weilandt et al., 2023) findings, we added experiments with a temporal and spatial split of the data in Appendix B, showing that the results of a model trained on other seasons and other parcels still reach high results.

We validated the models using metrics calculated for four level of aggregation: (i) with all the labels, (ii) with the aggregation of the labels using the Automatic Hierarchical Label Aggregation, (iii) with a set of crop of interest from the aggregation that were recognized important by a Food Security expert and with or (iv) without classes *others* and *grassland* (majority class in NL dataset). We used macro-average of the Precision, Recall and F1-score because the dataset is imbalanced. We also used micro-F1 score (m-F1) for the last setting, because accuracy is not possible when focusing on a subset of the classes. In addition, we computed the Accuracy as a general metric.

3.6. Implementation

We trained all the networks via mini-batch stochastic gradient descent using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 10^{-3} and a cross-entropy loss function. The number of neurons for the crop embedding layer (from 16 to 128, best 64), both the RNN internal layers (from 64 to 512, best 256) and the fully connected RS layer f_{rs} (from 32 to 256, best 128), as well as the number of stacked LSTM (from 1 to 4, best 3) were chosen using hyperparameters grid search of power of 2 on a subset of the NL dataset. The sizes of the layers f_{c1} and f_{c2} are the same as the one from the second RNN state \mathbf{h}_r .

3.7. Processing facility, data and code

The EO extraction and processing, the classification and the benchmarking were performed on the JRC Big Data Analytics Platform (BDAP) using an HTCondor environment (Soille et al., 2018). The platform¹³ has been built upon the near-data processing concept, which prescribes placing the computing facility close to the storage units to avoid the bottleneck of delaying or degrading interconnection. Experiments with the neural networks were run using PyTorch 1.4.0 (Paszke et al., 2019) on a GPU Nvidia RTX-8000 using CUDA 12.0. The training phase allows to process between roughly 5k and 30k examples per seconds, with each example containing 4 seasons of data.

The data extracted and used for this study are openly available on the public FTP.¹⁴ The code for the data processing, the labels aggregation, and deep learning experiments will be freely available after publication.¹⁵

4. Experiments and results

4.1. Feature extractions

4.1.1. Crop rotation

The crop label categories for the 2020 test set correspond to a long-tailed class distributions, as shown for the 32-class and 24-class aggregations for the French and the Dutch data sets in Figs. 5 and 6, respectively. The models are finally validated on a set of crops of interest from the 32-class and 24-class aggregation. Those 8- and 12-class, respectively for the NL and FR, were identified by experts as essential for crop production monitoring and food security. They are shown in green in Figs. 5 and 6.

¹³ <https://jeodpp.jrc.ec.europa.eu/bdap/>

¹⁴ The data can be downloaded on <https://jeodpp.jrc.ec.europa.eu/ftp/jrc-opendata/DRLL/CropDeepTrans/>.

¹⁵ Add URL after publication.

¹² i.e. when the domain changes as per with significantly new labels.




Fig. 5. Distributions of the crop types in the NL dataset for the test season (2020), after aggregation. Green bars are the selected crops for the 8-class evaluation; Red bars refer to the grasslands, and blue bars to the remaining crops.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Distributions of the crop types in the France dataset for the test season (2020), after aggregation. Green bars are the selected crops for the 12-class evaluation; Red bars refer to the grasslands, and blue bars to the remaining crops.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.1.2. RS-based features

First, the RS data retrieved as described in 2.3 were obtained. To support our development, we have released a SITS dataset, which is significantly larger in terms of number of parcel and length of the time series than any other existing SITS datasets (Selea, 2023). The processing of RS data, which incorporates a smoothing algorithm, sets this dataset apart from others in its category. The final dataset which consists in the full time series of more than 7M FOI, for a total of more than 35M FOI-season (NL : 5 seasons \times 596k FOI; FR: 5 seasons \times 6,49M) are available for download, as well as the extracted features used in the experiments.¹⁶

4.2. In-country end-of-season results

The results of the various models on France and the NL are shown in Tables 4 and 5 with respect to four distinct classification schemes, ranging from a fine-grained 141-class scheme for NL (resp. 151 for FR) to a coarse-grained 10-class scheme (resp. 14 for FR). The 8-class scheme (resp. 12 for FR) is the same than the 10-class one, in order to only focus the performances on the crops of interest. Tables 4 and 5 aim at showing the interest of our multi-modal method compared to

what is generally used in the field, by using only the remote sensing of the season, in an independent way.

4.3. In-country early-season results

We compared the performance of our best model trained without the data-augmentation technique, i.e. trained solely on end-of-season classification examples, with an early-season model of the same architecture trained using the proposed data-augmentation technique (see Section 3.3). Comparisons were also made with a unimodal model using only the CR, and an unimodal IntraYE_{RS} enhanced with the data-augmentation technique.

Fig. 7 shows the performances of the model in terms of micro-F1 on the set of 10 crops. For sake of clarity, we focused on the results of the model trained using this data-augmentation over the NL.

4.4. Cross-country transfer learning results

Table 6 shows the performance of our best architecture model on the NL. The neural network was trained either from scratch or using a pre-trained model. We compared its performances on a few-shot setting against using the full NL dataset. The table shows that pre-training the model led to improved performance in terms of almost in terms of all metrics for all tasks. Furthermore, the performance increases

¹⁶ <https://jeodpp.jrc.ec.europa.eu/ftp/jrc-opendata/DRLL/CropDeepTrans/>

Table 4

Results over Netherlands of the end-of-season classification models with different modalities: Remote Sensing (RS), Crop Rotations as embeddings (C) or BoC, and Spatial Crop Distribution. The metrics shown are macro Precision (P), Recall (R) and F1 score, as well as accuracy and micro-F1 score (m-F1).

Labels	# Modalities	141-class				24-class				10-class				8-class			
		P	R	F1	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1	m-F1
InterYE _{Crop}	1 (C)	36.0	25.5	27.4	76.2	53.3	37.2	39.1	76.5	51.8	43.0	43.5	77.7	43.3	35.5	34.9	53.6
IntraYE _{RS} 2019a	1 (RS)	27.4	20.9	20.4	89.8	64.0	60.9	60.4	90.3	78.8	75.9	74.5	92.9	76.1	72.6	70.8	87.8
InterYE _{RS}	1 (RS)	22.8	17.7	17.1	89.1	59.2	58.5	57.3	89.6	71.2	73.4	72.0	92.1	67.0	69.6	68.0	85.6
HierE _{RS}	1 (RS)	20.7	17.5	16.7	90.2	64.3	61.0	61.2	90.9	80.5	74.4	74.3	93.5	78.0	70.4	70.3	88.3
IntraYE _{MM} 2021	2 (RS+BoC)	55.6	39.7	43.2	92.8	76.6	69.8	72.1	93.1	83.0	80.5	80.9	94.7	80.2	77.9	78.0	90.0
InterYE _{MM}	2 (RS+C)	41.1	33.0	33.6	92.2	70.8	70.5	69.9	92.6	82.2	79.7	80.4	94.5	80.2	76.3	77.5	89.5
HierE _{MM}	2 (RS+C)	47.3	38.7	39.7	93.3	74.7	75.5	74.7	93.7	85.2	81.9	83.1	95.2	83.6	78.8	80.6	91.1
HierE _{final}	3 (All)	47.1	39.3	40.2	93.6	76.6	75.8	75.7	94.0	86.7	81.9	83.6	95.5	85.3	78.7	81.1	91.6

Table 5

Results over France of the end-of-season classification models with different modalities: Remote Sensing (RS), Crop Rotations as embeddings (C) or Bag-of-Crops (BoC), and Spatial Crop Distribution. The metrics shown are macro precision, recall and F1 score, as well as accuracy and micro-F1 score (m-F1).

Labels	# Modalities	151-class				32-class				14-class				12-class			
		P	R	F1	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1	m-F1
InterYE _{Crop}	1 (C)	35.6	31.0	31.7	66.0	43.7	38.8	38.7	66.2	38.9	34.3	31.7	69.1	30.9	26.4	23.0	42.7
IntraYE _{RS} 2019a	1 (RS)	22.9	15.7	15.2	64.0	51.1	46.0	44.6	64.5	69.8	62.2	64.7	75.7	69.3	59.7	63.1	74.6
InterYE _{RS}	1 (RS)	21.3	13.2	12.6	54.9	46.5	41.5	39.2	55.5	63.9	59.6	60.2	72.2	62.7	57.4	58.5	71.2
HierE _{RS}	1 (RS)	25.3	19.0	18.8	66.3	55.5	50.7	50.1	66.9	72.5	65.5	67.7	76.9	71.9	63.2	66.1	76.5
IntraYE _{MM} 2021	2 (RS+BoC)	52.7	32.4	35.9	82.7	70.1	59.3	61.8	82.8	78.1	68.7	71.0	86.6	76.2	65.6	68.0	80.3
InterYE _{MM}	2 (RS+C)	45.9	35.2	36.4	82.4	67.7	60.5	62.4	82.7	72.7	67.4	69.2	86.1	70.0	63.6	65.8	77.5
HierE _{MM}	2 (RS+C)	50.2	41.9	43.2	84.8	70.7	67.6	68.2	85.0	77.0	73.4	74.9	88.4	75.0	70.2	72.3	81.8
HierE _{final}	3 (All)	45.1	37.3	38.1	85.4	72.1	68.8	69.2	85.7	79.8	76.1	77.6	89.1	78.1	73.5	75.4	83.6

Table 6

Results over Netherlands of the few-shot final classification models, with or without pre-training over France. N represent the number of examples shown per aggregated class on the target dataset. The metrics shown are macro precision, recall and F1 score, as well as accuracy and micro-F1 score (m-F1). N stands for the Few-Shot size.

Labels	N	141-class				24-class				10-class				8-class			
		P	R	F1	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1	m-F1
✗	0	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅	∅
	16	5.8	5.1	4.8	70.8	23.7	21.4	20.4	71.1	38.5	37.4	36.3	73.6	38.5	37.4	36.3	45.3
	64	2.7	2.5	2.2	69.2	17.1	13.1	12.5	69.4	27.3	25.7	23.3	69.6	27.3	25.7	23.3	34.7
	256	4.2	4.8	2.9	66.5	18.2	16.9	14.1	66.8	25.0	23.2	20.5	68.1	25.0	23.2	20.5	20.4
	1024	19.6	13.3	12.4	80.8	53.6	39.8	37.2	80.3	69.7	60.4	61.5	84.0	69.7	60.4	61.5	76.3
✓	0	5.7	4.8	4.2	47.3	14.7	15.1	11.1	46.6	20.6	19.7	16.6	46.9	12.3	7.4	8.4	24.5
	16	12.2	7.8	7.6	70.3	30.5	23.8	24.5	70.4	37.9	33.9	34.0	72.3	37.9	33.9	34.0	45.2
	64	16.7	13.6	13.5	74.7	41.9	38.7	38.1	75.0	51.6	45.4	46.6	76.4	51.6	45.4	46.6	54.4
	256	25.8	21.4	20.8	82.5	55.6	51.1	50.6	82.7	67.3	58.0	60.1	84.6	67.3	58.0	60.1	69.2
	1024	32.7	27.3	26.0	84.9	61.3	57.3	54.3	84.9	73.8	72.0	71.6	87.0	73.8	72.0	71.6	80.9
✗	All	47.1	39.2	40.2	93.7	76.6	75.8	75.8	94.0	86.7	81.9	83.6	95.5	85.3	78.7	81.1	91.6
✓	All	42.5	35.3	36.0	92.8	67.3	53.4	55.9	94.2	89.9	82.2	85.3	95.7	88.8	77.6	82.3	91.8




Fig. 7. Comparison of early classification using different modalities, with/out data augmentation (macro-F1 with 10 classes) on Netherlands.




Fig. 8. Comparison of the F1-scores by crops of the best hierarchical multimodal model and the model using different modalities (Crop rotation, Remote sensing only and all) on the Netherlands. We used the InterYE_{Crop} (in blue), IntraYE_{RS} (in orange) and HierE_{final} (in green) models. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with the number of training examples, and the highest performance was achieved for 1024 training examples for each aggregated class. Notice that the non-pretrained model (i.e., red-cross) using $N = All$ corresponds to HierE_{final} from Table 4.

5. Discussion

The results are discussed in the five first subsections of the section, then followed by a presentation of the limitations (Section 5.6) and recommendations for future research (Section 5.7).

5.1. The label aggregation method

In line with second objective, we introduced a distinct benchmarking approach that utilizes HCATv2, the hierarchical crop type classification system of EuroCrops. This is a method to output pertinent metrics on a dataset that has many classes with a long tail distribution, by grouping together the minoritarian classes that are similar (i.e. close in the HCATv2 graph). In Table 4 (NL) and Table 5 (FR), the aggregation level is displayed in the 24-class and 32-class columns, respectively. This means that the performance results for both countries incorporate the internal distributions. Overall, the accuracy scores for both countries did not exhibit significant variation when accounting for all EuroCrops classes (93.6 and 85.4, for NL and FR, respectively) or after aggregation (94.0 and 85.7), suggesting that the models performed well on the very dominant classes. However, the macro F1-scores experienced an improvement due to the merging of classes with a limited number of samples: it reveals the interest of the aggregation method for classifier performances evaluation on a dataset with a high number of crop classes.

5.2. Hierarchical multimodal models: a way to gain in performances

The third objective of the study (Section 1.2) is to evaluate the performances of models relying on diverse model configurations involving different modalities.

For the NL, the results of Table 4 indicate that the best performance was achieved by the final model (HierE_{final}) combining the three modalities. It achieved a macro-F1 score of 40.2% on the 141-class scheme and a micro-F1 score of 91.6% on the 8-class scheme. With the exception of HierE_{RS} , which failed to learn during the optimization process even on the training set (achieving an F1 score of 0 for both the 8- and 12-class settings), all models based on LSTM to process RS time series achieve higher performances compared to the approach

of concatenating statistical functional vectors of each sliding window. Notably, the unimodal IntraYE_{RS} outperformed InterYE_{RS} , while the multimodal IntraYE_{MM} and HierE_{MM} outperform InterYE_{MM} . More details are found in the next Section 5.3.

Several similarities and notable differences can be observed when comparing the models between the NL 4 and FR 5. First, the general performance metrics were lower for FR than for the NL. This difference may be attributed to FR having more classes than the NL. Second, the performances of the RS model's performance using only one season of context (InterYE_{RS}) were significantly lower than for the same one on NL, relatively to the other models. The higher variance of the RS data in FR, resulting from its larger size and greater diversity compared to the NL, may contribute to the lower performances observed in this scenario.

5.3. Performances per crops

The utilization of various modalities and their combinations for each of the primary crops in both countries is depicted in Figs. 8 and 9. It demonstrated an upward trend in the level of improvement with an increase in the number of modalities employed. The benefits of crop rotation were more significant for certain crops such as pasture, while others such as beetroot are harder to predict without RS signal. This suggested that the crop rotation for FR, limited to only starting from 2015, might not offer enough information to accurately predict crops in complex or irregular crop rotation sequences.

Fig. 10 provides more detailed information on the performance of the best model in the NL. It displays the F1 scores for eight crops of interest in the NL, based on the time of the season used for prediction. Furthermore, for the same crops, we analyzed the time series data for the 2020 cropping season (i.e., the test season) averaged at the country level for each of the four remotely sensed variables on Fig. A.2, accompanied by the standard error shown on Fig. A.3. These visualizations highlight both the variability between crops and the potential confusion that can arise between different crop types. As the season progresses, there is a noticeable trend of increasing F1-scores, highlighting the improved accuracy of crop discrimination. However, there are evident disparities among crops. Notably, grassland demonstrates a high F1-score at the season's onset and maintains this level from the end of spring onwards. This is attributed to the fact that the unique temporal signature of grassland phenology does not provide significant discriminatory power as the season advances. Contrarily, crops with distinct morphological and phenological characteristics, such as potatoes, green silo maize, and sugar beet, achieve higher




Fig. 9. Comparison of the F1-scores by crops of the best hierarchical multimodal model and the model using different modalities (Crop rotation, Remote sensing only and all) on France We used the InterYE_{Crop} , IntraYE_{RS} and HierE_{final} models.




Fig. 10. F1-score for each crop group along the season for NL. When crop have winter and spring varieties, the spring varieties are represented as dashed lines.

F1-scores earlier in the season. The F1-score of barley reaches a peak of 0.75 by the end of July. An intriguing observation is that during the early spring, spring barley exhibits a higher F1-score. However, by June, winter barley surpasses it. This shift can be elucidated by the differential crop calendars that spring and winter barley adhere to.

The end-of-season performances of the models in the NL and FR, denoted as HierE_{final} , are presented using 5 km grid cell maps in Figs. 11 and 12 respectively. The maps reveal notable regional effects, particularly in FR. For instance, in Brittany (located in the north-west of FR), lower accuracy was observed for most crops, especially *winter barley*. It worth noting that despite labeling the crop types uniformly by country, variations in crop varieties, climates, and agricultural practices among farmers have an impact on phenology, consequently affecting the RS signal. Consequently, using a single dataset for training the model resulted in heterogeneous performances over a large country like FR. The development of regional models is thus highly recommended. Furthermore, the model lacks information regarding

the specific locations, except through the proxy of crop distribution. Incorporating geographic coordinates and/or weather variables as model inputs could contribute to accounting for spatial variations over large areas, such as FR.

The confusion matrices for both countries can be seen in Fig. 13. Because of the highly imbalanced distribution of the labels, we classically normalized each matrix regarding the predicted values (i.e. by row) for visibility reasons. Hence, the diagonal cells represent the precision of the model on each class. We display the values in percentage, meaning they go from 0 to 100. These matrices allow for grouping crops based on the observed confusions. In general, for both countries, there are instances of misclassification between crops, indicating difficulties in distinguishing between certain crop types: (i) between *green silo maize* and *grain maize corn popcorn*, (ii) between winter cereals (*winter common soft wheat*, *winter barley*, *winter triticale*), (iii) between spring cereals (*spring barley*, *spring common soft wheat*), (iv) between summer crops *sunflower*, *millet/sorghum*. These confusions

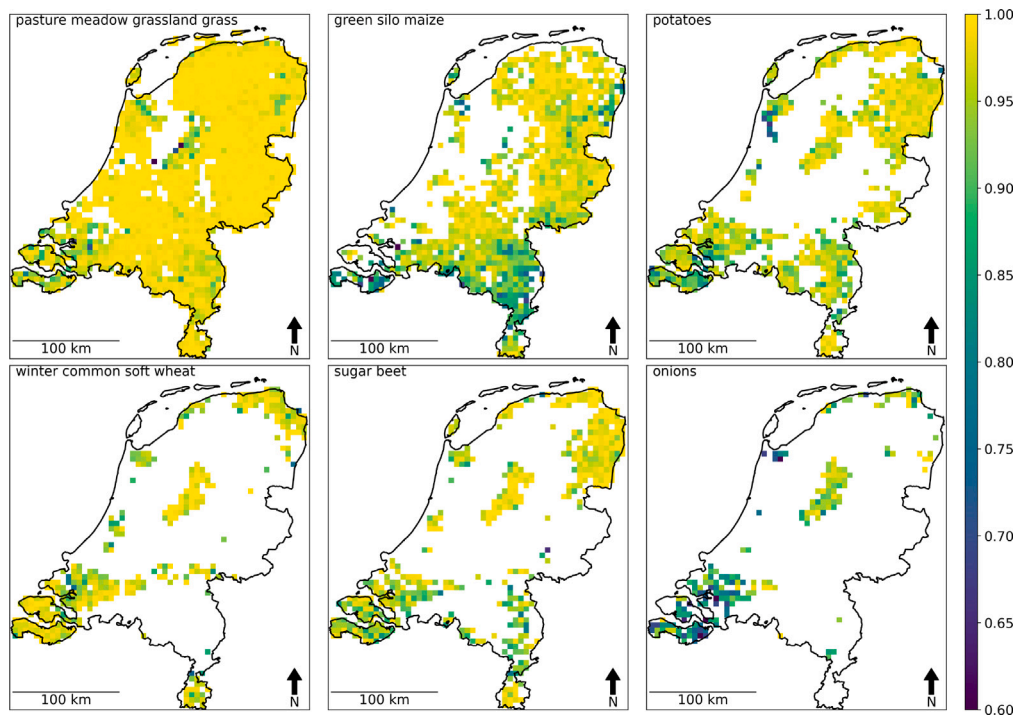


Fig. 11. Map of F1 for the six most important crops over The Netherlands. The F1-score is computed for each crop and for each 5 km grid cell. Grid cells with less than 50ha (i.e., 2% of the land) of the given crops are not plotted. Map projection is EPSG:3035.

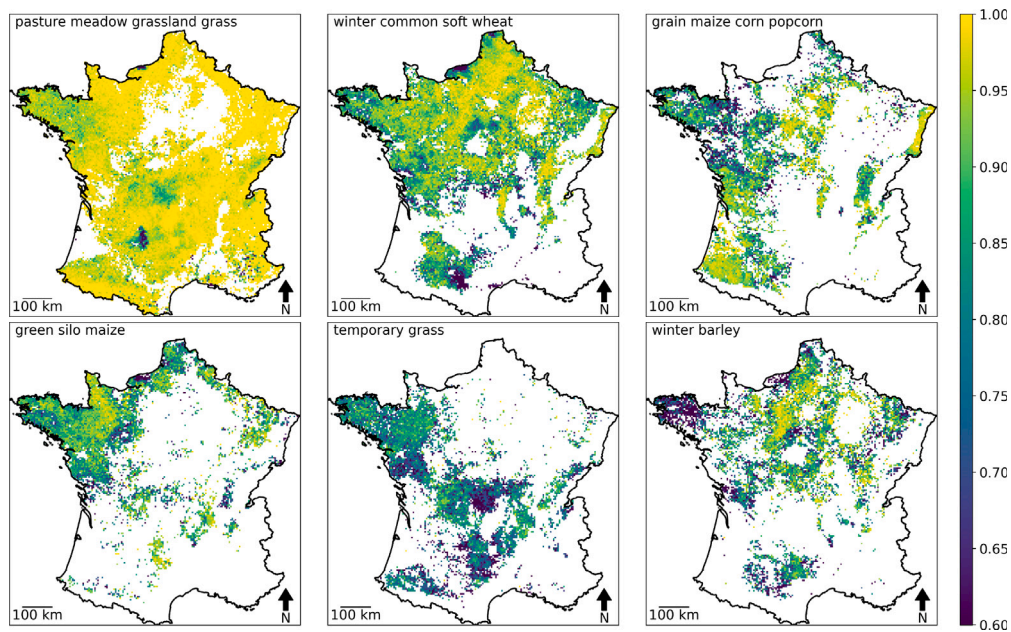


Fig. 12. Map of F1 for the six most important crops over Metropolitan France. See legend of Fig. 11.

are anticipated as they occur with synchronous phenologies of the crops that could differs significantly from one region to the other in Europe (d’Andrimont et al., 2020, 2021; Meroni et al., 2021).

5.4. Early-season models

The use of sub-setting technique, intended for early-season classification, was found to be ineffective in improving the performance of a model solely trained on remote sensing data, as shown in Fig. 7.

However, when the multimodal model including crop rotations was applied, it resulted in improved performances as early as May. It is worth noting that the overall performances of the multimodal model was observed to be inferior to that of a unimodal crop-only model. This was due to the model overemphasizing the RS modality as the season progressed. To address this issue, a gate mechanism could be incorporated, as proposed in Arevalo et al. (2017) and Chen et al. (2017), which selectively discards noisy modalities. By utilizing the multimodal hierarchical configuration, the model achieved around

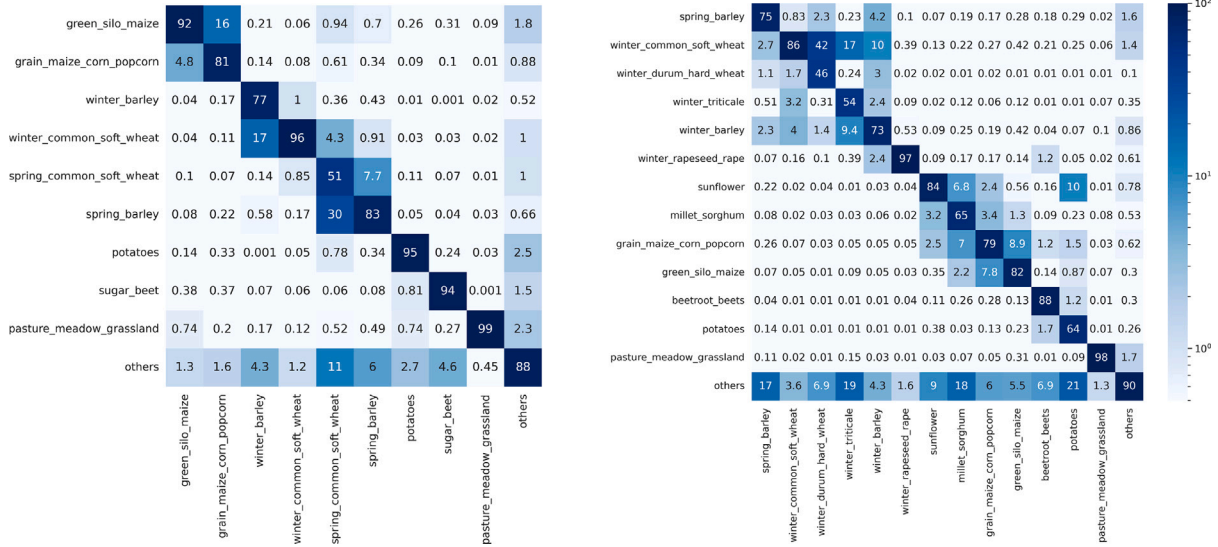


Fig. 13. Confusion matrices for the 10- and 14-class settings for Netherlands and France, respectively. True labels are in rows and predicted label are in columns, with a column-level normalization.

mid-July 95% of the end-of-season overall accuracy. This corresponds to the period when the winter crops are harvested, and the summer crops reach their peak vegetation.

In Fig. 10, the HierE_{MM} model data-augmented model was assessed through time. F-score are provided for the 10 classes of interest for NL, as defined in Section 4. Overall, there was an offset of the curve rise-up. The shift is earlier (from mid-April to mid-June), which was consistent with other studies (Rußwurm et al., 2023).

5.5. Cross-country transfer learning results

According to Table 6, there were two notable distinctions between the models pre-trained on FR and the ones trained from scratch. The first distinction pertained to few-shot setting, where the pre-training enables not only superior but also more consistent results in few-shot classification, reducing dependence on the specific examples used during few-shot training. The second distinction lies in the full-data setting, where we observed that the from-scratch model performs better when validated with 141 and 24 classes (with an m-F1 score of 40.2 and 75.8, respectively, compared to 36.0 and 55.9), while the pre-trained model was better when using fewer classes. These results highlighted an intriguing behavior, suggesting that the pre-trained neural network was more general and less prone to over-fitting on the prominent classes of the NL dataset. We can conclude that transfer learning proves beneficial when limited labeled examples are available, while in the full-dataset training mode, it enhances the model’s performance for general classes at the expense of specific classes in the dataset.

5.6. Limitations

The primary constraint of this study, particularly in developing countries and numerous other nations, is the requirement for digitized parcel boundary data. Another limitation is the necessity to obtain information regarding the crops cultivated in the previous season. The impact of noisy input data, such as past crop information derived from a prediction system rather than ground-truth data, remains unknown in terms of the system’s response. Exploring this aspect constitutes an intriguing avenue for future research.

To encode the RS signal, we employed a backbone that involved several stages: EO preprocessing steps, averaging the pixels of the FOI, smoothing the time series signal, and ultimately applying a

temporal context window with statistical functionals. Although it is straightforward to evaluate the computational resources needed for training the neural network (see Appendix D), quantifying the time required for this sequence of preprocessing steps can be challenging, as it is not easily achievable on the fly. Studies proved that other methods were more efficient in terms of performances (Sainte Fare Garnot et al., 2020). An inherent enhancement to encode the RS signal in a more effective manner would involve employing an end-to-end approach. This approach entails learning the aggregation of RS data and integrating its representation into a neural network, similar to architectures such as CNN-temporal or CNN-RNN (Pelletier et al., 2019; Sainte Fare Garnot et al., 2019) or more advanced structures like PSE-LTAE (Sainte Fare Garnot et al., 2020; Quinton and Landrieu, 2021; Weilandt et al., 2023). By adopting these powerful architectures, the encoding of RS signals can be optimized, thereby potentially improving the overall performance and accuracy of the system. Furthermore, Rußwurm and Körner (2020) have demonstrated that preprocessing steps in RS can enhance classifier accuracy, although the improvement tends to be marginal when applied to advanced learning models.

Finally, we would like to try a meta-learning method to tackle the few-shot learning more effectively, by using specific algorithm like MAML (Finn et al., 2017) and MetaNorm (Du et al., 2020). Also, a domain adaptation method like the one proposed by Capliez et al. (2023), but spatially and not temporally, or the one of Nyborg et al. (2022b) would be very useful for the few-shot setting.

5.7. Recommendations

Regarding future research directions, there are several avenues for further exploration. One potential direction is to integrate knowledge from the EuroCrops ontology graph inside the learning model, for example by creating multi-level embeddings of each crop. This could improve the ability of the model to capture the complex spatiotemporal variability of crops. It would be also possible to integrate knowledge at the loss level, in a way similar to what Turkoglu et al. (2021) proposed.

A more complex way to fuse the modalities together could also be explored, such as using a Gated Multimodal Unit (Arevalo et al., 2017). This could lead to better integration of the different data modalities and improved performance of the model.

It would also be valuable to investigate the results at a more regional/local level, especially for FR with its large landmass, crop

diversity, and meteorological conditions. Local hierarchical clustering and the performance of the model at the regional level could be examined to gain a deeper understanding of how the model performs in different regions. We saw that the results for FR were lower, possibly due to the diversity of crops and the distribution vector used. Investigating the effect of a region-specialized model, for example by fine-tuning using Adapter layers (Poth et al., 2020), could be a potential solution. In addition, adding meteorological features and investigating their impact could be worthwhile, especially in the case of extreme events. Methods such as (Tseng et al., 2021b) or using learned embeddings that represent the time of the thermals (Nyborg et al., 2022a) could be explored.

Another area to investigate is the potential of a specific loss function for the early season model, like the one proposed in Rußwurm et al. (2023), as per our simple data-augmentation technique. This could lead to better performance in the early season, which is an ongoing challenge for crop classification.

Other potential avenues for future work include adding more countries to the experiments, but also testing the system with different backbones, allowing ingestion of the EO raw time series as they are. Furthermore, assessing the effects of different combinations of bands and sensors, encompassing various specifications such as spatial, temporal, angular, and spectral aspects, would be necessary to determine their influence on performance. This last step would prevent reliance on man-made filters like Hampel or Whittaker and man-made features like FAPAR and LAI, as they contain filtered information, filtering some that may be useful for the final task (Trigeorgis et al., 2016).

Overall, these future directions could further improve the accuracy and generalization of the proposed multimodal approach for crop classification.

6. Conclusions

In conclusion, we proposed a multimodal hierarchical approach for crop classification that leverages crop rotation history, optical remote sensing signals, and local crop distributions. We released a large harmonized time series dataset of 7M Feature Of Interest (FOI) for a total of around 35M FOI-season. We introduced a new dataset-agnostic method relying on data and expert knowledge for aggregating crops, allowing to evaluate a classifier on a specific region in a meaningful way. Finally, we propose a data-augmentation method to boost the results in early-season setting. Our approach achieved high accuracy without in-situ data from the test season and showed promising results for cross-domain generalization through transfer learning and few-shot learning experiments. Pre-training on a dataset improves transferability between countries, allowing for cross-domain and label prediction and stabilization of the learning in a few-shot setting. Our approach can contribute significantly to agriculture management and policy monitoring.

CRedit authorship contribution statement

Valentin Barriere: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Martin Claverie:** Writing – review & editing, Writing – original draft, Visualization, Resources, Project administration, Methodology, Investigation, Data curation, Conceptualization. **Maja Schneider:** Writing – review & editing. **Guido Lemoine:** Writing – review & editing. **Raphaël d’Andrimont:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Project administration, Methodology, Investigation, Formal analysis.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data is already available using a link.

Acknowledgments

The authors would like to thank Momtchil Iordanov for his support for visuals and Loïc Landrieu for the useful comments on the manuscript. They also would like to thank the Big Data Analytics project for their continuous support. V.B. has been funded by the grant National Center for Artificial Intelligence CENIA FB210017, Basal ANID.

Appendix A

See Figs. A.1–A.3

Appendix B

In our main paper, all the experiments were ran using a temporal split between the train, dev and test. In Table A.1 we show the results of experiments using a spatio-temporal split. It means the parcels used for testing are from different years than the year used for training, but also they are different parcels. We used 90% of the parcels for training, and 10% of the parcels for testing. Our results show a slight deterioration of the performances, coinciding with the ones of Weilandt et al. (2023) as we stated in the main paper.

Appendix C

In our main paper, all the experiments were ran using a set of features comprising FAPAR, LAI, B4 and B8 A. As we stated in the main paper, the focus of this work is not on the input features sets, as our architecture is modular and the encoder we used can be changed for another one, using different features. Our work on modeling focused on showing the interest of an hierarchical architecture.

In Table A.2 we show the results of experiments using more features obtained from more spectral bands, over data from NL. In addition to the initial feature set, we added B2, B3, B8 given at 10 m resolution, and also B11 and B12 that were interpolated at 10 m. With 9 time-series instead of 4, we obtained a feature vector of size 63 instead of 28. We also ran the same experiments using the 4 10 m bands of S2: B2, B3, B4 and B8. Our results show an amelioration of all the models that are using more features. We also observe the same improvement between the different architectures proposed, confirming that the hierarchy helps even with a broader feature set. It is noteworthy that the performances of the InterYE_{MM} are dropping a bit when using only the 10 m bands of S2. We explain this by the fact that there is no sequential processing of the RS time series for this model: the RS vectors of functionals from the different sliding windows are concatenated into one. For this reason, higher-level features like FAPAR and LAI are helping to reach higher results. This is not true for all the other models containing a IntraYE, which process the time series with an LSTM and can extract the useful features (Trigeorgis et al., 2016).

FR / NL crops harmonization (40 crops selected)

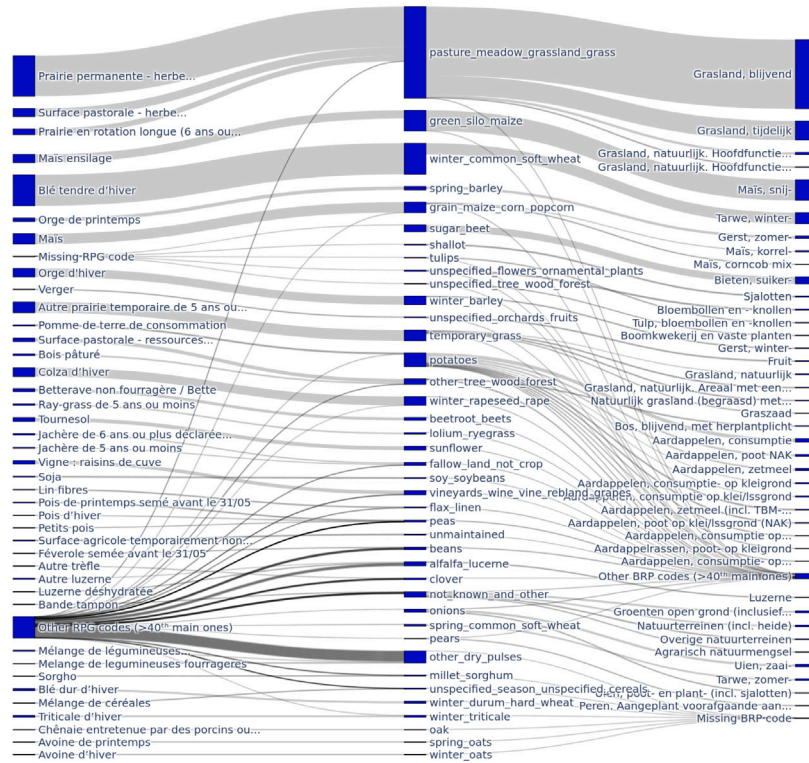


Fig. A.1. Sankey diagram of the crop harmonization, linking the French RPG (left) and the Dutch BRP (right), using HCATv2 from EuroCrops (center). The bars represent a relative share of the surface for each country. For sake of presentation (in order to fit in one page), only the 40 main crop types for each country are represented. Other crop types are grouped in “Other” classes. An interactive version of the diagram without class limitation is available on https://jeodpp.jrc.ec.europa.eu/ftp/jrc-opendata/DRLL/CropDeepTrans/data/sankey_All_crops.html.

Table A.1

Results over the Netherlands and France of the best end-of-season classification architecture on 10% of the parcels of the dataset for the test season 2020. One split is purely temporal (T) and the other one is temporal and spatial (T+S). The T split lines contain the results of our model on this specific 10% subpart of the dataset. The T+S lines contain the results of a model with the same architecture trained only over the 90% remaining. None of the models have seen timeseries from season 2020 during training. The metrics shown are macro Precision (P), Recall (R) and F1 score, as well as accuracy and micro-F1 score (m-F1).

Labels	Split	141/151-class				24/32-class				10/14-class				8/12-class			
		P	R	F1	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1	m-F1
NL	T	47.2	41.9	42.7	93.7	77.2	75.9	76.0	94.1	87.0	82.1	83.8	95.6	85.7	78.9	81.4	91.7
	T+S	47.3	42.5	42.7	93.3	75.8	74.8	74.6	93.7	86.6	81.2	82.9	95.3	85.3	77.8	80.4	91.0
FR	T	44.4	38.7	39.4	85.4	72.0	68.6	69.1	85.6	79.5	75.9	77.4	89.1	77.8	73.2	75.1	83.5
	T+S	44.8	38.4	38.4	85.1	71.6	68.0	68.8	85.4	79.1	74.4	76.3	88.8	77.3	71.5	73.9	82.7

Table A.2

Results over the Netherlands of several end-of-season classification architectures, using different features sets. The first one is composed of B4, B8A, FAPAR and LAI, the second one is an extension composed of first one with B2, B3, B8, B11 and B12, the third one only contains the 10 m bands of S2. The metrics shown are macro Precision (P), Recall (R) and F1 score, as well as accuracy and micro-F1 score (m-F1).

Labels	Features	141-class				24-class				10-class				8-class			
		P	R	F1	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1	m-F1
IntraYE _{RS}	4	27.4	20.9	20.4	89.8	64.0	60.9	60.4	90.3	78.8	75.9	74.5	92.9	76.1	72.6	70.8	87.8
IntraYE _{MM}		55.6	39.7	43.2	92.8	76.6	69.8	72.1	93.1	83.0	80.5	80.9	94.7	80.2	77.9	78.0	90.0
InterYE _{MM}		41.1	33.0	33.6	92.2	70.8	70.5	69.9	92.6	82.2	79.7	80.4	94.5	80.2	76.3	77.5	89.5
HierE _{MM}		47.3	38.7	39.7	93.3	74.7	75.5	74.7	93.7	85.2	81.9	83.1	95.2	83.6	78.8	80.6	91.1
HierE _{final}		47.1	39.3	40.2	93.6	76.6	75.8	75.7	94.0	86.7	81.9	83.6	95.5	85.3	78.7	81.1	91.6
IntraYE _{RS}	9	36.0	27.4	27.4	92.5	73.3	68.9	69.6	92.9	86.5	82.3	82.2	95.2	84.7	79.9	79.6	91.7
IntraYE _{MM}		61.0	45.6	49.0	94.3	81.0	75.9	77.7	94.6	88.0	85.9	86.3	96.1	85.9	84.2	84.2	92.7
InterYE _{MM}		47.1	35.8	37.9	94.0	77.2	74.6	75.5	94.4	88.6	84.5	86.0	96.1	87.3	82.2	84.0	92.8
HierE _{MM}		48.6	40.1	41.0	95.0	80.5	80.3	80.1	95.4	90.8	87.7	88.9	96.8	89.7	85.8	87.3	94.2
HierE _{final}		52.5	46.3	46.7	95.3	81.5	81.2	81.1	95.6	90.6	87.9	89.0	96.9	8.96	85.7	87.3	94.2
IntraYE _{RS}	S2-10m	26.7	17.4	18.4	88.6	64.1	53.3	54.2	88.9	81.5	72.1	73.2	91.9	80.0	68.6	69.8	86.3
IntraYE _{MM}		55.2	41.8	44.6	92.8	76.2	70.5	72.5	93.2	85.2	81.9	82.8	94.9	83.0	79.5	80.4	90.4
InterYE _{MM}		40.5	30.6	31.8	90.8	65.9	63.6	63.8	91.0	80.6	78.7	79.1	93.2	78.1	75.4	76.0	86.2
HierE _{MM}		48.5	39.3	40.5	93.5	75.8	75.0	75.0	93.8	88.0	83.7	85.4	95.5	86.9	81.0	83.4	91.8
HierE _{final}		49.6	41.1	42.1	93.5	76.0	75.8	75.4	93.9	88.8	83.7	85.8	95.5	88.3	80.7	84.0	91.9

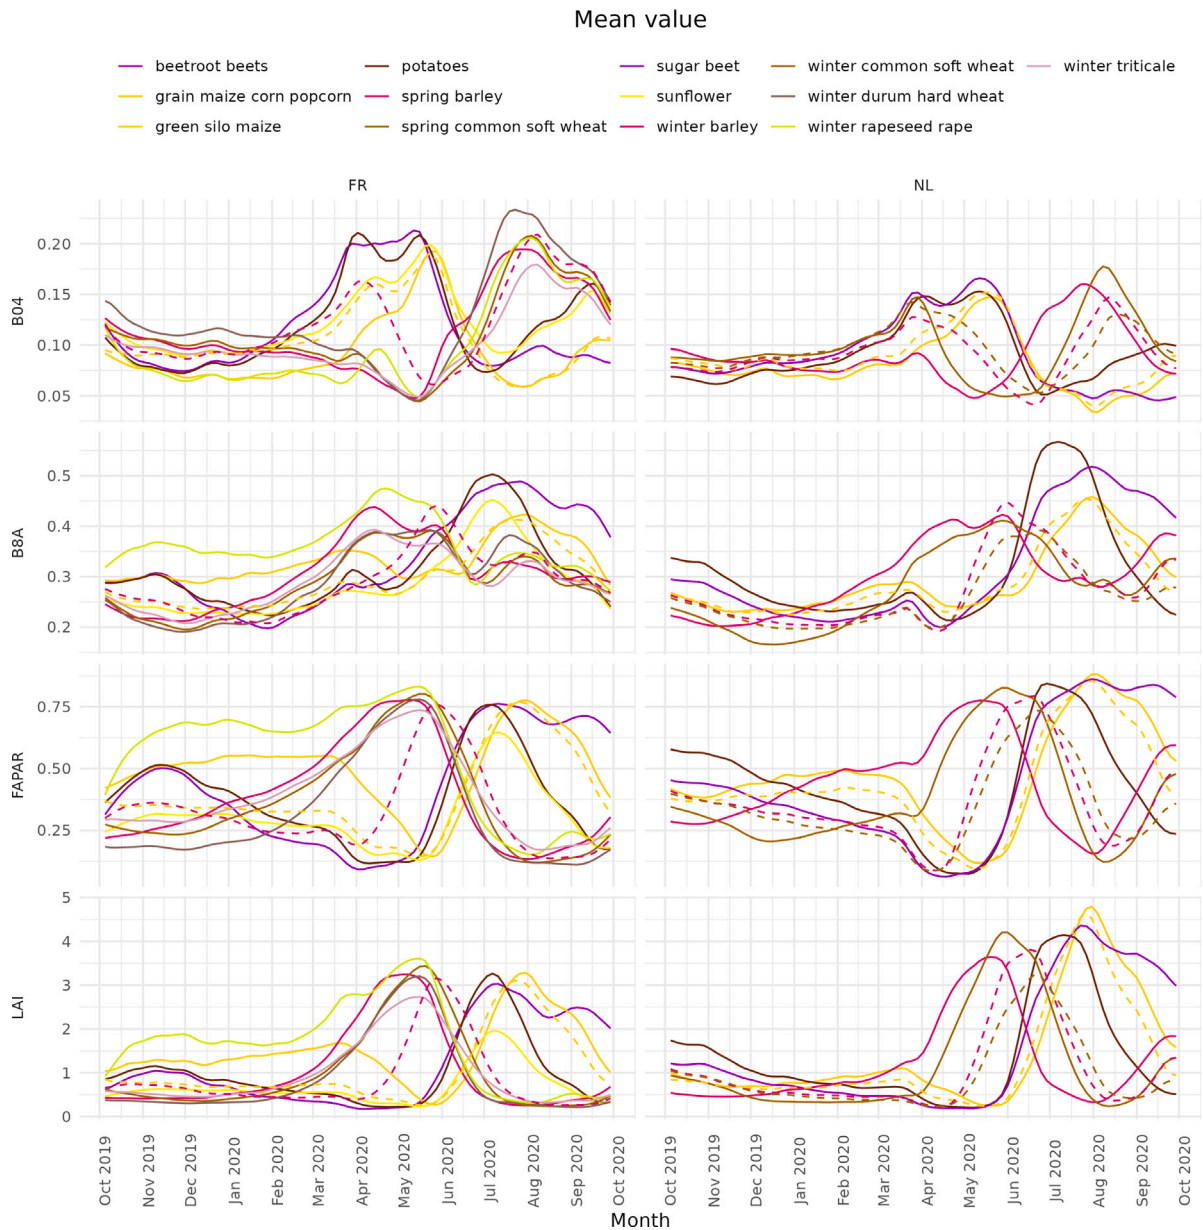


Fig. A.2. Mean of the EO-derived variable for each country 2019–2020. When crop have winter and spring varieties, the spring varieties are represented as dashed lines.

Table A.3

Computation for each of the models: time in seconds and GPU RAM memory (Mem.) in Gigabytes.

Models	CR	RS	CD	Modelisation-level		Hierar.	Training time (s)			Mem.
				Within season	Between seasons		Batch/s	NL	FR	
IntraYE _{RS}	✗	✓	✗	✓	✗	✗	49	3575	12 300	1.03
IntraYE _{MM}	✓	✓	✗	✓	✗	✗	47	3725	12 810	1.03
InterYE _{Crop}	✓	✗	✗	✗	✓	✗	116	500	1720	1.00
InterYE _{RS}	✗	✓	✗	✗	✓	✗	162	350	1200	1.00
InterYE _{MM}	✓	✓	✗	✗	✓	✗	93	625	2150	1.00
HierE _{RS}	✗	✓	✗	✓	✓	✓	11	5300	18 230	1.59
HierE _{MM}	✓	✓	✗	✓	✓	✓	11	5300	18 230	1.59
HierE _{final}	✓	✓	✓	✓	✓	✓	11	5300	18 230	1.60

Appendix D

The computational power in terms of training time per batch and per dataset, and GPU memory used during the training of the different models is shown in Table A.3. We used a batch size of 256. We did

not use any stopping strategy, as we chose the model giving the best performance on the validation set. We used a maximum number of 25 epochs for NL and 10 epochs for FR. We used float16 as precision for the weights and neurons of the neural network.




Fig. A.3. Standard deviation of the EO-derived variable for each country for the growing season 2019–2020. When crop have winter and spring varieties, the spring varieties are represented as dashed lines.

References

Arevalo, J., Solorio, T., Montes-Y-Gómez, M., González, F.A., 2017. Gated multimodal units for information fusion. In: 5th International Conference on Learning Representations, ICLR 2017 - Workshop Track Proceedings. [arXiv:1702.01992](https://arxiv.org/abs/1702.01992).

Bahdanau, D., Chorowski, J., Serdyuk, D., Brakel, P., Bengio, Y., 2016. End-to-end attention-based large vocabulary speech recognition. In: ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, vol. 2016-May, pp. 4945–4949. [http://dx.doi.org/10.1109/ICASSP.2016.7472618](https://doi.org/10.1109/ICASSP.2016.7472618), [arXiv:1508.04395](https://arxiv.org/abs/1508.04395).

Barriere, V., 2017. Hybrid Models for Opinion Analysis in Speech Interactions. In: ICML, pp. 647–651.

Barriere, V., Clavel, C., Essid, S., 2017. Opinion Dynamics Modeling for Movie Review Transcripts Classification with Hidden Conditional Random Fields. In: INTERSPEECH.

Barriere, V., Clavel, C., Essid, S., 2018. Attitude Classification in Adjacency Pairs of a Human-Agent Interaction with Hidden Conditional Random Fields. In: ICASSP.

Barriere, V., Claverie, M., 2022. Multimodal Crop Type Classification Fusing Multi-Spectral Satellite Time Series with Farmers Crop Rotations and Local Crop Distribution. In: Proceedings of 2nd Workshop on Complex Data Challenges in Earth Observation, IJCAI, vol. 3207, pp. 50–57, [arXiv:2208.10838](https://arxiv.org/abs/2208.10838).

Benamara, F., Taboada, M., Mathieu, Y., 2016. Evaluative Language Beyond Bags of Words: Linguistic Insights and Computational Applications. *Comput. Linguist.* [http://dx.doi.org/10.1162/COLI_a_00278](https://doi.org/10.1162/COLI_a_00278), URL http://www.mitpressjournals.org/doi/10.1162/COLI_a_00278. [arXiv:1309.4408](https://arxiv.org/abs/1309.4408).

Bohan, D.A., Schmucki, R., Abay, A.T., Termansen, M., Bane, M., Charalabidis, A., Cong, R.-G., Derocles, S.A., Dorner, Z., Forster, M., Gibert, C., Harrower, C., Oudoire, G., Therond, O., Young, J., Zalai, M., Pocock, M.J., 2021. Designing farmer-acceptable rotations that assure ecosystem service provision in the face of climate change. In: Bohan, D.A., Dumbrell, A.J., Vanbergen, A.J. (Eds.), *The Future of Agricultural Landscapes, Part III*. In: *Advances in Ecological Research*, vol. 65, Academic Press, pp. 169–244. [http://dx.doi.org/10.1016/bs.aecr.2021.01.002](https://doi.org/10.1016/bs.aecr.2021.01.002), (Chapter Five). URL <https://www.sciencedirect.com/science/article/pii/S0065250421000027>.

Capliez, E., Ienco, D., Gaetano, R., Baghdadi, N., Salah, A.H., 2023. Temporal-Domain Adaptation for Satellite Image Time-Series Land-Cover Mapping with Adversarial Learning and Spatially Aware Self-Training. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 16, 3645–3675. [http://dx.doi.org/10.1109/JSTARS.2023.3263755](https://doi.org/10.1109/JSTARS.2023.3263755).

Chen, M., Wang, S., Liang, P.P., Baltrušaitis, T., Zadeh, A., Morency, L.-P., 2017. Multimodal sentiment analysis with word-level fusion and reinforcement learning. In: *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pp. 163–171.

- Chung, J., Gulcehre, C., Cho, K., Bengio, Y., 2015. Gated feedback recurrent neural networks. Proceedings of the 32nd International Conference on Machine Learning, {ICML} 2015, Proceedings of the 32nd International Conference on Machine Learning, {ICML} 2015, vol. 37, 2067–2075. <http://dx.doi.org/10.1145/2661829.2661935>. arXiv:1502.02367.
- Claverie, M., Ju, J., Masek, J.G., Dungan, J.L., Vermote, E.F., Roger, J.-C., Skakun, S.V., Justice, C., 2018. The Harmonized Landsat and Sentinel-2 surface reflectance data set. *Remote Sens. Environ.* 219, 145–161.
- Claverie, M., Vermote, E.F., Weiss, M., Baret, F., Hagolle, O., Demarez, V., 2013. Validation of coarse spatial resolution LAI and FAPAR time series over cropland in southwest France. *Remote Sens. Environ.* 139, 216–230.
- d'Andrimont, R., Claverie, M., Kempeneers, P., Muraro, D., Yordanov, M., Peressutti, D., Batič, M., Waldner, F., 2023. AI4Boundaries: an open AI-ready dataset to map field boundaries with Sentinel-2 and aerial photography. *Earth Syst. Sci. Data* 15 (1), 317–329.
- d'Andrimont, R., Taymans, M., Lemoine, G., Ceglar, A., Yordanov, M., van der Velde, M., 2020. Detecting flowering phenology in oil seed rape parcels with Sentinel-1 and-2 time series. *Remote Sens. Environ.* 239, 111660.
- d'Andrimont, R., Verhegghen, A., Lemoine, G., Kempeneers, P., Meroni, M., Van der Velde, M., 2021. From parcel to continental scale—A first European crop type map based on Sentinel-1 and LUCAS Copernicus in-situ observations. *Remote Sens. Environ.* 266, 112708.
- Delloy, C., Weiss, M., Defourny, P., 2018. Retrieval of the canopy chlorophyll content from Sentinel-2 spectral bands to estimate nitrogen uptake in intensive winter wheat cropping systems. *Remote Sens. Environ.* 216, 245–261.
- Dogliotti, S., Rossing, W., Van Ittersum, M., 2003. ROTAT, a tool for systematically generating crop rotations. *Eur. J. Agron.* 19 (2), 239–250.
- Du, Y., Zhen, X., Shao, L., Snoek, C.G., 2020. Metanorm: Learning to normalize few-shot batches across domains. In: International Conference on Learning Representations.
- Eilers, P.H., Pesendorfer, V., Bonifacio, R., 2017. Automatic smoothing of remote sensing data. In: 2017 9th International Workshop on the Analysis of Multitemporal Remote Sensing Images. MultiTemp, IEEE, pp. 1–3.
- Finn, C., Abbeel, P., Levine, S., 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In: 34th International Conference on Machine Learning. ICML 2017, vol. 3, pp. 1856–1868. arXiv:1703.03400.
- Giordano, S., Bailly, S., Landrieu, L., Chehata, N., 2020. Improved crop classification with rotation knowledge using sentinel-1 and -2 time series. *Photogramm. Eng. Remote Sens.* 86 (7), 431–441. <http://dx.doi.org/10.14358/PERS.86.7.431>.
- Harris, Z.S., 1954. Distributional Structure. *Word* 10 (2–3), 146–162. http://dx.doi.org/10.1007/978-94-009-8467-7_1, URL <http://psycnet.apa.org/psycinfo/1956-02807-001>.
- Hochreiter, S., Schmidhuber, J., 1997. LONG SHORT-TERM MEMORY. *Neural Comput.* 9 (8), 1735–1780. <http://dx.doi.org/10.1162/neco.1997.9.8.1735>, URL <http://www7.informatik.tu-muenchen.de/~hochreit/%5Cnhttp://www.idsia.ch/~juergen>. arXiv:1206.2944.
- Jia, X., Nayak, G., Khandelwal, A., Karpatne, A., Kumar, V., 2019. Classifying heterogeneous sequential data by cyclic domain adaptation: An application in land cover detection. In: Proceedings of the 2019 SIAM International Conference on Data Mining. SDM, pp. 540–548. <http://dx.doi.org/10.1137/1.9781611975673.61>, arXiv:https://epubs.siam.org/doi/pdf/10.1137/1.9781611975673.61. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611975673.61>.
- Johnson, D.M., Mueller, R., 2021. Pre-and within-season crop type classification trained with archival land cover information. *Remote Sens. Environ.* 264, 112576.
- Kingma, D., Ba, J., 2014. Adam: A Method for Stochastic Optimization. In: International Conference on Learning Representations. pp. 1–13, URL <http://arxiv.org/abs/1412.6980>. arXiv:1412.6980.
- Lin, C., Zhong, L., Song, X.-P., Dong, J., Lobell, D.B., Jin, Z., 2022. Early-and in-season crop type mapping without current-year ground truth: Generating labels from historical information via a topology-based approach. *Remote Sens. Environ.* 274, 112994.
- Lucas, B., Pelletier, C., Schmidt, D., Webb, G.I., Petitjean, F., 2023. A Bayesian-inspired, deep learning-based, semi-supervised domain adaptation technique for land cover mapping. *Mach. Learn.* 112 (6), 1941–1973. <http://dx.doi.org/10.1007/s10994-020-05942-z>.
- Merlos, F.A., Hijmans, R.J., 2020. The scale dependency of spatial crop species diversity and its relation to temporal diversity. *Proc. Natl. Acad. Sci.* 117 (42), 26176–26182.
- Meroni, M., d'Andrimont, R., Vrieling, A., Fasbender, D., Lemoine, G., Rembold, F., Seguin, L., Verhegghen, A., 2021. Comparing land surface phenology of major European crops as derived from SAR and multispectral data of Sentinel-1 and-2. *Remote Sens. Environ.* 253, 112322.
- Mikolov, T., Karafiát, M., Burget, L., Černocký, J., Khudanpur, S., 2010. Recurrent neural network based language model. In: Eleventh Annual Conference of the International Speech Communication Association.
- Nyborg, J., Pelletier, C., Assent, I., 2022a. Generalized Classification of Satellite Image Time Series with Thermal Positional Encoding. IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, vol. 2022-June, 1391–1401. <http://dx.doi.org/10.1109/CVPRW56347.2022.00145>. arXiv:2203.09175.
- Nyborg, J., Pelletier, C., Lefèvre, S., Assenta, I., 2022b. TimeMatch: Unsupervised Cross-Region Adaptation by Temporal Shift Estimation. *ISPRS J. Photogramm. Remote Sens.* arXiv:2111.02682v3.
- Osman, J., Inglada, J., Dejoux, J.F., 2015. Assessment of a Markov logic model of crop rotations for early crop mapping. *Comput. Electron. Agric.* 113, 234–243. <http://dx.doi.org/10.1016/j.compag.2015.02.015>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32 (NeurIPS), arXiv:1912.01703.
- Pelletier, C., Webb, G.I., Petitjean, F., 2019. Temporal convolutional neural network for the classification of satellite image time series. *Remote Sens.* 11 (5), 1–22. <http://dx.doi.org/10.3390/rs11050523>, arXiv:1811.10166.
- Peng, K.C., Wu, Z., Ernst, J., 2018. Zero-Shot Deep Domain Adaptation. *ECCV* 793–810. http://dx.doi.org/10.1007/978-3-030-01252-6_47, arXiv:1707.01922.
- Porter, J., Xie, L., Challinor, A., Cochrane, K., Howden, S., Iqbal, M., Lobell, D., Isabel Travasso, M., Chhetri, N., Garrett, K., et al., 2014. Food security and food production systems.
- Poth, C., Pfeiffer, J., Andreas, R., Kamath, A., Vuli, I., Ruder, S., Cho, K., Gurevych, I., 2020. AdapterHub : A framework for adapting transformers. arXiv preprint arXiv:2007.07779.
- Quinton, F., Landrieu, L., 2021. Crop rotation modeling for deep learning-based parcel classification from satellite time series. *Remote Sens.* 13 (22), <http://dx.doi.org/10.3390/rs13224599>, arXiv:2110.08187.
- Rußwurm, M., Courty, N., Emonet, R., Lefèvre, S., Tuia, D., Tavenard, R., 2023. End-to-end learned early classification of time series for in-season crop type mapping. *ISPRS J. Photogramm. Remote Sens.* 196 (October 2022), 445–456. <http://dx.doi.org/10.1016/j.isprsjprs.2022.12.016>.
- Rußwurm, M., Körner, M., 2018. Multi-temporal land cover classification with sequential recurrent encoders. *ISPRS Int. J. Geo-Inf.* 7 (4), <http://dx.doi.org/10.3390/ijgi7040129>, arXiv:1802.02800.
- Rußwurm, M., Körner, M., 2020. Self-attention for raw optical Satellite Time Series Classification. *ISPRS J. Photogramm. Remote Sens.* 169 (November), 421–435. <http://dx.doi.org/10.1016/j.isprsjprs.2020.06.006>, arXiv:1910.10536.
- Rußwurm, M., Lefèvre, S., Körner, M., 2019a. BreizhCrops: A satellite time series dataset for crop type identification. *CoRR*, abs/1905.11893. URL <http://arxiv.org/abs/1905.11893>. arXiv:1905.11893.
- Rußwurm, M., Tavenard, R., Lefèvre, S., Körner, M., 2019b. Early Classification for Agricultural Monitoring from Satellite Time Series. In: ICML Ai4socialgood Workshop. URL <http://arxiv.org/abs/1908.10283>. arXiv:1908.10283.
- Rußwurm, M., Wang, S., Körner, M., Lobell, D., 2019c. Meta-Learning for Few-Shot Land Cover Classification. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops.
- Sainte Fare Garnot, V., Landrieu, L., 2021. Leveraging Class Hierarchies with Metric-Guided Prototype Learning. In: Proceedings of the British Machine Vision Conference 2021. BMVC 2021.
- Sainte Fare Garnot, V., Landrieu, L., Giordano, S., Chehata, N., 2019. Time-Space Tradeoff in Deep Learning Models for Crop Classification on Satellite Multi-Spectral Image Time Series. In: IGARSS, pp. 6247–6250. <http://dx.doi.org/10.1109/igarss.2019.8900517>, arXiv:1901.10503.
- Sainte Fare Garnot, V., Landrieu, L., Giordano, S., Chehata, N., 2020. Satellite image time series classification with pixel-set encoders and temporal self-attention. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 12322–12331. <http://dx.doi.org/10.1109/CVPR42600.2020.01234>, arXiv:1911.07757.
- Sanh, V., Wolf, T., Ruder, S., Court, H., Row, H., 2018. A Hierarchical Multi-task Approach for Learning Embeddings from Semantic Tasks. In: AAAI. arXiv:arXiv:1811.06031v1.
- Schneider, M., Broszeit, A., Körner, M., 2021. Eurocrops: A pan-european dataset for time series crop type classification. arXiv preprint arXiv:2106.08151.
- Schneider, M., Schelte, T., Schmitz, F., Körner, M., 2023. EuroCrops: The largest harmonized open crop dataset across the European union. *Sci. Data* 10 (1), 612. <http://dx.doi.org/10.1038/s41597-023-02517-0>.
- Schuller, B., Steidl, S., Batliner, A., Hirschberg, J., Burgoon, J.K., Baird, A., Elkins, A., Zhang, Y., Coutinho, E., Evani, K., 2016. The INTERSPEECH 2016 Computational Paralinguistics Challenge: Deception, Sincerity & Native Language. In: Proceedings of the Annual Conference of the International Speech Communication Association. INTERSPEECH.
- Selea, T., 2023. AgriSen-COG, a multicountry, multitemporal large-scale Sentinel-2 benchmark dataset for crop mapping using deep learning. *Remote Sens.* 15 (12), URL <https://www.mdpi.com/2072-4292/15/12/2980>.
- Serban, I.V., Sordoni, A., Bengio, Y., Courville, A., Pineau, J., 2015. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. <http://dx.doi.org/10.1017/CBO9781107415324.004>, URL <http://arxiv.org/abs/1507.04808>. arXiv:1507.04808.
- Soille, P., Burger, A., De Marchi, D., Kempeneers, P., Rodriguez, D., Syrris, V., Vasilev, V., 2018. A versatile data-intensive computing platform for information retrieval from big geospatial data. *Future Gener. Comput. Syst.* 81, 30–40.

- Trigeorgis, G., Ringeval, F., Brueckner, R., Marchi, E., Nicolaou, M.A., Schuller, B., Zafeiriou, S., 2016. Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network. In: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, pp. 5200–5204. <http://dx.doi.org/10.1109/ICASSP.2016.7472669>, URL <http://ieeexplore.ieee.org/document/7472669/>.
- Tseng, G., Kerner, H., Nakalembe, C., Becker-Reshef, I., 2021a. Learning to predict crop type from heterogeneous sparse labels using meta-learning. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops. pp. 1111–1120. <http://dx.doi.org/10.1109/CVPRW53098.2021.00122>.
- Tseng, G., Zvonkov, I., Nakalembe, C., Kerner, H., 2021b. CropHarvest: a global satellite dataset for crop type classification. In: Neurips. (NeurIPS), pp. 1–14, URL <https://github.com/nasaharvest/cropharvest>.
- Turkoglu, M.O., D'Aronco, S., Perich, G., Liebisch, F., Streit, C., Schindler, K., Wegner, J.D., 2021. Crop mapping from image time series: Deep learning with multi-scale label hierarchies. *Remote Sens. Environ.* 264, 1–37. <http://dx.doi.org/10.1016/j.rse.2021.112603>, arXiv:2102.08820.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention Is All You Need. arXiv:1706.03762 [cs].
- Wang, Z., Zhang, H., He, W., Zhang, L., 2021. Phenology Alignment Network : A Novel Framework for Cross-Regional Time Series Crop Classification. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. CVPRW, IEEE, pp. 2934–2943. <http://dx.doi.org/10.1109/CVPRW53098.2021.00329>.
- Weilandt, F., Behling, R., Goncalves, R., Madadi, A., Richter, L., Sanona, T., Spengler, D., Welsch, J., 2023. Early crop classification via multi-modal satellite data fusion and temporal attention. *Remote Sens.* 15 (3), <http://dx.doi.org/10.3390/rs15030799>, URL <https://www.mdpi.com/2072-4292/15/3/799>.
- Weiss, M., Baret, F., 1999. Evaluation of canopy biophysical variable retrieval performances from the accumulation of large swath satellite data. *Remote Sens. Environ.* 70 (3), 293–306.
- Weiss, M., Jacob, F., Duveiller, G., 2020. Remote sensing for agricultural applications: A meta-review. *Remote Sens. Environ.* 236, 111402.
- Wöllmer, M., Kaiser, M., Eyben, F., Schuller, B., Rigoll, G., 2013a. LSTM-modeling of continuous emotions in an audiovisual affect recognition framework. *Image Vis. Comput.* 31 (2), 153–163. <http://dx.doi.org/10.1016/j.imavis.2012.03.001>.
- Wöllmer, M., Weninger, F., Knaup, T., Schuller, B., Sun, C., Sagae, K., Morency, L.-P., 2013b. YouTube Movie Reviews: Sentiment Analysis in an Audio-Visual Context. *IEEE Intell. Syst.* 28 (3), 46–53. <http://dx.doi.org/10.1109/MIS.2013.34>, URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6487473>.
- Xiao, Y., Mignolet, C., Mari, J.-F., Benoit, M., 2014. Modeling the spatial distribution of crop sequences at a large regional scale using land-cover survey data: A case from France. *Comput. Electron. Agric.* 102, 51–63.
- Xu, C., Tao, D., Xu, C., 2013. A Survey on Multi-view Learning. pp. 1–59. <http://dx.doi.org/10.1145/1553374.1553391>, URL <http://arxiv.org/abs/1304.5634>. arXiv:1304.5634.
- Yaramasu, R., Bandaru, V., Pnvr, K., 2020. Pre-season crop type mapping using deep neural networks. *Comput. Electron. Agric.* 176, 105664. <http://dx.doi.org/10.1016/j.compag.2020.105664>, URL <https://www.sciencedirect.com/science/article/pii/S0168169920307742>.
- Zhong, L., Hu, L., Zhou, H., 2019. Deep learning based multi-temporal crop classification. *Remote Sens. Environ.* 221, 430–443. <http://dx.doi.org/10.1016/j.rse.2018.11.032>, URL <https://www.sciencedirect.com/science/article/pii/S0034425718305418>.